

1 **Ancient biases in phenotype production drove the functional evolution of a protein family**

2 Santiago Herrera-Álvarez¹†, Jaeda E. J. Patton²†, Joseph W. Thornton^{1,3}*

3 ¹Department of Ecology and Evolution; ²Committee on Genetics, Genomics, and Systems
4 Biology; ³Department of Human Genetics, University of Chicago; Chicago, IL, USA.

5 †These authors contributed equally to this work; the order of their names was determined by
6 coin flip.

7 *Corresponding author. Email: joet1@uchicago.edu

8 ABSTRACT

9 Biological systems may be biased in the phenotypes they can access by mutation¹⁻⁷, but the
10 extent of these biases and their causal role in the evolution of extant phenotypic diversity
11 remains unclear. There are three major challenges: it is difficult to isolate the effect of bias in the
12 genotype-phenotype (GP) map from that of natural selection in producing natural diversity^{6,8-11},
13 the universe of possible genotypes and phenotypes is so vast and complex that a direct
14 characterization has been impossible, and most extant phenotypes evolved long ago in species
15 whose GP maps cannot be recovered. Here we develop exhaustive multi-phenotype deep
16 mutational scanning to experimentally characterize the complete GP maps of two reconstructed
17 ancestral steroid receptor proteins, which existed during an ancient phylogenetic interval when a
18 new phenotype—specific binding of a new DNA response element—evolved¹². We measured all
19 possible DNA specificity phenotypes encoded by all possible amino acid combinations at sites in
20 the protein's DNA binding interface. We found that the ancestral GP maps are structured by
21 strong global bias—unequal propensity to encode the various phenotypes—and extreme
22 heterogeneity in the phenotypes accessible around each genotype, which strongly affect
23 evolution on both long and short timescales. Distinct biases in the two ancestral maps steered
24 evolution toward the lineage-specific functional phenotypes that evolved during history. Our
25 findings establish that ancient biases in the GP relationship were causal factors in the
26 evolutionary process that produced the present-day patterns of phenotypic conservation and
27 diversity in this protein family.

28 MAIN TEXT

29 Countless conceivable lifeforms have evolved rarely or never, and those that exist are mostly
30 restricted to specific lineages^{13–16}. No flying vertebrates have two pairs of wings, for example,
31 and no turtles or frogs fly. What explains the biased distribution of phenotypes in nature?
32 Classical explanations focus on the influence of selection^{17,18}, but it is possible that the
33 propensities of biological systems to produce phenotypic variation could also shape evolutionary
34 outcomes. A phenotype can become fixed in an evolving population only if it is first generated
35 by mutation. If biological systems are more likely to produce some phenotypes than others^{1–7},
36 and if these propensities change over time as lineages diverge^{19,20}, then some phenotypes will be
37 more likely to evolve in some taxa than in others.

38 Whether phenotype production has been an important cause of evolutionary outcomes is
39 unclear, because most patterns of phenotypic variation observed in nature could arise from
40 production biases, natural selection, or both, and disentangling their past influences is extremely
41 challenging^{6,8–11}. Ideally, we would isolate the phenotype production process by directly
42 characterizing the complete genotype-phenotype (GP) map, which maps all possible
43 combinations of mutations to the phenotypes they encode. This would allow us to precisely
44 quantify the ability of a system to produce phenotypic variation, both on a global scale and by
45 mutation from each particular genotype. The total space of genotypes and phenotypes is vast, but
46 we reasoned that by combining three recent technical advances, this goal could become tractable
47 for proteins and their biochemical phenotypes. The first technique is deep mutational scanning
48 (DMS), which allows huge libraries of protein variants to be characterized experimentally²¹. The
49 scope of genetic variation to be measured in a DMS study can be defined as all combinations of
50 all 20 possible amino acid states at the sequence sites that determine the protein's phenotype of
51 interest, thus encompassing all potential genetic variation at those sites^{22–26}. Complete
52 combinatorial DMS studies to date, however, have assayed only one or a few phenotypes that
53 exist in extant proteins. Although this approach can reveal all genetic variants that encode these
54 phenotypes, it cannot address why those particular phenotypes evolved in the first place; to
55 understand why evolution turned out as it did, we must characterize the propensity of mutations
56 to produce not only the phenotypes that evolved historically but also all the phenotypes that did
57 not. The second technique—comprehensive multi-phenotype profiling—addresses this limitation
58 by quantifying all possible phenotypes that a single protein can perform, such as binding of all
59 possible substrates or DNA elements in a defined class^{27–31}. We reasoned that by combining
60 DMS with comprehensive multi-phenotype assays, we could map all possible phenotypes onto
61 all possible genotypes within a defined scope. This would allow us to characterize the total
62 capacity of a protein system to produce and access phenotypic variation by genetic change.

63 The phenotypes of extant lineages evolved long ago, so understanding the causal role of
64 phenotype production in historical evolution requires GP maps to be characterized as they
65 existed in the deep past. The third technique—ancestral protein reconstruction³²—can address
66 this problem by providing the protein backgrounds on which a comprehensive combinatorial
67 multi-phenotype DMS study is performed. Moreover, characterizing such GP maps across a
68 phylogenetic time series of reconstructed ancestral proteins^{24,33} would reveal how biases in
69 phenotype production may have changed over time and whether these biases are congruent with
70 the trajectories of phenotypic evolution that actually unfolded during history.

71 Here we apply this approach to assess how phenotype production shaped the functional
72 diversification of the steroid hormone receptor protein family. We use comprehensive multi-
73 phenotype DMS to experimentally characterize GP maps of the binding interface of two
74 reconstructed ancestral steroid hormone receptor DNA binding domains (SR DBDs) and their
75 ability to encode specific recognition of all possible DNA response elements. We then analyze

76 these maps to understand 1) how they could shape potential phenotypic outcomes of evolution on
77 short and long timescales, 2) characterize the mechanisms that changed key features of the maps
78 across evolutionary time, and 3) assess the impact of the maps on the historical evolutionary
79 processes that yielded the lineage-specific patterns of DNA specificity in extant steroid hormone
80 receptors.

81 **Two complete ancestral GP maps**

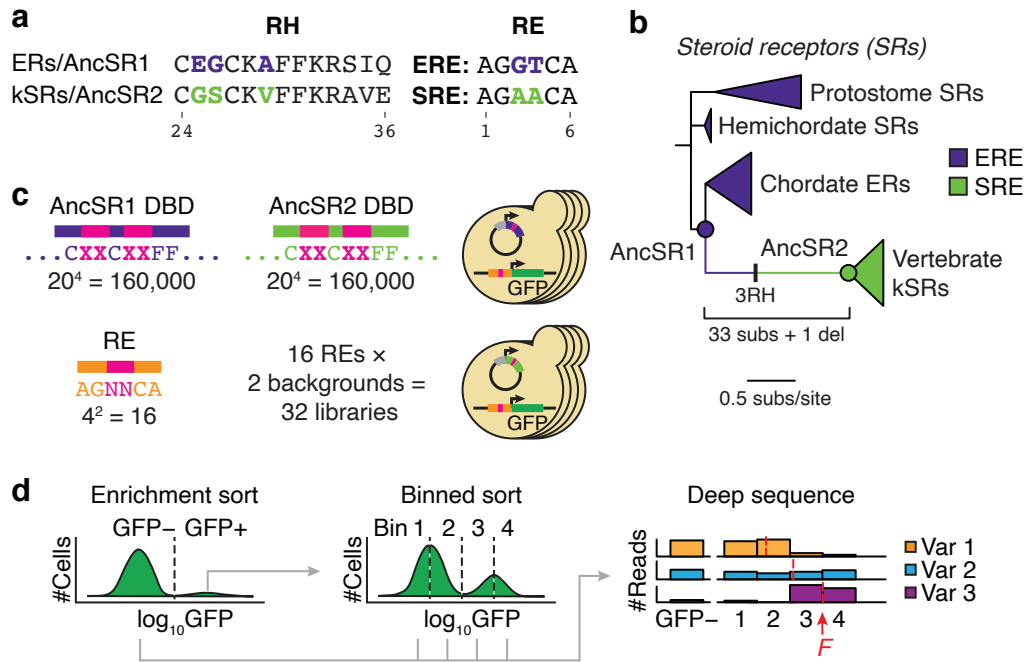
82 SRs are a family of transcription factors that regulate physiological and reproductive biology in
83 bilaterian animals. Most bilaterian taxa have a single SR, which specifically binds to inverted
84 palindromes of the motif AGGTCA, called the estrogen response element (ERE; Fig. 1a). In
85 chordates, a gene duplication of the ancestral SR (AncSR1) produced two major SR classes,
86 which have different DNA specificity phenotypes: chordate estrogen receptors (ERs) retain the
87 ancestral ERE specificity, but a novel specificity for a palindrome of AGAACA, called the
88 steroid response element (SRE), evolved in the lineage leading to AncSR2, the common ancestor
89 of the chordate ketosteroid receptors (kSRs; Fig. 1a, b)¹². Specificity for DNA is determined
90 primarily by the amino acid sequence of a recognition helix (RH) that binds in the DNA major
91 groove^{34,35}. AncSR1 and AncSR2 DBDs differ by 34 amino acid replacements, but experiments
92 on the reconstructed proteins established that three amino acid changes in the RH were the
93 primary cause of the evolution of SRE specificity¹².

94 To understand how phenotype production may have shaped the evolution of SR-DBD
95 specificity, we characterized combinatorially complete GP maps of the DBD-response element
96 (RE) interface at the key ancestral timepoints AncSR1 and AncSR2. The scope of genotypes is
97 all possible $20^4 = 160,000$ amino acid variants at four variable sites in the recognition helix—the
98 three that changed between AncSR1 and AncSR2, plus one other that varies in the broader
99 nuclear receptor family (Fig. 1c). The scope of specificity phenotypes consists of all $4^2 = 16$
100 possible RE sequences that can be produced by all combinations of nucleotides at the two base
101 positions that vary between ERE and SRE. These two maps of the recognition helix-RE interface
102 can be thought of as submaps within the much larger GP map of the entire DBD, which are
103 connected by the 31 other “background” substitutions that occurred between the AncSR1 and
104 AncSR2 proteins (Fig. 1b).

105 We engineered two protein libraries, each containing all 160,000 variants of the
106 recognition helix in the background of either the AncSR1 or AncSR2 DBD, along with 16 yeast
107 strains, each containing a GFP reporter driven by one of the REs (Fig. 1c, Extended Data Fig.
108 1a–e). We transformed each RE strain separately with the two protein libraries, with barcodes to
109 mark the strain and the ancestral background, for a total of 5.12 million protein-DNA complexes.
110 , we used an initial round of fluorescence-activated cell sorting to enrich the yeast libraries for
111 GFP-positive cells, pooled the enriched libraries, sorted cells in three replicates by their
112 fluorescence, and sequenced the sorted bins (Fig. 1d, Extended Data Fig. 1f, g). Using this
113 strategy, we obtained empirical fluorescence estimates for the majority of complexes with good
114 replicability ($r^2 = 0.92$ across replicates, excluding complexes at the lower bound of
115 fluorescence; Extended Data Fig. 2). Fluorescence of the remaining complexes was predicted
116 using a generalized linear model trained on the experimental data (Methods, Extended Data Fig.
117 3a–d)^{36,37}.

118 Each protein variant was assigned a DNA specificity phenotype based on these
119 experiments. A protein variant is classified as specific if it is functional in complex with only one
120 RE, promiscuous if it is functional on multiple REs, or nonfunctional if it is not functional on
121 any RE. We defined functional complexes as those having fluorescence at least as great as the

122 wild-type complex in each background (*i.e.* EGKA:ERE for the AncSR1 library and GSKV:SRE
 123 for AncSR2) (Methods, Extended Data Fig. 3e–g).



124

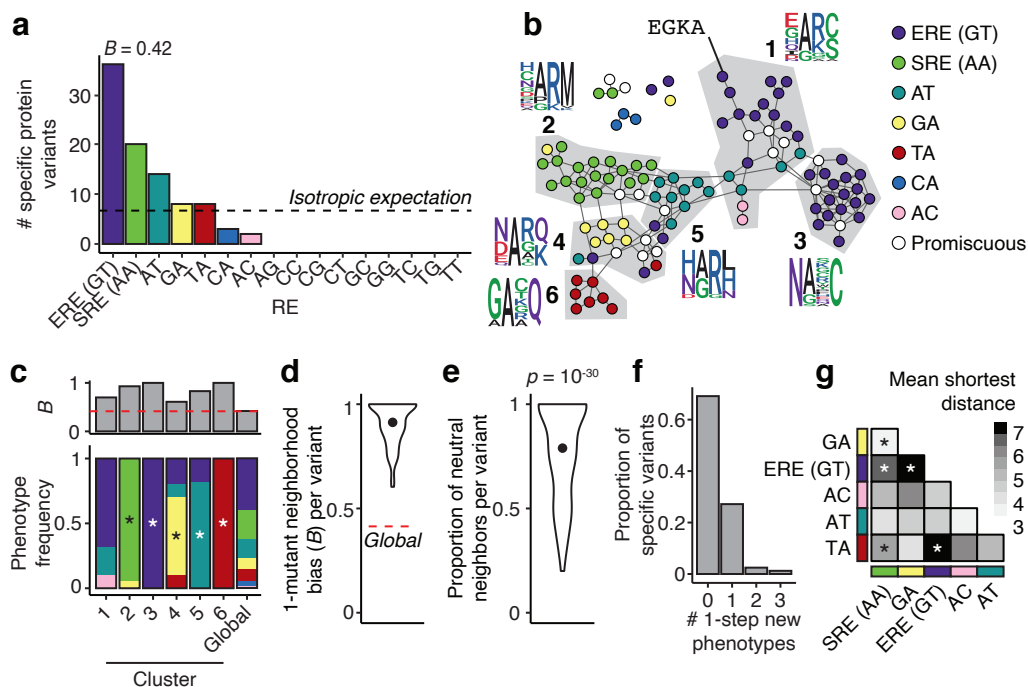
125 **Fig. 1 | Characterizing ancestral GP maps using multi-phenotype DMS.** **a**, Amino acid
 126 sequence of the recognition helix (RH) in extant and ancestral steroid receptor (SR) proteins and
 127 the sequence of the RE they bind to. Colored residues are responsible for differences in protein-
 128 RE specificity. **b**, Phylogeny of SRs. Each clade of proteins is colored by the RE sequence it
 129 recognizes. In chordates, a historical transition from ERE to SRE specificity occurred along the
 130 branch between AncSR1 (the common ancestor of all chordate SRs) and AncSR2 (the common
 131 ancestor of vertebrate kSRs). The number of historical sequence changes along the AncSR1-
 132 AncSR2 branch is shown; three of these in the recognition helix (RH) caused the specificity
 133 switch¹². **c, d**, DMS experiment to assay effects of RH genotype on binding to variable REs. **c**,
 134 We built combinatorial libraries of all combinations of 20 amino acid states at four variable sites
 135 in the RH (pink Xs), using the rest of the AncSR1 and AncSR2 DBDs as backgrounds (top left).
 136 These were transformed into 16 *S. cerevisiae* strains, each containing one of the 16 possible RE
 137 motifs (pink Ns, bottom left) genomically integrated upstream of a GFP reporter gene (right). **d**,
 138 We assayed binding of DBD-RE complexes using FACS coupled with deep sequencing. For
 139 each library, we performed an initial enrichment sort to select for GFP⁺ cells. We then grew up
 140 the selected cells, pooled them across the 32 libraries, and resorted them into four fluorescence
 141 bins in triplicate (binned sort). Sorted cells were deep sequenced to estimate the mean \log_{10} GFP
 142 (F) of each combination of protein and RE genotypes.

143

144 Global bias in the AncSR1 GP map

145 The probability that a phenotype will evolve equals the probability that it will be produced by
 146 mutation times the probability that, once produced, it will be fixed. The GP map
 147 would have no effect on evolutionary outcomes if and only if it had two properties: isotropy—

148 encoding all phenotypes with equal probability—and homogeneity—producing the same
 149 distribution of phenotypes from all starting genotypes in the map^{38–41}. If the map is anisotropic,
 150



151

152 **Fig. 2 | Global and local bias in the AncSR1 GP map.** **a**, Global production distribution in the
 153 AncSR1 GP map. Bars represent the number of protein variants that bind specifically to each
 154 RE. The dashed line shows the expected frequencies if the distribution were unbiased. B ,
 155 phenotype bias, calculated as one minus the entropy (base 16) of the distribution. **b**, Sequence
 156 space network of the AncSR1 GP map. Nodes represent functional protein variants, colored by
 157 their RE specificity; white nodes, promiscuous genotypes. Edges connect protein variants that
 158 can be interconverted by a single nucleotide change. Genotype clusters (1–6, ordered by
 159 decreasing size) identified by a community structure detection algorithm are shown in gray.
 160 Sequence logos show amino acid frequencies at the variable RH sites in each cluster. **c**, Bottom:
 161 Frequencies of specificity phenotypes within each genotype cluster; the global production
 162 distribution is shown for comparison. Asterisks, phenotypes significantly enriched within a
 163 cluster relative to the global production distribution (Fisher’s exact test, $p < 0.05$ after Bonferroni
 164 correction). Top: strength of phenotype bias (B) in each cluster. Red line, B of global production
 165 distribution. **d**, Distribution of phenotype bias (B) of the 1-mutant neighborhood of every RE-
 166 specific protein variant in the main network component. Dot shows the mean. Dashed red line,
 167 global phenotype bias. **e**, Proportion of neutral neighbors per RE-specific protein variant in the
 168 main component of the AncSR1 map. Dot shows the mean. P -value, probability that the mean
 169 would be at least as great as observed if phenotypes were randomly reassigned in the main
 170 component ($n = 91$). **f**, Distribution of the number of new phenotypes accessible within one
 171 mutation, across all RE-specific variants in the AncSR1 main component. **g**, Mean distance
 172 between pairs of phenotypes in the AncSR1 main component. The color of each cell shows the
 173 mean of the length of the most direct path from every genotype encoding one phenotype to every
 174 genotype encoding the other. Bonferroni corrected p -values for a two-sided permutation test
 175 where phenotype associations were shuffled within the main component: * $p < 0.001$.
 176

177 then phenotypes more likely to be produced would be more likely to evolve; if the map is
178 heterogeneous, then the probability that each phenotype will be produced—and hence evolve—
179 would change as lineages diverge from each other across the map.

180 We assessed the isotropy of the AncSR1 GP map by characterizing the frequency
181 distribution of DNA specificity phenotypes encoded by all functional protein variants. Only 107
182 out of 160,000 total genotypes in the library were functional (0.07%). Of these, the majority (91)
183 were specific for a single RE. We calculated the bias (B) of this global phenotype distribution,
184 defined as 1 minus the Shannon entropy (base 16); B can range from 0 when specificity for all 16
185 REs is encoded with equal frequency to 1 when only a single phenotype is encoded. We found
186 that the distribution is strongly anisotropic ($B = 0.42$). Two specificity phenotypes—ERE and
187 SRE—together account for >60% of all specific genotypes, and only five others can be produced
188 at all; nine phenotypes are not encoded by any protein variant (Fig. 2a).

189 We refer to this anisotropy as global bias in the GP map⁴⁰. Global bias in the AncSR1
190 map imposes hard limits on phenotypic evolution—the majority of conceivable phenotypes
191 could never evolve in this map, even if they conferred strong fitness advantages. The global bias
192 is also congruent with evolutionary history—the phenotypes that evolved historically in the two
193 lineages descending from AncSR1 are also the most frequently encoded.

194 **Local bias in the AncSR1 GP map**

195 We next assessed the homogeneity of the AncSR1 GP map using Maynard-Smith's classic
196 network model of sequence space⁴². Each functional protein variant is a node with its
197 experimentally defined phenotype. Nodes are connected by edges if their amino acid sequences
198 can be interconverted by a single nucleotide change given the standard genetic code.
199 Nonfunctional variants are excluded from the network, based on the assumption that they will be
200 removed quickly from evolving populations by purifying selection.

201 We found that the distribution of phenotypes in AncSR1 sequence space is strongly
202 heterogeneous. Although the majority of functional genotypes (91%) and phenotypes (6 of 7) are
203 mutually connected in a single main network component, each phenotype tends to be sequestered
204 in a local region (Fig. 2b). Using a community structure detection algorithm⁴³, we found that the
205 main network component can be partitioned into six clusters of genotypes that have dense
206 connectivity within clusters and weak connectivity between (Fig. 2b). The phenotype bias B
207 within every single cluster is higher than the global bias of the map, and 5 of 6 clusters are
208 significantly enriched for a single specificity phenotype, which differ among all 5 clusters (Fig.
209 2c). The clumpy distribution of phenotypes in sequence space arises from the simple fact that
210 similar genotypes, which are connected to each other in sequence space, are likely to encode
211 similar phenotypes (Fig. 2b, logos).

212 This heterogeneity creates local bias⁴⁰: the propensity to produce phenotypes depends
213 strongly on the particular genotype occupied at the time. The one-mutant neighborhood around
214 every genotype has extremely high bias (mean $B = 0.91$; Fig. 2d), indicating that individual
215 genotypes can access much less phenotypic variation than is encoded across genotype space as a
216 whole. Most mutations are phenotypically neutral (79% of edges; Fig. 2e), and most genotypes
217 can directly access at most one new phenotype (Fig. 2f). The historical starting genotype
218 (EGKA), for example, has access to only one functional neighbor, which also has ERE
219 specificity. Another consequence of heterogeneity is that phenotypes, aggregated over the
220 genotypes that encode them, are differentially accessible to each other, with substantial variation
221 in the number of mutations required to transform each phenotype into the others (Fig. 2g). For
222 example, SRE-specific protein genotypes are directly accessible from nodes encoding specificity

223 for AT and GA, but they are multiple substitutions away from all ERE-specific genotypes (Fig.
224 2b).

225 The GP map of AncSR1 is therefore both anisotropic and heterogeneous, and these
226 properties impose global and local biases on the production of phenotypes. Global bias favors
227 production of the historical phenotypes ERE and SRE and entirely prevents the production of
228 most conceivable phenotypes. Local bias further restricts the number of accessible phenotypes
229 from each particular genotype, favoring conservation over the evolution of new phenotypes,
230 including from the historical genotype EGKA.

231 **Biases in the GP map affect phenotypic outcomes of evolution**

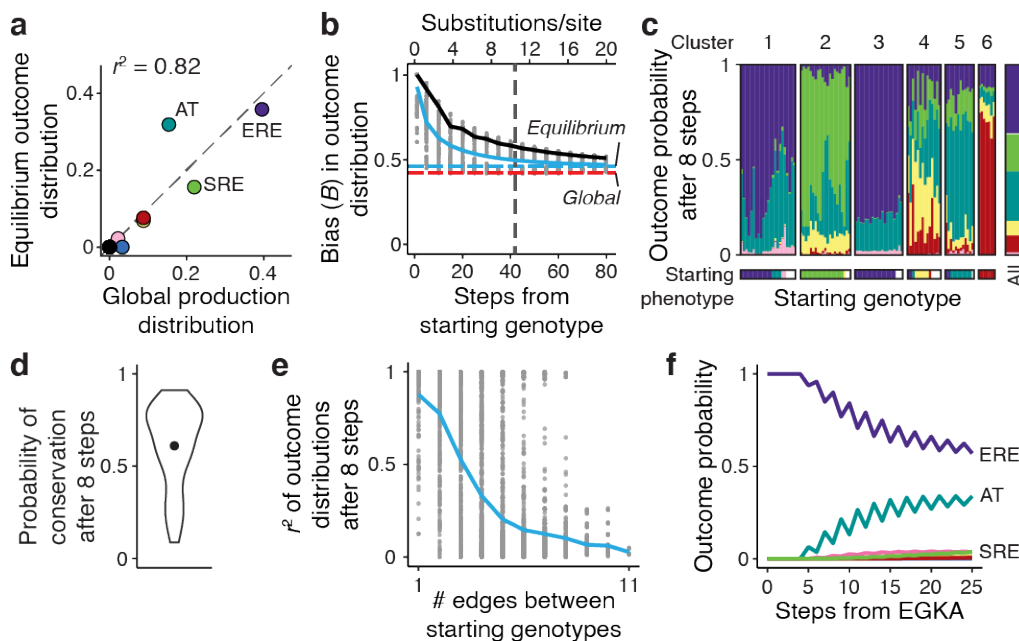
232 To characterize the potential influence of the AncSR1 GP map on the outcomes of evolution, we
233 modeled evolution on the network of functional amino acid genotypes as a discrete-time Markov
234 chain from every possible starting genotype given a variable trajectory length. Each time-step in
235 a trajectory is an amino acid substitution, the probability of which is weighted by the number of
236 single-nucleotide mutations that can mediate it; the relative probability of evolving a given
237 phenotype at the end of the trajectory is the sum of the probabilities of evolving all genotypes
238 that encode it. This model, which corresponds to neutral molecular evolution in which all
239 functional genotypes have equal fitness^{42,44}, represents a null scenario: the fixation process
240 imposes no biases on evolutionary outcomes except to prevent the loss of function via purifying
241 selection, thus allowing us to isolate the influence of biases imposed by the GP map on
242 evolutionary outcomes.

243 We first computed the equilibrium distribution of phenotypic outcomes after an infinite
244 number of substitutions. This represents the limiting case at which the distribution of outcomes is
245 insensitive to the starting genotype and does not change with additional substitutions. The
246 equilibrium outcome distribution is well correlated with the global production distribution (Fig.
247 3a, $r^2 = 0.82$), reflecting the constraints imposed by the global production bias. However, there
248 are differences: the equilibrium distribution is more biased ($B = 0.46$), and whereas ERE and SRE
249 specificity are the two most frequently encoded phenotypes, ERE and AT specificity are the
250 most likely equilibrium outcomes (Fig. 3a). This difference arises because most AT-specific
251 genotypes are located centrally within the network, while SRE-specific genotypes are in a more
252 peripheral cluster (Fig. 2b) and are therefore less likely to be occupied. The heterogeneous
253 connectivity of the GP network and global production bias therefore affect evolutionary
254 outcomes, even over infinitely long timescales.

255 On finite timescales, local bias strongly affects evolutionary outcomes. After 3
256 substitutions, for example—the shortest path between the historical ancestral and derived
257 genotypes—the outcome distributions are very strongly biased (mean $B = 0.8$ across starting
258 genotypes, Fig. 3b), because most genotypes can reach only a few new specificity phenotypes by
259 a path of this length (Extended Data Fig. 4). The bias in outcomes gradually decays as
260 trajectories get longer, but it takes 42 substitutions (10.5 per site) for the mean bias to decrease to
261 within 0.05 units of the equilibrium (Fig. 3b, vertical dashed line). By comparison, the maximum
262 root-to-tip branch length in the steroid receptor DBD phylogeny (Fig. 1b), which spans over 500
263 million years of evolution, is just 2.2 substitutions per site. The phenotypes likely to evolve on
264 phylogenetically relevant timescales are therefore strongly affected by local bias in the GP map.

265 Another consequence of local bias is that outcomes are strongly contingent on the genetic
266 starting point. Consider a trajectory length of 8 substitutions—long enough for new phenotypes
267 to become accessible from most starting points, but not so long that the influence of local bias is
268 lost. At this timescale, genotypes differ dramatically in the distribution of phenotypes that evolve

269 from them (Fig. 3c). Much of this variation is explained by the genotype cluster to which the
 270 starting node belongs (Fig. 3c), because evolutionary trajectories rarely jump between weakly
 271



272

273 **Fig. 3 | The AncSR1 GP map biases evolutionary outcomes towards phenotype**

274 **conservation.** **a**, Comparison between the global production distribution and the long-term
 275 equilibrium distribution of phenotypic outcomes in the AncSR1 main network. Each dot shows
 276 the frequency of one specificity phenotype in the two distributions. Black dot at the origin
 277 represents nine phenotypes not encoded in the map. Dashed gray line, $y = x$. Squared Pearson's
 278 correlation coefficient is shown. **b**, Strength of bias (B) in evolutionary outcomes as a function of
 279 the length of evolutionary trajectories. Each gray dot shows the B of the outcome distribution for
 280 trajectories of a given number of substitutions starting from one node on the main network
 281 component. Solid blue and black lines show the mean across all starting genotypes and from
 282 EGKA, respectively. Dashed horizontal red and cyan lines show B of the global production
 283 distribution and the equilibrium distribution, respectively. Vertical dashed line shows the number
 284 of substitutions required for mean B to reach within 0.05 units of the equilibrium value. The
 285 secondary x -axis (above) shows the trajectory length as substitutions per site. **c**, Distribution of
 286 evolutionary outcomes after 8 substitution steps from every starting genotype in the AncSR1
 287 main network component, organized by the cluster of the starting genotype (top). Bottom bar
 288 shows the phenotype of each starting genotype. Bars at right show the average outcome
 289 distribution for all starting genotypes. **d**, Distribution of the probability of phenotype
 290 conservation after 8 substitution steps across all specific starting genotypes in the AncSR1 main
 291 network component. Dot shows the mean. **e**, Evolutionary outcomes become less similar as
 292 starting genotypes diverge from each other. Each dot shows the similarity of the distributions of
 293 phenotypic outcomes (Pearson's r^2) of 8-step trajectories starting from a pair of genotypes,
 294 versus the number of network edges between the pair. Blue line, mean similarity across all pairs
 295 of starting genotypes. **f**, Probability of evolving each specificity phenotype starting from EGKA
 296 as a function of the number of substitutions.

297

298 connected clusters and clusters are strongly enriched for individual phenotypes. Even at this

299 timescale, the direction of phenotypic evolution on average favors conservation of the starting
300 phenotype (Fig. 3d), but when new phenotypes evolve, these too differ strongly among starting
301 genotype (Fig. 3c).

302 A final consequence of local bias is that as lineages diverge from each other across the
303 map, the distributions of phenotypic outcomes likely to evolve from them become increasingly
304 dissimilar. The correlation between the distributions of phenotypic outcomes after eight-step
305 evolutionary trajectories from pairs of starting genotypes depends strongly on the distance
306 between those genotypes in the network. For pairs of genotypes that are one substitution apart,
307 the average r^2 is 0.88, but this correlation drops to 0.50 when the genotypes are three steps apart
308 and is entirely lost at 11 steps ($r^2 = 0.02$, the maximum distance on the network) (Fig. 3e). Biases
309 in the outcomes of phenotypic evolution therefore become distinct among lineages as they
310 traverse the GP map.

311 **The AncSR1 GP map favored historical conservation of ERE specificity**

312 Local and global bias have a particularly strong and long-lasting impact on the outcomes of
313 evolutionary trajectories that begin from the historical genotype of the recognition helix in
314 AncSR1 (EGKA). It takes 80 substitutions for the bias in phenotypic outcomes from this starting
315 point to decay to within 0.05 units of equilibrium, almost double the average across genotypes
316 (Fig. 3b, blue vs. black solid lines). It takes at least 5 substitutions for any new specificity
317 phenotype to be accessed, and even after 8 substitutions the probability of conserving ERE
318 specificity is still 0.90 (Fig. 3f). The AncSR1 GP map heavily favors phenotypic conservation
319 from the historical starting genotype across phylogenetically relevant timescales. Bias imposed
320 by the GP map is therefore congruent with the long-term historical conservation of ERE
321 specificity in the lineages that descend from AncSR1 and lead to modern-day estrogen receptors.

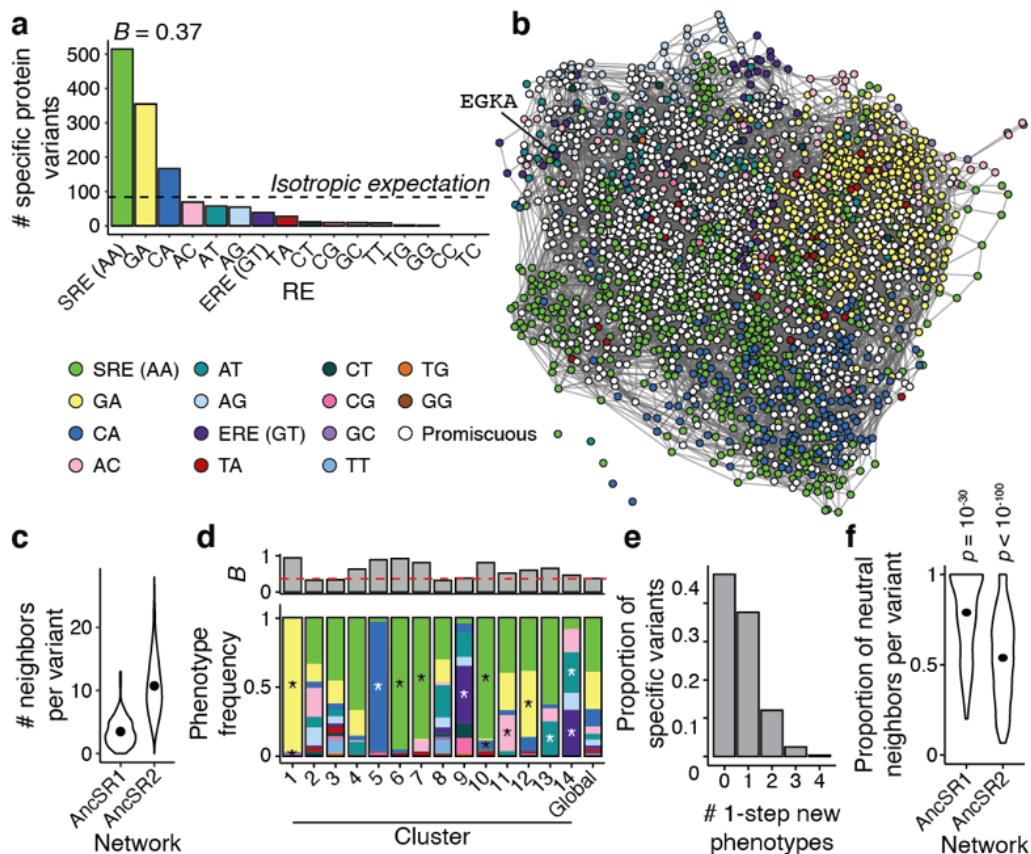
322 The historical outcome that evolved in AncSR1's other descendant lineage—acquisition of
323 SRE specificity in the kSR clade—was very unlikely on phylogenetic timescales. SRE-specific
324 genotypes are distant from EGKA (Fig. 2b), so the probability of evolving SRE specificity after
325 eight substitutions is only 0.0008 (Fig. 3f), despite the fact that this is the second-most frequently
326 encoded specificity phenotype in the network overall. Strong local bias around EGKA therefore
327 overrides the global bias towards SRE specificity, making the historical outcome in the kSR
328 clade extremely unlikely.

329 **Evolution of a different GP map in AncSR2**

330 Given that local bias made SRE specificity unlikely to evolve from the ancestral genotype in the
331 AncSR1 map, how could this phenotype have historically evolved in the kSRs? We reasoned that
332 the GP map must have changed along the branch leading to AncSR2 when SRE specificity was
333 acquired. Previous experiments showed that the background substitutions that occurred outside
334 the recognition helix during this interval had a nonspecific permissive effect on both ERE and
335 SRE activation, allowing the protein to tolerate the historical substitutions and other mutations in
336 the RH (Fig. 1b)^{12,24}. We predicted that the background substitutions had a similarly permissive
337 effect across all REs, increasing the number of functional genotypes in the map and the number
338 of phenotypes they encode, including SRE specificity and others.

339 To assess this hypothesis, we characterized the GP map of the RH sites in AncSR2 and
340 compared it to the map in the AncSR1 background. As predicted, the number of functional
341 genotypes and phenotypes both massively increased (Fig. 4a, b). There are 2,407 functional
342 protein genotypes in the AncSR2 map, an increase of >20-fold over the AncSR1 background.
343 Fourteen of the 16 possible specificity phenotypes are now encoded in the map, twice as many as

344 in AncSR1 (Fig. 2a, 4a). The background substitutions therefore dramatically expanded the
 345 functional genetic and phenotypic variation that can be produced within the recognition helix.
 346



347
 348 **Fig. 4 | Global and local bias and connectivity changed in the AncSR2 GP map.** **a**, Global
 349 production distribution and global B of the AncSR2 GP map. **b**, Sequence space network of the
 350 AncSR2 GP map. **c**, Number of one-step neighbors per protein variant in each network. Dots
 351 show the mean of each distribution. **d**, Bottom: Frequencies of specificity phenotypes within
 352 each genotype cluster (1–14, ordered by decreasing size); the global production distribution is
 353 shown for comparison. Only the 14 largest clusters, which contain >90% of genotypes, are
 354 shown. Asterisks, phenotypes significantly enriched within a cluster relative to the global
 355 production distribution (Fisher’s exact test, $p < 0.05$ after Bonferroni correction). Top: strength
 356 of phenotype bias (B) in each cluster. Red line, B of global production distribution. **e**,
 357 Distribution of the number of new phenotypes accessible within one mutation, across all RE-
 358 specific protein variants in the AncSR2 main component. **f**, Proportion of neutral neighbors per
 359 RE-specific variant in the main network component of the AncSR1 and AncSR2 maps. Dots
 360 show the mean. p -value, probability that the mean would be at least as great as observed if
 361 phenotypes were randomly reassigned in the main component of each map (AncSR1 $n = 91$,
 362 AncSR2 $n = 2,402$).
 363

364 Connectivity between genotypes in the map increased, reducing local bias and facilitating
 365 access to new phenotypes. In the AncSR2 network, all but five of the 2,407 functional nodes are
 366 connected in a single main component (Fig. 4b), and the mean number of edges per node is 10.7,
 367 a three-fold increase compared to the AncSR1 network (Fig. 4c). Genotype clusters are still
 368 present, but bias within clusters is weaker than in the AncSR1 map (Fig. 2c, 4d). As a

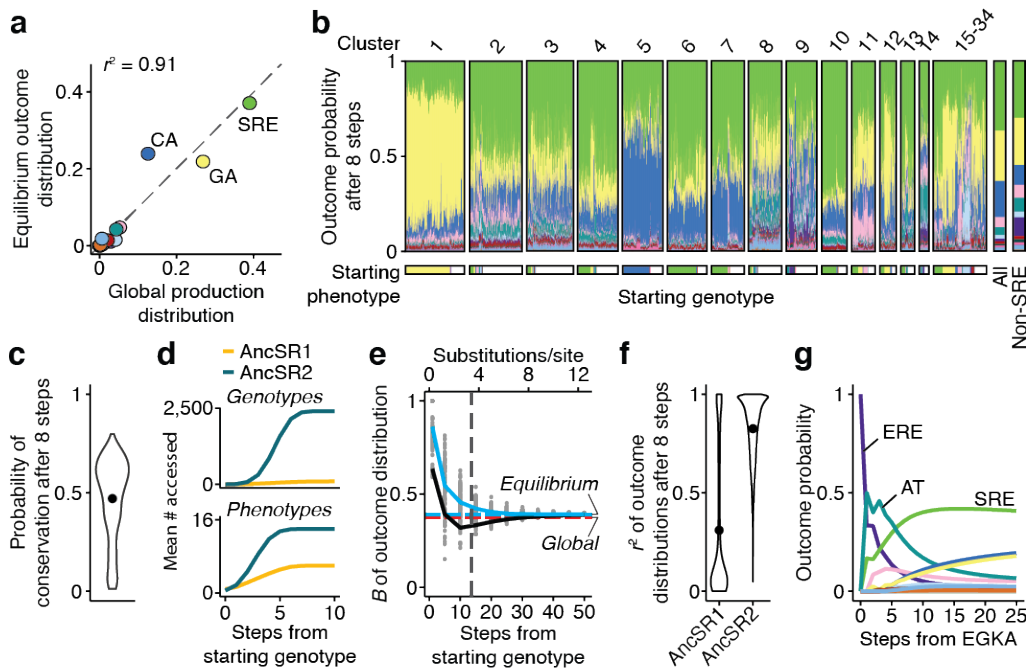
369 consequence, genotypes have more access to new phenotypes: >50% of genotypes in the
 370 AncSR2 map can access between 1 and 4 new phenotypes within a single mutation (Fig. 4e,
 371 compare to Fig. 2e), because genotypes are typically connected to far more non-neutral
 372 neighbors (Fig. 4f).

373 Finally, the global production distribution of phenotypes also changed across this
 374 interval. In the AncSR2 map, SRE became the most frequently encoded phenotype (39% of
 375 specific variants), and ERE's rank declined from first to seventh (encoding just 3% of specific
 376 variants) (Fig. 2a, 4a). The background substitutions therefore realigned the global phenotype
 377 bias from favoring the ancestral specificity to producing the derived specificity.

378 **The AncSR2 GP map favored evolution of SRE specificity**

379 These changes in the AncSR2 GP map dramatically altered the likely phenotypic outcomes of
 380 evolution. At long-term equilibrium using our Markov model and the AncGR2 map, the most
 381 likely evolutionary outcome is now SRE specificity, with a probability close to 40% (Fig. 5a,
 382 compared to <20% in the AncSR1 map). At moderate timescales as well, SRE specificity is the
 383 most likely outcome across the majority of starting genotypes (Fig. 5b). The probability of
 384 evolving new phenotypes overall is considerably higher in the AncSR2 network compared to
 385 AncSR1 (mean probability of conservation after 8 steps 0.47 in AncSR2 but 0.61 in AncSR1,
 386 Fig. 3d, 5c).

387



388

389 **Fig. 5 | The AncSR2 GP map biases evolutionary outcomes towards SRE specificity.** a,
 390 Comparison between the global production distribution and the long-term equilibrium
 391 distribution of phenotypic outcomes in the AncSR2 main network. Dashed gray line, $y = x$. b,
 392 Distribution of evolutionary outcomes after 8 substitution steps from every starting genotype in
 393 the AncSR2 main network component, organized by the cluster of the starting genotype (top).
 394 Bottom bar shows the phenotype of each starting genotype. Bars at right show the average
 395 outcome distribution for all starting genotypes and all non-SRE-specific starting genotypes,
 396 respectively. c, Distribution of the probability of phenotype conservation after 8 substitution
 397 steps across all specific starting genotypes in the AncSR2 main network component. Dot shows

398 the mean. **d**, Number of genotypes (top) and phenotypes (bottom) accessible as a function of the
399 length of evolutionary trajectories. Lines show the mean across all starting genotypes in each
400 network. Gold, AncSR1 network; teal, AncSR2 network. **e**, Strength of bias (B) in evolutionary
401 outcomes as a function of the length of evolutionary trajectories. Lines and colors are the same
402 as in Fig. 3b. **f**, Distribution of the similarity in outcome distributions (Pearson's r^2) for 8-step
403 trajectories starting from all pairs of genotypes in the AncSR1 and AncSR2 main networks. Dots
404 show means. **g**, Probability of evolving each specificity phenotype starting from EGKA as a
405 function of the number of substitutions.

406

407 These changes in evolutionary outcomes are attributable to the increased connectivity of
408 the AncSR2 network and the shift in the global production distribution. From any starting point,
409 the increase in functional nodes and connectivity allows access to far more genotypes and new
410 phenotypes (Fig. 5d). As a result, the influence of local bias is lost faster, and trajectories more
411 rapidly converge on the equilibrium distribution (Fig. 5e), which more closely resembles the
412 production distribution than in the AncSR1 background (Fig. 5a). Evolutionary outcomes are
413 also more similar across pairs of starting points than they were in the AncSR1 map (Fig. 5f).
414 Combined with the shift in the global production distribution, this causes SRE specificity—
415 which was already the second-most likely outcome in the AncSR1 map—to become the most
416 likely outcome from a majority of starting points in the AncSR2 background.

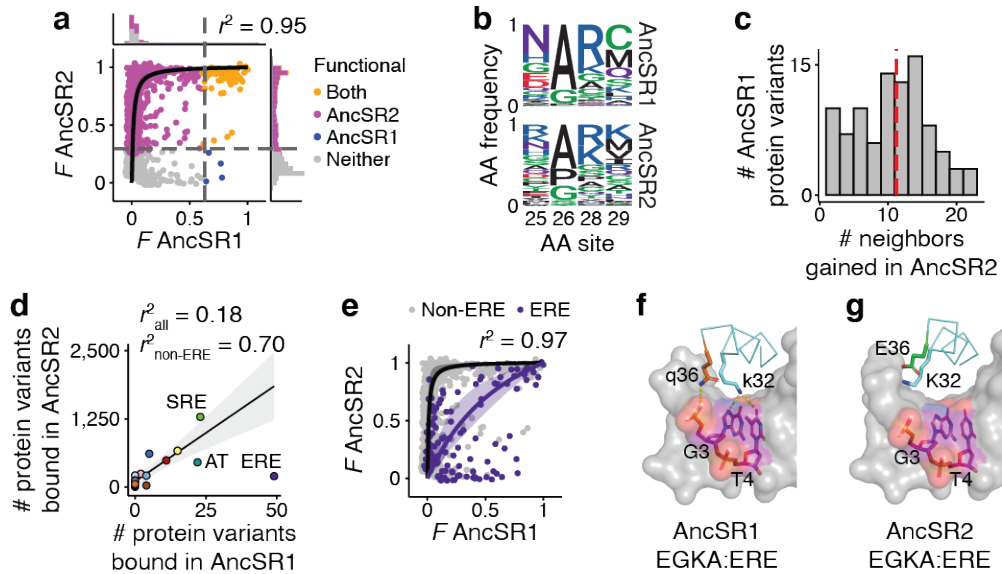
417 From the historical RH genotype EGKA (the AncSR2 protein with the RH states reverted
418 to their ancestral states), the likely outcomes of phenotypic evolution are dramatically different
419 than in the AncSR1 map. EGKA is much less mutationally isolated in the AncSR2 network, so
420 the probability of conserving ERE specificity after 8 substitutions drops from 0.9 in the AncSR1
421 map (Fig. 3f) to 0.07 in the AncSR2 map (Fig. 5g). The probability of evolving new specificity
422 phenotypes on moderate timescales increases accordingly: after just three steps, two new
423 phenotypes—including SRE specificity—are more likely than conservation of ERE. By six
424 steps, SRE specificity becomes the most likely of all phenotypic outcomes.

425 The background substitutions that occurred along the branch to AncSR2 therefore
426 changed the GP map of the RH in a way that dramatically changed the probable phenotypic
427 outcomes of evolution. This map strongly favors phenotypic diversification, and it makes the
428 particular phenotype that historically evolved in the kSR lineage the most likely of all possible
429 outcomes.

430 **Simple biophysical mechanisms changed the GP map**

431 Finally, we sought insight into the biophysical mechanisms that changed the GP map of the
432 recognition helix between AncSR1 and AncSR2. Although our experiments provide a functional
433 rather than biophysical readout, different biophysical mechanisms predict different patterns of
434 functional change between the AncSR1 and AncSR2 maps. We therefore analyzed the change in
435 fluorescence of each protein-DNA complex variant between the two backgrounds to identify
436 potential biophysical mechanisms and considered them in light of existing crystal structures. We
437 found evidence for two major mechanisms.

438 First, the background substitutions between AncSR1 and AncSR2 appear to have caused
439 a universal increase in affinity across all protein-DNA complexes. Previous experiments and
440 crystal structures showed that the background substitutions improve nonspecific DNA contacts
441 and binding cooperativity to both ERE and SRE^{12,24}; we therefore hypothesized that affinity
442 increased universally for all amino acid variants across all 16 REs. To test this hypothesis, we fit
443 a simple model in which fluorescence in each ancestral background is a function of a complex's



444

445 **Fig. 6 | Nonspecific effects of background substitutions on DBD-RE affinity.** **a**, Fluorescence
 446 of each complex in the AncSR1 vs. AncSR2 background, scaled between the upper and lower
 447 bounds for each background. Curve shows best-fit model assuming that the affinity of every
 448 complex in the AncSR2 background is related to its affinity in the AncSR1 background by the
 449 same scaling factor. Shaded region around the curve (barely visible) shows bootstrapped 95%
 450 confidence interval (CI). The Pearson's r^2 between the data and model predictions is shown ($n =$
 451 $2,627$). Histograms show distribution of F in each background. Dashed lines show the
 452 fluorescence of the wild type complex in each background (AncSR1-EGKA:ERE or AncSR2-
 453 GSKV:SRE). Colors indicate the backgrounds in which each genotype is functional. **b**, Amino
 454 acid frequencies at the variable RH sites across all functional protein variants in the AncSR1 and
 455 AncSR2 maps. **c**, Distribution of the number of neighbors gained in the AncSR2 background
 456 across all functional protein variants in the AncSR1 background that remain functional in the
 457 AncSR2 background. Dashed line, mean. **d**, Correlation between the number of protein variants
 458 bound per RE in each background. Black line, linear fit to all REs except ERE; shaded region,
 459 95% CI. **e**, Same as **a**, but fitting a model in which the background substitutions affect affinity of
 460 all variants for ERE by one scaling factor and for all other REs by a different scaling factor.
 461 Purple, observed fluorescence and best-fit model predictions for ERE complexes; gray, for non-
 462 ERE complexes. **f**, Crystal structure of the AncSR1-EGKA protein in complex with ERE (PDB
 463 4OLN). The RH backbone is shown as a ribbon, with key side chains shown as sticks. The gray
 464 surface shows ERE, with variable bases and backbone as sticks. In this complex, glutamine (q) at
 465 site 36 forms a hydrogen bond (yellow dashed line) with the DNA backbone, and lysine (k) at
 466 site 32 forms two hydrogen bonds to the ERE-specific bases G and T. **g**, Same as **f**, but with the
 467 AncSR2-EGKA crystal structure (PDB 4OND). Substitution to glutamic acid (E) at site 36
 468 abolishes the ancestral hydrogen bond to the DNA backbone and results instead in electrostatic
 469 repulsion from the backbone. This deforms the recognition helix, abolishing the hydrogen bonds
 470 between K32 and the G and T bases. In **F** and **G**, lowercase letters represent ancestral amino acid
 471 states, and uppercase derived.

472

473 affinity, and affinity is scaled by a constant factor in AncSR2 relative to AncSR1. The model fits
 474 the data very well ($r^2 = 0.95$; Fig. 6a), with an estimated 70-fold universal improvement in
 475 affinity in the AncSR2 background. This apparent increase in affinity explains the vast increase

476 in functional genotypes and specificity phenotypes between AncSR1 and AncSR2, because many
477 protein-DNA complexes that had weak affinity in the AncSR1 background—and were therefore
478 nonfunctional—bind strongly enough in AncSR2 to produce functional levels of fluorescence.
479 The number of promiscuous protein variants also increases, because many variants cross the
480 threshold for functionality on multiple REs (Extended Data Fig. 5a). A universal improvement in
481 affinity explains not only the increased size but also the greater connectivity of the AncSR2
482 network: the background substitutions do not qualitatively change the amino acid determinants
483 of binding but instead make them less stringent (Fig. 6b), so many of the newly functional nodes
484 in AncSR2 are close neighbors of those that were already functional in AncSR1, with an average
485 gain of 11 new neighbors per node (Fig. 6c).

486 The second apparent mechanism is that the background substitutions negatively affect
487 specific binding to ERE, shifting the global production bias away from ERE and leaving SRE as
488 the most-encoded phenotype in the AncSR2 background. A universal affinity increase predicts
489 that the number of variants with every specificity phenotype should increase proportionally
490 across the AncSR1-AncSR2 interval; this pattern holds, but ERE is an outlier, with far fewer
491 variants than would be expected given the pattern for other phenotypes (Fig. 6d). Moreover, ERE
492 complexes exhibit notably lower fluorescence in the AncSR2 background than predicted by a
493 universal increase in affinity (Extended Data Fig. 5b). We estimated the effect of the background
494 substitutions on ERE affinity by incorporating a background-by-ERE interaction term into our
495 affinity-fluorescence model; adding this parameter improves the fit to the data ($r^2 = 0.97$), with
496 the background substitutions improving ERE affinity by an estimated 2.3-fold, compared to 99-
497 fold for all other REs (Fig. 6e). The extent of the relative reduction in fluorescence differs among
498 protein variants, however, suggesting additional specific interactions between background
499 substitutions and amino acids in the recognition helix (Extended Data Fig. 5c). Crystal structures
500 of the EGKA:ERE complex¹² suggest a possible structural basis for the global reduction in ERE
501 affinity: one of the background substitutions (q36E) deforms the protein backbone of the
502 recognition helix, abolishing two hydrogen bonds that are formed between a conserved residue
503 and bases in the ERE (Fig. 6f, g). Corroborating this mechanism, the background substitutions
504 also shift the global bias away from AT specificity (Fig. 6d), and this is the only other RE that
505 can form these hydrogen bonds.

506 The structure of the GP map therefore changed between AncSR1 and AncSR2 via two
507 simple biophysical mechanisms. By increasing all proteins' affinity for all REs, while also
508 impairing their affinity for ERE, the background substitutions reduced local bias and changed the
509 direction of global bias, facilitating the evolution of many new genotypes and phenotypes and
510 shifting the protein's global propensity away from conserving ERE specificity to evolving the
511 new specificity for SRE.

512 **Robustness to assumptions**

513 To assess whether our conclusions are sensitive to assumptions that we made in our analysis, we
514 reanalyzed our experimental data under different models and assumptions. First, we applied
515 different thresholds to classify genotypes as functional or nonfunctional, included promiscuous
516 genotypes when characterizing global production distributions, and characterized these
517 distributions using only genotypes with experimentally measured phenotypes. In every case, we
518 observed similar forms of bias in both the AncSR1 and AncSR2 GP maps to those reported
519 above (Extended Data Fig. 6).

520 Second, instead of treating the protein as an evolutionary unit independent of the RE, we
521 repeated our analyses using an alternative sequence space network in which the protein and RE

522 coevolve as a complex. In this model, evolution may occur via single-step amino acid mutations
523 in the protein or nucleotide mutations in the RE. Our main conclusions again hold: global and
524 local biases impact phenotypic evolution over long and short timescales, favoring ERE
525 conservation in the AncSR1 map and evolution of SRE specificity in AncSR2 (Extended Data
526 Fig. 7).

527 Finally, we addressed uncertainty about the ancestral sequences. AncSR1 and AncSR2
528 DBD reconstructions have very high confidence, containing just five and zero ambiguously
529 reconstructed sites, respectively³³. Experimental data from a prior single-mutant DMS study
530 show that the effects of mutations in the RH are virtually identical when they are introduced into
531 the AncSR1 background or into an alternative reconstruction of AncSR1 that incorporates all
532 plausible alternative amino acids at the ambiguously reconstructed sites ($r^2 > 0.99$; Extended
533 Data Fig. 8)³³. The very limited uncertainty about the AncSR1 ancestral sequence is therefore
534 likely to have little or no effect on our conclusions.

535 **The GP map was a cause of historical phenotypic evolution**

536 Our data establish that global and local biases in the two ancestral GP maps we studied were
537 causal factors in the historical lineage-specific evolution of DNA specificity. Establishing
538 causality in a multifactorial framework requires 1) evidence that a putative cause increases the
539 probability of the outcome(s) of interest, and 2) evidence for a specific mechanism by which the
540 cause affects the outcome's probability⁴⁵. Concerning the first requirement, our experiments
541 show that biases in the AncSR1 map increased the probability that ERE specificity would be
542 evolutionarily conserved, and biases in the AncSR2 map increased the probability that SRE
543 specificity would be acquired. The second requirement is satisfied by a simple axiom of
544 population genetics: the probability that a phenotype will evolve is the product of its probability
545 of production and its probability of fixation under the influence of selection and drift. If biases in
546 the GP map increase the production probability, then evolutionary outcomes will in turn be
547 biased.

548 A cause must precede its effect. The biases that favored the conservation of ERE
549 specificity in the AncSR1 map are ancestral to the ER lineage in which that outcome occurred
550 (Fig. 1b). This map persisted unchanged for hundreds of millions of years of phenotype
551 conservation, because zero amino acid changes anywhere in the DBD occurred along the
552 descendant branches leading from AncSR1 to ER α in the ancestor of all bony vertebrates. Even
553 most present-day ER α DBDs contain zero or at most a single substitution relative to AncSR1
554 (Extended Data Fig. 9). As for the acquisition of SRE specificity in the AncSR2 lineage, the
555 global bias that favors production of SRE specificity as the second-most encoded phenotype was
556 already present in the AncSR1 map. Further, the massive increase in connectivity of the AncSR2
557 map, which dramatically increased the propensity for new phenotypes to evolve, must have been
558 acquired before SRE specificity actually evolved, because the recognition helix substitutions that
559 conferred SRE specificity during history cannot be tolerated unless the background substitutions
560 that nonspecifically increased DNA affinity occurred first¹². Our experiments do not resolve
561 whether the third major property of the AncSR2 map—a shift in global bias away from ERE
562 specificity that further enhanced the propensity to encode SRE specificity—occurred before or
563 after this phenotype was historically acquired.

564 We do not argue that selection played no historical role in the evolution of specificity. It
565 seems likely that purifying selection would have favored conservation of ERE specificity in the
566 chordate ERs, and positive selection could have contributed to fixation of SRE specificity in the

567 AncSR2 lineage. If so, however, selection would have further increased the probability of
568 outcomes that were already favored by the biases imposed by the GP map.

569 Our data show that the GP map's influence is strong enough to override the influence of
570 selection in many cases. For example, some global biases we observed are absolute. There are 9
571 specificity phenotypes that cannot be encoded at all in the AncSR1 GP map, and two cannot be
572 encoded in AncSR2; these could phenotypes never evolve, no matter how large a fitness benefit
573 they might confer. Local bias is also absolute in many cases: from every starting point, the vast
574 majority of phenotypes are impossible to produce directly by mutation, and most require many
575 substitutions before they become accessible. Selection would therefore be powerless to fix these
576 phenotypes over short or medium timescales. The GP map limited evolution to a small subset of
577 possible phenotypes; history, further influenced by selection and chance, played out within this
578 set.

579 There is evidence that features similar to those we observed in the steroid receptor GP
580 map affect biological systems and their evolution across levels of organization. Global bias is
581 apparent in other molecular^{46,47} and developmental systems⁴⁸⁻⁵¹, and the resulting biases are
582 often congruent with natural patterns of diversity⁵²⁻⁵⁵. Local bias also appears to be widespread,
583 because most random mutations are phenotypically neutral if they are tolerated^{51,56-59}, and long-
584 term phenotype conservation is widespread in the fossil record⁶⁰. When new phenotypes are
585 acquired, identical perturbations often yield different phenotypes in different lineages⁶¹⁻⁶⁴, and
586 convergent evolution becomes less likely among distantly related lineages⁶⁵. As lineages evolve
587 across their GP maps, their biology inevitably changes, imposing new biases on the production
588 and future evolution of genotypes and phenotypes. It therefore seems likely that anisotropy and
589 heterogeneity are near-universal characteristics of GP maps^{2,39,41}, and that the biases these
590 properties create have shaped large-scale patterns of phenotype conservation and lineage-specific
591 evolutionary change across the tree of life.

592 Our study differs in kind from previous combinatorial DMS studies, which have
593 addressed the distribution in sequence space of just one or a few phenotypes that are encoded by
594 extant proteins, rather than the space of all possible phenotypes^{22,24,25,63,66-68}. These studies have
595 shown that sequence landscapes are rugged, so the probability of reaching particular genotypes
596 encoding those phenotype may depend on the starting point and intermediate mutational steps.
597 Because those studies take the phenotypic "destination" for granted, they cannot address why
598 those phenotypes, rather than all the other conceivable outcomes, exist at all.

599 Our work shows that as a protein or other biological system moves through sequence
600 space, the set of phenotypes that it can produce changes at every step. Life is astonishing in its
601 diversity, but an even deeper puzzle lies in the fact that only a tiny fraction of conceivable
602 phenotypes have ever evolved, and those which have evolved are mostly limited to particular
603 taxa^{14-16,69}. Chance and selection are likely important factors in explaining the patchy
604 distribution of phenotypes on Earth. But the very particular biology we observe today must also
605 reflect the constantly changing potential of biological systems, as they vary and diverge, to
606 generate new forms of life at every moment in time.

607 **Methods**

608 ***RE reporter strains***

609 To measure binding of SR DBD to the 16 RE variants, we adapted a yeast GFP reporter system
610 previously developed to measure binding to ERE and SRE, where GFP expression is well
611 correlated with DNA affinity over a range of at least 2 M^{-2} ($r^2 = 0.74$)³³. We engineered 16 yeast
612 strains, each of which reports on binding of the DBD to one RE. We modified the yeast strain
613 CM997 (YPS1000 MATa ho::KMX)⁷⁰ to replace the *KMX* gene at the *HO* locus with a construct
614 containing yeast-enhanced GFP downstream of a minimal *CYCI* promoter with an array of four
615 palindromic RE sites (tcaAGN^uNCacagTGNNCTga), each separated by a 19-nt sequence, along
616 with a *HygR* gene. To ensure a consistent dynamic range of fluorescence across strains, we made
617 changes to two RE strains in the nucleotide sequences flanking the palindromes at sites that do
618 not affect specificity^{34,35} (see Supplementary Methods for details). These constructs were
619 transformed into yeast using the lithium acetate method⁷¹ and selected for resistance to
620 hygromycin and susceptibility to G418; integration was confirmed by Sanger sequencing.

621 To validate this reporter system, we measured fluorescence of each strain in the presence and
622 absence of a DBD variant with universally high affinity to all REs (AncSR1+11P+GGKA)^{12,29}.
623 We used a low-copy yeast vector (pDBD) to express this DBD variant as a C-terminal fusion
624 with an SV40 nuclear localization signal and a *S. cerevisiae* Gal4 activation domain (Gal4AD)
625 under control of a pGAL1 promoter. We transformed this construct into each yeast strain using
626 the lithium acetate method followed by G418 selection (50 $\mu\text{g}/\text{mL}$). Single colonies were
627 inoculated in YPD+G418 and transferred to YPGal+G418 media for 6 hours to induce DBD
628 expression. GFP fluorescence was measured on a BD LSRFortessa flow cytometer using a 488
629 nm laser with 505 nm long pass and 525/50 nm band pass filter. We used as the metric of
630 fluorescence $\log_{10}(\text{GFP}/\text{FSC-A}^{1.5})$, which normalizes fluorescence to cell volume. All 16 strains
631 showed DBD-dependent fluorescence across a similar dynamic range (Extended Data Fig. 1a–c).

632 ***AncSR1 and AncSR2 combinatorial library construction***

633 We used as the wild-type protein sequences the maximum *a posteriori* AncSR1 and AncSR2
634 DBD sequences inferred from a maximum likelihood phylogeny of nuclear receptors³³.

635 We optimized codon usage for yeast and cloned the ancestral DBDs into the pDBD2.1
636 expression vector, which is modified from the pDBD vector^{24,33} to express GFP at a level within
637 the dynamic range of fluorescence for the wild type AncSR1:ERE and AncSR2:SRE complexes.
638 A bidirectional pGAL1/GAL10 promoter simultaneously drives DBD and mCherry expression,
639 which allowed us to monitor plasmid retention in yeast (Extended Data Fig. 1d).

640 Combinatorial mutant libraries were created by synthesizing oligos (IDT) with degenerate NNS
641 codons to encode all 20 amino acids and a stop codon at four recognition helix sites of each
642 ancestral protein (Extended Data Fig. 1e). To distinguish sequencing reads coming from different
643 RE strains, 16 synonymously barcoded versions of the library were designed for each
644 background (Extended Data Fig. 1e, Supplementary Table 1). Each barcode (REBC) differed by
645 at least three nucleotides to ensure accurate read assignment despite sequencing errors. The
646 oligos were cloned into the pDBD2.1 vector using the BsaI-HF Golden Gate Assembly kit
647 (NEB), transformed into Invitrogen ElectroMAX DH5 α -E *E. coli*, and maxiprepmed
648 (Supplementary Methods). Transformation yields exceeded 1.08×10^7 cfu per barcoded library,
649 providing 56-fold coverage of the amino acid library size (Supplementary Table 2). Assemblies

650 were validated by Sanger sequencing of independent transformants and PCR of the plasmid
651 libraries to confirm the correct insert size.

652 Maxiprepped libraries (GenElute HP, Sigma-Aldrich) were transformed into the yeast reporter
653 strains using an optimized yeast electroporation protocol (Supplementary Methods).
654 Transformation yields exceeded 10^7 cfu per library (50-fold coverage), estimated by dilution
655 plating (Supplementary Table 2). Yeast libraries were flash-frozen in liquid N₂ in 200 OD₆₀₀-mL
656 aliquots with 25% glycerol and stored at -80°C . Multiple transformant rates estimated from
657 Sanger sequencing of individual colonies⁷² were estimated to result in 0.03% or fewer cells with
658 multiple plasmid copies at time of sorting.

659 *Cell sorting*

660 We used fluorescence-activated cell sorting (FACS) to separate cells based on their GFP
661 expression. We performed two rounds of sorting: an initial “enrichment sort” to enrich for GFP+
662 variants in the full libraries, and a second, higher resolution “binned sort” on the enriched
663 libraries to generate quantitative fluorescence estimates for each variant. Enrichment sorting was
664 performed in batches of 8 libraries. Two glycerol stocks per library were thawed on ice, after
665 which cells were recovered for 2 hours in 400 mL YPD+chloramphenicol (chlor) per library at
666 30°C and 225 rpm. After recovery, G418 was added to the culture and a sample of cells was
667 taken for dilution plating. We recovered a minimum of 1.6×10^7 cfu per library (82-fold
668 coverage). After 15 hours of overnight growth, libraries were washed once in PBS, resuspended
669 to OD₆₀₀ 0.25 in 50 mL YPGal+G418, and grown for 6 hours to induce DBD expression. Cells
670 were then spun down, washed once in PBS, resuspended in 5 mL PBS, and kept on ice for
671 sorting.

672 Sorting was performed at the University of Chicago Cytometry and Antibody Technology
673 Facility on a BD FACSARIA Fusion machine. We used a 488 nm laser with 495 nm long pass
674 filter and 515/20 nm band pass filter for GFP detection, and a 561 nm laser with 595 nm long
675 pass filter and 610/20 nm band pass filter for mCherry detection. After gating on homogeneous
676 single cells and mCherry expression, we sorted cells into GFP⁻ and GFP⁺ populations (Extended
677 Data Fig. 1f). To normalize fluorescence to cell volume, GFP gates were drawn to have a slope
678 of 1.5 on a log(FSC-A)-log(GFP) plot. We sorted 2.5×10^7 cells per library in the enrichment
679 stage (129-fold coverage, Supplementary Table 2).

680 Enriched cells from different libraries were pooled by GFP bin and grown in either 700 mL
681 (GFP⁺) or 2 L (GFP⁻) of YPD+G418+chlor. Cultures were grown overnight at 225 rpm and 22–
682 30°C , depending on the ratio of cells to media, until they were at least OD₆₀₀ 3 but not yet
683 saturated. 200 OD₆₀₀-mL 25% glycerol stocks were then made for both the GFP⁺ and GFP⁻
684 cultures. 10 OD₆₀₀-mL of the GFP⁻ culture was used for plasmid extraction using a previously
685 described protocol²¹.

686 The binned sort was performed to yield three replicates per library. For each replicate, two 200
687 OD₆₀₀-mL glycerol stocks of GFP⁺ cells per enrichment sort batch were thawed on ice,
688 recovered in 400 mL YPD+chlor for 2 hours, and sampled for dilution plating. After adding
689 G418, cultures were grown overnight, achieving a recovery rate at least 4X the number of GFP+
690 cells collected during the enrichment sort (Supplementary Table 3). Overnight cultures were
691 pooled proportionally to the GFP⁺ cell counts from the enrichment sort, yielding a total of 100
692 OD₆₀₀-mL. The pooled cells were washed with PBS, induced for DBD expression in 400 mL
693 YPGal+G418 for 6 hours, washed again, resuspended in 40 mL PBS, and kept on ice for sorting.
694 Binned sorting followed the enrichment sort protocol but used four GFP bins instead of two

695 (Extended Data Fig. 1g), with $\sim 1.6 \times 10^8$ cells collected per replicate. The number of sorted cells
696 and recovered reads was consistent across libraries and replicates (Supplementary Table 4).

697 *Deep sequencing*

698 After sorting, cells were grown in 100 mL YPD+G418+chlor per 10^7 sorted cells, or at least 100
699 mL per bin. Cultures were grown overnight to at least OD₆₀₀ 3.0 but not yet saturated, and 50
700 OD₆₀₀-mL was collected per 10^7 sorted cells for plasmid extraction.

701 Sequencing libraries were constructed from plasmids extracted from the enrichment sort GFP–
702 population and the four binned sort populations using two rounds of amplification. In the first
703 round, the RH scanning and REBC regions of the DBD were amplified with primers that added a
704 6-nt barcode for bin and replicate identification (BRBC)⁷³. For every 10 OD₆₀₀-mL of yeast used
705 for plasmid extraction, 3 μ L of plasmid template was used in a 10 μ L Q5 PCR reaction (NEB).
706 AncSR1- and AncSR2-specific primers were mixed proportionally to background-specific cell
707 counts (estimated from flow cytometry) to minimize amplification bias. To introduce nucleotide
708 diversity for improved cluster identification during Illumina sequencing, eight unique forward
709 and reverse primer pairs were used per bin and background to encode frameshift diversity and
710 attach read 1 primer sequences in both directions. PCR conditions included 52°C annealing for
711 13 cycles. Reactions were then pooled by bin/replicate and purified using the Zymo DNA Clean
712 & Concentrator Kit. In the second round, half of the first-round product was amplified with
713 primers to add Illumina P5 and P7 adapter sequences. PCR was performed in 50 μ L Q5 reactions
714 (NEB) per 10 μ L round 1 product reaction at 68.4°C annealing for 12 cycles. The final product
715 was size-selected on a 2% agarose gel, excised, purified using the Qiagen Gel Extraction Kit, and
716 re-purified with the Zymo DNA Clean & Concentrator Kit.

717 Final sequencing library concentrations were quantified by Qubit. Libraries were pooled
718 according to the number of cells sorted per bin/replicate, and 1.8 pM dilutions were prepared
719 according to Illumina’s standard protocol. Replicate 1 of the binned sort libraries was sequenced
720 on a NextSeq High Output run. The remaining replicates were sequenced on a NovaSeq S1 run
721 at the University of Chicago Genomics Facility. We used standard read primers and 86 cycles for
722 read 1 and 80 cycles for read 2. This enabled us to bidirectionally sequence the region containing
723 the variable RH codons and REBC.

724 *Mean fluorescence estimation, data cleaning and validation*

725 Sequencing reads were processed using a custom pipeline. We used *sickle* v1.33⁷⁴ to filter reads
726 based on their quality: we kept reads with a Phred score ≥ 30 and a minimum length of 79
727 nucleotides. We then used *PEAR* v0.9.6⁷⁵ to merge the trimmed paired-end reads (minimum
728 assembly length 100 nucleotides). Finally, we used Biopython toolkit v1.79⁷⁶ to demultiplex the
729 assembled reads by DBD background, REBC, and BRBC. We only considered reads that
730 mapped exactly to the DBD background and allowed reads with at most one mismatch in the
731 REBC and one in the BRBC.

732 The mean fluorescence for protein:RE complexes observed in the binned sort data was estimated
733 as previously described³³. We first estimated the proportion of cells of each complex g in each
734 bin b ($c_{g,b}$) from the proportion of reads in b that mapped to g . The mean fluorescence estimate
735 F_g for each complex was then estimated by taking the weighted mean fluorescence across bins
736 (mean fluorescence of each bin was measured during sorting), with weights $c_{g,b} / \sum_b c_{g,b}$.

737 We applied several filtering and correction steps to reduce global measurement error and
738 normalize fluorescence estimates between replicates. First, complexes with fewer than 27 reads
739 per replicate were removed to ensure >95% had a standard error (SE) of ≤ 0.1 (5% of the assay
740 range; Extended Data Fig. 2a). Second, complexes observed in only one replicate were excluded.
741 Third, batch effects were corrected by fitting I-splines to normalize fluorescence between
742 replicates (Extended Data Fig. 2b). Finally, SE was recalculated and complexes with $SE > 0.1$
743 were removed (Extended Data Fig. 2c). The final dataset had a mean pairwise Pearson's $r^2 =$
744 0.55 across replicates. The poor correlation arises primarily because the vast majority of
745 complexes are at the lower fluorescence bound, so r^2 is dominated by measurement noise; for
746 variants with fluorescence above the lower bound (roughly $F \geq -4.0$), r^2 improved to 0.92.
747 Altogether, we obtained fluorescence estimates for 628,732 AncSR1 and 658,475 AncSR2
748 variants, covering 24.6% and 25.7% of possible variants, respectively (excluding nonsense
749 variants).

750 Many variants were observed at high read depth in the GFP- bin of the enrichment sort but not in
751 the binned sort. We assigned these a null phenotype (lower-bound fluorescence) using a
752 statistical procedure based on read depth (see Supplementary Methods), resulting in 859,171
753 AncSR1 and 638,762 AncSR2 protein:RE null complexes (FDR = 0.1; Extended Data Fig. 2d).
754 This increased the total phenotyped variants to 1,487,903 in AncSR1 and 1,297,237 in AncSR2,
755 covering 58% and 51% of all possible variants, respectively.

756 To evaluate the accuracy of the sort-seq fluorescence values, we measured the fluorescence of 5
757 isogenic variants by flow cytometry, which were also spiked into the DMS libraries prior to the
758 binned sort. We found a high correlation between the fluorescence estimates from flow
759 cytometry and sorting (Pearson's $r^2 = 0.87$, Extended Data Fig. 2e). We additionally compared
760 the fluorescence estimates of the same variants that were contained in the DMS libraries and
761 again observed a strong correlation with flow cytometry measurements (Pearson's $r^2 = 0.97$,
762 Extended Data Fig. 2e).

763 To evaluate whether the REBC mutations affected fluorescence, we constructed AncSR1 and
764 AncSR2 “mini-libraries” consisting of each of the 16 REBCs engineered into the respective
765 wild-type protein variant. These were transformed via electroporation into the ERE or SRE
766 reporter strain, respectively, at 1:16 the scale of the full libraries, and spiked into the full-scale
767 libraries before sorting. The fluorescence of the mini-library variants did not differ significantly
768 by REBC ($p = 0.98$ AncSR1, $p = 0.99$ AncSR2, one-way ANOVA), indicating that fluorescence
769 estimates are directly comparable between libraries with different REBC mutations.

770 *Fluorescence inference for missing complexes*

771 To predict the fluorescence of the remaining complexes for which we did not obtain
772 experimental estimates, we fit a generalized linear model based on reference-free analysis
773 (RFA)^{36,37} to the experimental data. The model estimates a sigmoid function to capture the
774 measurement bounds of the assay, plus additive and interaction effects (specific epistasis) for all
775 amino acid states at the four variable sites in the DBD and all nucleotide states at the two
776 variable sites in the RE. All possible intramolecular interactions up to third order amino acid
777 interactions in the DBD and second order nucleotide interactions in the RE, and intermolecular
778 interactions up to third order amino acid-by-second order nucleotide interactions were included.
779 L2 regularization with 10-fold cross validation was used to reduce overfitting (Extended Data
780 Fig. 3a; Supplementary Methods). We fit separate RFA models for each ancestral background

781 using the *glmnet* v4.1-6 R package⁷⁷. Model fits to the observed data were $R^2 = 0.96$ for AncSR1
782 active complexes (0.31 all complexes) and $R^2 = 0.99$ for AncSR2 active complexes (0.88 all
783 complexes) (Extended Data Fig. 3b). These models were used to predict fluorescence values for
784 unobserved protein-RE complexes. We also used the fitted models to correct the predictions for
785 complexes in one of the modified strains that had systematically lower fluorescence
786 (Supplementary Methods; Extended Data Fig. 3c, d).

787 ***Classification of functional complexes***

788 We classified complexes as functional if their fluorescence was not significantly lower than the
789 wild type complex, *i.e.* EGKA:ERE in the AncSR1 background and GSKV:SRE in the AncSR2
790 background. Complexes inferred as null from the enrichment sort were classified as
791 nonfunctional. For complexes observed in the binned sort, we used a *t*-test to account for
792 measurement error. For complexes with predicted fluorescence from the RFA models, we
793 performed a nonparametric bootstrap test using the distribution of model residuals concatenated
794 over the ten cross-validation fits to account for model prediction error (Supplementary Methods;
795 Extended Data Fig. 3e). For both tests, we used a Benjamini-Hochberg FDR threshold of 0.25 to
796 classify variants as nonfunctional if they were significantly less fluorescent than the wild type
797 complex (Extended Data Fig. 3f). The low stringency of the FDR threshold was chosen to reduce
798 the false positive rate for calling variants functional. The majority of complexes classified as
799 functional in both backgrounds had fluorescence estimates obtained from the binned sort
800 experiment (59.3% AncSR1, 75.4% AncSR2; Extended Data Fig. 3g).

801 ***Protein genotype networks***

802 Following Maynard Smith's sequence space formalism⁴², we built genotype networks consisting
803 of all functional RH variants in each DBD background. RH genotypes are connected by an edge
804 if they differ by a single amino acid mutation that can be produced via a single nucleotide
805 mutation given the standard genetic code. Genotype networks for joint protein-DNA models
806 follow a similar logic (Supplementary Methods). We used the R package *igraph* v1.5.1⁷⁸ to build
807 and analyze the genotype networks, and the software *gephi* v0.10.1⁷⁹ for network visualization.
808 To identify clusters of densely connected genotypes within the networks, we used the
809 `cluster_edge_betweenness` function from the R *igraph* package.

810 ***Model of evolution on GP maps***

811 We modeled evolution on the genotype networks as an origin-fixation process under a strong
812 selection-weak mutation regime^{80,81}. To isolate the effect of the GP map's structure on evolution,
813 we considered a scenario in which all functional genotypes have equal fitness, so the fixation
814 probability is affected only by drift, and nonfunctional variants are removed by purifying
815 selection. The relative probability $P(i,j)$ of substitution from protein genotype *i* to genotype *j* is
816 therefore equal to the amino acid mutation rate μ_{ij} , normalized over all single-step neighbors of *i*
817 in the network. We assumed that there are no biases in the nucleotide mutation process (*e.g.*
818 transition vs. transversion rate), so μ_{ij} is affected only by unequal mutational access between
819 amino acids imposed by the genetic code. To incorporate this effect, we scaled μ_{ij} by the number
820 of possible nucleotide mutations that can change any nucleotide sequence that encodes *i* to any
821 nucleotide sequence that encodes *j*:

$$\mu_{ij} = \eta_{ij}^{o^*} \times \prod_{o \neq o^*} c_o \quad (1)$$

823 where o indexes the amino acid position, o^* is the position at which the amino acid change
824 occurs, $\eta_{ij}^{o^*}$ is the number of possible single nucleotide changes that can produce the state in j
825 from the state in i at site o^* , and c_o is the number of possible codons for the invariant amino acid
826 state at site o .

827 We used these transition probabilities to specify a discrete time Markov model for each ancestral
828 genotype network, where each step is a single amino acid substitution. Genotypes that are more
829 than one nucleotide change apart cannot access each other in a single time step, and the
830 probability of staying in the same genotype across a single step in the Markov chain is also zero.
831 We only considered functional genotypes within the main component of each network (the
832 largest connected component). With this model, we computed the probability distribution $\pi_{(k)}$ of
833 evolving all possible genotypes after k substitution steps given any specified set of starting
834 genotypes:

$$\pi_{(k)} = \pi_{(0)} \times P^k \quad (2)$$

836 where P is the transition matrix with entries $P(i, j)$, $k > 0$, and $\pi_{(0)}$ is the vector of the probability
837 distribution of genotypes at time step $k = 0$. Setting a single element i of $\pi_{(0)}$ to 1 and all others to
838 zero corresponds to evolution from a single starting genotype; setting all elements of $\pi_{(0)}$ to $1/n$,
839 where n is the number of functional genotypes in the network, averages over all possible starting
840 genotypes. We calculated the relative probability of evolving a given specificity phenotype at
841 time step k by summing over all elements of $\pi_{(k)}$ that encode that specificity and normalizing by
842 the total probability across all specific protein genotypes.

843 *Effects of background substitutions*

844 To estimate the effect of the background substitutions between AncSR1 and AncSR2 on binding
845 affinity, we first considered a model where the background substitutions have a universal
846 nonspecific effect on affinity across all RH and RE genotypes. We assumed that fluorescence is
847 proportional to the fraction of protein bound to DNA. If a complex g has dissociation constant
848 $K_d(g)$ in the AncSR1 background, then its AncSR1 fluorescence (normalized to scale between 0
849 and 1) is:

$$850 \quad F(g)_{AncSR1} = \frac{1}{1 + \frac{K_d(g)}{[RE]}} \quad (3)$$

851 where $[RE]$ is the concentration of RE. If the background substitutions scale $K_d(g)$ by a factor α ,
852 then fluorescence in the AncSR2 background is

$$853 \quad F(g)_{AncSR2} = \frac{1}{1 + \alpha \left(\frac{K_d(g)}{[RE]} \right)} \quad (4)$$

854 Rearranging these equations gives an expression for fluorescence in the AncSR2 background as
855 a function of fluorescence in the AncSR1 background and α :

$$856 \quad F(g)_{AncSR2} = \frac{1}{1 + \alpha \left(\frac{1 - F(g)_{AncSR1}}{F(g)_{AncSR1}} \right)} \quad (5)$$

857 We fit this model to the AncSR1 and AncSR2 fluorescence data using orthogonal regression,
858 which accounts for measurement error in both backgrounds. We used only complexes that had
859 fluorescence measurements from the binned sort in both backgrounds, and whose fluorescence
860 was significantly greater than that of nonsense variants in either background ($n = 2,627$).
861 Fluorescence was normalized in each background to scale between the upper and lower bounds
862 inferred from the RFA models. Confidence intervals (CI) were constructed by bootstrapping the
863 data and refitting the model. The effect of the background substitutions was estimated to be $\alpha =$
864 0.014 (95% CI: 0.010–0.014), corresponding to a 70-fold increase in affinity (95% CI: 70–99).

865 We next considered a model where the background substitutions have a different effect on ERE
866 affinity than they do on other REs. We modified the model such that α_1 represents the ERE-
867 specific effect of the background substitutions and α_2 the effect on the other 15 REs. We fit this
868 model as before and obtained parameter estimates of $\alpha_1 = 0.43$ (95% CI: 0.19–0.76) and $\alpha_2 =$
869 0.010 (95% CI: 0.0028–0.010), corresponding to fold-increases in affinity of 2.3 (95% CI: 1.3–
870 5.2) on ERE and 99 (95% CI: 99–361) on other REs.

871 ***Code availability***

872 Scripts for analysis are available at www.github.com/JoeThorntonLab/RH-RE_scanning.

873 References

- 874 1. Wagner, G. P. & Altenberg, L. Perspective : Complex Adaptations and the Evolution of Evolvability.
875 *Evolution* **50**, 967–976 (1996).
- 876 2. Stoltzfus, A. & Yampolsky, L. Y. Climbing Mount Probable: Mutation as a Cause of Nonrandomness
877 in Evolution. *J. Hered.* **100**, 637–647 (2009).
- 878 3. Hodgins-Davis, A., Duveau, F., Walker, E. A. & Wittkopp, P. J. Empirical measures of mutational
879 effects define neutral models of regulatory evolution in *Saccharomyces cerevisiae*. *Proc. Natl. Acad.*
880 *Sci.* **116**, 21085–21093 (2019).
- 881 4. Gould, S. J. & Lewontin, R. C. The Spandrels of San Marco and the Panglossian Paradigm: A
882 Critique of the Adaptationist Programme. *Proc. R. Soc. B Biol. Sci.* **205**, 581–598 (1979).
- 883 5. Schluter, D. Adaptive radiation along genetic lines of least resistance. *Evolution* **50**, 1766–1774
884 (1996).
- 885 6. Arthur, W. The interaction between developmental bias and natural selection: From centipede
886 segments to a general hypothesis. *Heredity* **89**, 239–246 (2002).
- 887 7. Maynard-Smith, J. *et al.* Developmental constraints and evolution. *Q. Rev. Biol.* **60**, 265–287 (1985).
- 888 8. Wake, D. B. & Larson, A. Multidimensional analysis of an evolving lineage. *Science* **238**, 42–48
889 (1987).
- 890 9. Fay, J. C. & Wittkopp, P. J. Evaluating the role of natural selection in the evolution of gene
891 regulation. *Heredity* **100**, 191–199 (2008).
- 892 10. Dugand, R. J., Aguirre, J. D., Hine, E., Blows, M. W. & McGuigan, K. The contribution of mutation
893 and selection to multivariate quantitative genetic variance in an outbred population of *Drosophila*
894 *serrata*. *Proc. Natl. Acad. Sci.* **118**, e2026217118 (2021).
- 895 11. McGlothlin, J. W. *et al.* Adaptive radiation along a deeply conserved genetic line of least resistance
896 in *Anolis* lizards. *Evol. Lett.* 310–322 (2018) doi:10.1002/evl3.72.
- 897 12. McKeown, A. N. *et al.* Evolution of DNA specificity in a transcription factor family produced a new
898 gene regulatory module. *Cell* **159**, 58–68 (2014).
- 899 13. Raup, D. M. Geometric analysis of shell coiling: General Problems. *J. Paleontol.* **40**, 1178–1190
900 (1966).
- 901 14. Vermeij, G. J. Forbidden phenotypes and the limits of evolution. *Interface Focus* **5**, 20150028
902 (2015).
- 903 15. Deline, B. *et al.* Evolution of metazoan morphological disparity. *Proc Nat Acad Sci* E8909–E8918
904 (2018) doi:10.1073/pnas.1810575115.
- 905 16. Clark, J. W. *et al.* Evolution of phenotypic disparity in the plant kingdom. *Nat. Plants* (2023)
906 doi:10.1038/s41477-023-01513-x.
- 907 17. Dawkins, R. *Climbing Mount Improbable*. (WW Norton & Company., 1996).
- 908 18. Grant, P. R. & Grant, B. R. *40 Years of Evolution: Darwin's Finches on Daphne Major Island*.
909 (Princeton university press, Princeton, New Jersey, 2014). doi:10.5860/choice.52-0821.
- 910 19. Jablonski, D. Developmental bias, macroevolution, and the fossil record. *Evol. Dev.* 103–125 (2019)
911 doi:10.1111/ede.12313.
- 912 20. Stepan, S. J., Phillips, P. C. & Houle, D. Comparative quantitative genetics: evolution of the
913 Gmatrix. *Trends Ecol Evol* **17**, 320–327 (2002).
- 914 21. Fowler, D. M., Stephany, J. J. & Fields, S. Measuring the activity of protein variants on a large scale
915 using deep mutational scanning. *Nat. Protoc.* **9**, 2267–2284 (2014).
- 916 22. Podgornaia, A. I. & Laub, M. T. Pervasive degeneracy and epistasis in a protein-protein interface.
917 *Science* **347**, 673–677 (2015).
- 918 23. Wu, N. C., Dai, L., Olson, C. A., Lloyd-Smith, J. O. & Sun, R. Adaptation in protein fitness
919 landscapes is facilitated by indirect paths. *eLife* **5**, 1–21 (2016).
- 920 24. Starr, T. N., Picton, L. K. & Thornton, J. W. Alternative evolutionary histories in the sequence space
921 of an ancient protein. *Nature* **549**, 409–413 (2017).
- 922 25. Lite, T.-L. V. *et al.* Uncovering the basis of protein-protein interaction specificity with a
923 combinatorially complete library. *eLife* **9**, e60924 (2020).

- 924 26. Kemble, H., Nghe, P. & Tenaillon, O. Recent insights into the genotype–phenotype relationship from
925 massively parallel genetic assays. *Evol. Appl.* **12**, 1721–1742 (2019).
- 926 27. Bulyk, M. L., Huang, X., Choo, Y. & Church, G. M. Exploring the DNA-binding specificities of zinc
927 fingers with DNA microarrays. *Proc. Natl. Acad. Sci.* **98**, 7158–7163 (2001).
- 928 28. Newburger, D. E. & Bulyk, M. L. UniPROBE: an online database of protein binding microarray data
929 on protein-DNA interactions. *Nucleic Acids Res.* **37**, D77–D82 (2009).
- 930 29. Anderson, D. W., McKeown, A. N. & Thornton, J. W. Intermolecular epistasis shaped the function
931 and evolution of an ancient transcription factor and its DNA binding sites. *eLife* **4**, e07864–e07864
932 (2015).
- 933 30. Wheeler, L. C. & Harms, M. J. Were Ancestral Proteins Less Specific? *Mol. Biol. Evol.* **38**, 2227–
934 2239 (2021).
- 935 31. Patwardhan, R. P. *et al.* High-resolution analysis of DNA regulatory elements by synthetic saturation
936 mutagenesis. *Nat. Biotechnol.* **27**, 1173–1175 (2009).
- 937 32. Hochberg, G. K. A. & Thornton, J. W. Reconstructing Ancient Proteins to Understand the Causes of
938 Structure and Function. *Annu. Rev. Biophys.* **46**, 247–269 (2017).
- 939 33. Park, Y., Metzger, B. P. H. & Thornton, J. W. Epistatic drift causes gradual decay of predictability in
940 protein evolution. *Science* **376**, 823–830 (2022).
- 941 34. Carroll, J. S. *et al.* Genome-wide analysis of estrogen receptor binding sites. *Nat. Genet.* **38**, 1289–
942 1297 (2006).
- 943 35. Watson, L. C. *et al.* The glucocorticoid receptor dimer interface allosterically transmits sequence-
944 specific DNA signals. *Nat. Struct. Mol. Biol.* **20**, 876–883 (2013).
- 945 36. Park, Y., Metzger, B. P. H. & Thornton, J. W. The simplicity of protein sequence-function
946 relationships. *Nat. Commun.* **15**, 7953 (2024).
- 947 37. Metzger, B. P. H., Park, Y., Starr, T. N. & Thornton, J. W. Epistasis facilitates functional evolution in
948 an ancient transcription factor. *eLife* **12**, (2024).
- 949 38. Gerber, S. Not all roads can be taken: Development induces anisotropic accessibility in morphospace.
950 *Evol. Dev.* **16**, 373–381 (2014).
- 951 39. Stadler, B. M. R., Stadler, P. F., Wagner, G. P. & Fontana, W. The Topology of the Possible: Formal
952 Spaces Underlying Patterns of Evolutionary Change. *J. Theor. Biol.* **213**, 241–274 (2001).
- 953 40. Psujek, S. & Beer, R. D. Developmental bias in evolution: evolutionary accessibility of phenotypes in
954 a model evo-devo system. *Evol. Dev.* **10**, 375–390 (2008).
- 955 41. Salazar-Ciudad, I. Why call it developmental bias when it is just development? *Biol. Direct* **16**, 1–13
956 (2021).
- 957 42. Maynard-Smith, J. Natural Selection and the Concept of a Protein Space. *Nature* **225**, 726–734
958 (1970).
- 959 43. Newman, M. E. J. & Girvan, M. Finding and evaluating community structure in networks. *Phys. Rev.*
960 *E* **69**, 026113 (2004).
- 961 44. Kimura, M. *The Neutral Theory of Molecular Evolution*. (Cambridge University Press, 1983).
962 doi:10.1016/B978-1-55938-802-3.50013-4.
- 963 45. Russo, F. & Williamson, J. Interpreting Causality in the Health Sciences. *Int. Stud. Philos. Sci.* **21**,
964 157–170 (2007).
- 965 46. Schuster, P., Fontana, W., Stadler, P. F. & Hofacker, I. L. From sequences to shapes and back: A case
966 study in RNA secondary structures. *Proc R Soc Lond B* **255**, 279–284 (1994).
- 967 47. Fontana, W. & Schuster, P. Shaping space: The possible and the attainable in RNA genotype-
968 phenotype mapping. *J Theor Biol* **194**, 491–515 (1998).
- 969 48. Alberch, P. Ontogenesis and Morphological Diversification. *Am. Zool.* **20**, 653–667 (1980).
- 970 49. Chipman, A. D., Arthur, W. & Akam, M. A Double Segment Periodicity Underlies Segment
971 Generation in Centipede Development. *Curr. Biol.* **14**, 1250–1255 (2004).
- 972 50. Salazar-Ciudad, I. & Jernvall, J. A computational model of teeth and the developmental origins of
973 morphological variation. *Nature* **464**, 583–586 (2010).
- 974 51. Fuqua, T. *et al.* Dense and pleiotropic regulatory information in a developmental enhancer. *Nature*
975 **587**, 235–239 (2020).

- 976 52. Arthur, W. & Farrow, M. The Pattern of Variation in Centipede Segment Number as an Example of
977 Developmental Constraint in Evolution. *J. Theor. Biol.* **200**, 183–191 (1999).
- 978 53. Harjunmaa, E. *et al.* Replaying evolutionary transitions from the dental fossil record. *Nature* **512**, 44–
979 48 (2014).
- 980 54. Dingle, K., Ghaddar, F., Šulc, P. & Louis, A. A. Phenotype Bias Determines How Natural RNA
981 Structures Occupy the Morphospace of All Possible Shapes. *Mol. Biol. Evol.* **39**, 1–11 (2022).
- 982 55. Rohner, P. T. & Berger, D. Developmental bias predicts 60 million years of wing shape evolution.
983 *Proc. Natl. Acad. Sci.* **120**, e2211210120 (2023).
- 984 56. Galupa, R. *et al.* Enhancer architecture and chromatin accessibility constrain phenotypic space during
985 *Drosophila* development. *Dev. Cell* **58**, 51–62.e4 (2023).
- 986 57. Ferrada, E. & Wagner, A. Evolutionary innovations and the organization of protein functions in
987 genotype space. *PLoS ONE* **5**, (2010).
- 988 58. Ciliberti, S., Martin, O. C. & Wagner, A. Innovation and robustness in complex regulatory gene
989 networks. *Proc. Natl. Acad. Sci. U. S. A.* **104**, 13591–13596 (2007).
- 990 59. Matias Rodrigues, J. F. & Wagner, A. Evolutionary Plasticity and Innovations in Complex Metabolic
991 Reaction Networks. *PLoS Comput. Biol.* **5**, e1000613 (2009).
- 992 60. Gould, S. J. & Eldredge, N. Punctuated Equilibria : The Tempo and Mode of Evolution
993 Reconsidered. *Paleobiology* **3**, 115–151 (1977).
- 994 61. Alberch, P. & Gale, E. A. A developmental analysis of an evolutionary trend: digital reduction in
995 amphibians. *Evolution* **39**, 8–23 (1985).
- 996 62. Braendle, C., Baer, C. F. & Félix, M.-A. Bias and Evolution of the Mutationally Accessible
997 Phenotypic Space in a Developmental System. *PLoS Genet.* **6**, e1000877 (2010).
- 998 63. Phillips, A. M. *et al.* Binding affinity landscapes constrain the evolution of broadly neutralizing anti-
999 influenza antibodies. *eLife* **10**, 1–40 (2021).
- 1000 64. Starr, T. N. *et al.* ACE2 binding is an ancestral and evolvable trait of sarbecoviruses. *Nature* **603**,
1001 913–918 (2022).
- 1002 65. Ord, T. J. & Summers, T. C. Repeated evolution and the impact of evolutionary history on
1003 adaptation. *BMC Evol. Biol.* **15**, 137 (2015).
- 1004 66. Aakre, C. D. *et al.* Evolving New Protein-Protein Interaction Specificity through Promiscuous
1005 Intermediates. *Cell* **163**, 594–606 (2015).
- 1006 67. Poelwijk, F. J., Kiviet, D. J., Weinreich, D. M. & Tans, S. J. Empirical fitness landscapes reveal
1007 accessible evolutionary paths. *Nature* **445**, 383–386 (2007).
- 1008 68. Anderson, D. W., Baier, F., Yang, G. & Tokuriki, N. The adaptive landscape of a metallo-enzyme is
1009 shaped by environment-dependent epistasis. *Nat. Commun.* **12**, 3867 (2021).
- 1010 69. Lewontin, R. C. Four complications in understanding the evolutionary process. in *SFI Bulletin* vol. 18
1011 (2003).
- 1012 70. Maclean, C. J. *et al.* Deciphering the Genic Basis of Yeast Fitness Variation by Simultaneous
1013 Forward and Reverse Genetics. *Mol. Biol. Evol.* **34**, 2486–2502 (2017).
- 1014 71. R.D. Gietz & R.A. Woods. Yeast Transformation by the LiAc/SS Carrier DNA/PEG Method. in
1015 *Yeast Protocol*, W. Xiao, Ed. 107–120 (Humana Press, Totowa, NJ, 2006).
- 1016 72. Scanlon, T. C., Gray, E. C. & Griswold, K. E. Quantifying and resolving multiple vector
1017 transformants in *S. cerevisiae* plasmid libraries. *BMC Biotechnol.* **9**, 95 (2009).
- 1018 73. Mir, K., Neuhaus, K., Bossert, M. & Schober, S. Short Barcodes for Next Generation Sequencing.
1019 *PLOS ONE* **8**, e82933 (2013).
- 1020 74. N.A. Joshi & J.N. Fass. Sickle: A sliding-window, adaptive, quality-based trimming tool for FastQ
1021 files. (2011).
- 1022 75. J. Zhang, K. Kobert, T. Flouri, & A. Stamatakis. PEAR: A fast and accurate Illumina Paired-End
1023 reAd mergeR. (2015).
- 1024 76. Cock, P. J. A. *et al.* Biopython: Freely available Python tools for computational molecular biology
1025 and bioinformatics. *Bioinformatics* **25**, 1422–1423 (2009).
- 1026 77. Jerome Friedman *et al.* glmnet: Lasso and Elastic-Net Regularized Generalized Linear Models.

- 1027 78. Csardi, G. & Nepusz, T. The igraph software package for complex network research. *InterJournal*
1028 **Complex Systems**, 1695 (2006).
- 1029 79. Bastian, M., Heymann, S., & Jacomy, M. Gephi: an open source software for exploring and
1030 manipulating networks. *Proc. Int. AAAI Conf. Web Soc. Media* **3**, 361–362 (2009).
- 1031 80. Gillespie, J. Molecular Evolution Over the Mutational Landscape. *Evolution* **38**, 1116–1129 (1984).
- 1032 81. Mccandlish, D. M. & Stoltzfus, A. Modeling Evolution Using the Probability of Fixation: History
1033 and Implications. *Q. Rev. Biol.* **89**, 225–252 (2014).

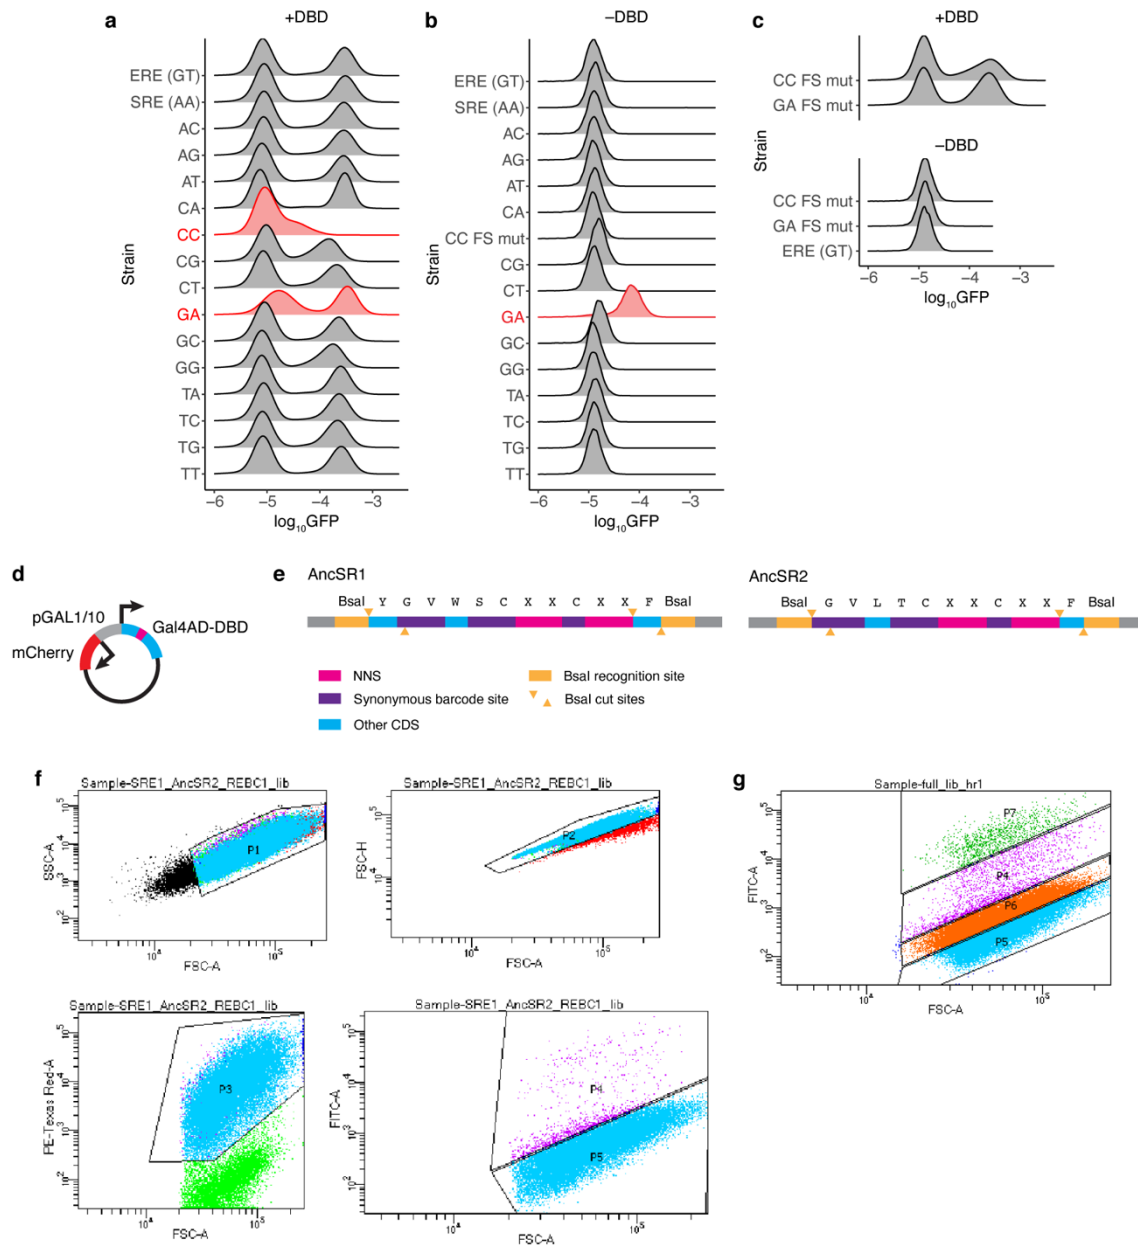
1034

1035 **Acknowledgements:** We thank Yeonwoo Park and Brian Metzger for advice throughout the
1036 project and all Thornton Lab members for comments on the manuscript. This work was
1037 supported by the University of Chicago's Research Computing Center, Cytometry Core, and
1038 Genomics Core. Funding was provided by the National Institutes of Health grants
1039 R01GM131128 (J.W.T.), R01GM121931 (J.W.T.), R35GM14533601 (J.W.T.), an NSF
1040 Graduate Research Fellowship (J.E.J.P.), and a Rosemary Grant Award from the Society for the
1041 Study of Evolution (S.H.A.).

1042 **Author contributions:** All authors conceived the project; S.H.A. and J.E.J.P. performed
1043 experiments and analyzed data; all authors interpreted results and wrote the manuscript.

1044 **Competing interests:** The authors declare that they have no competing interests.

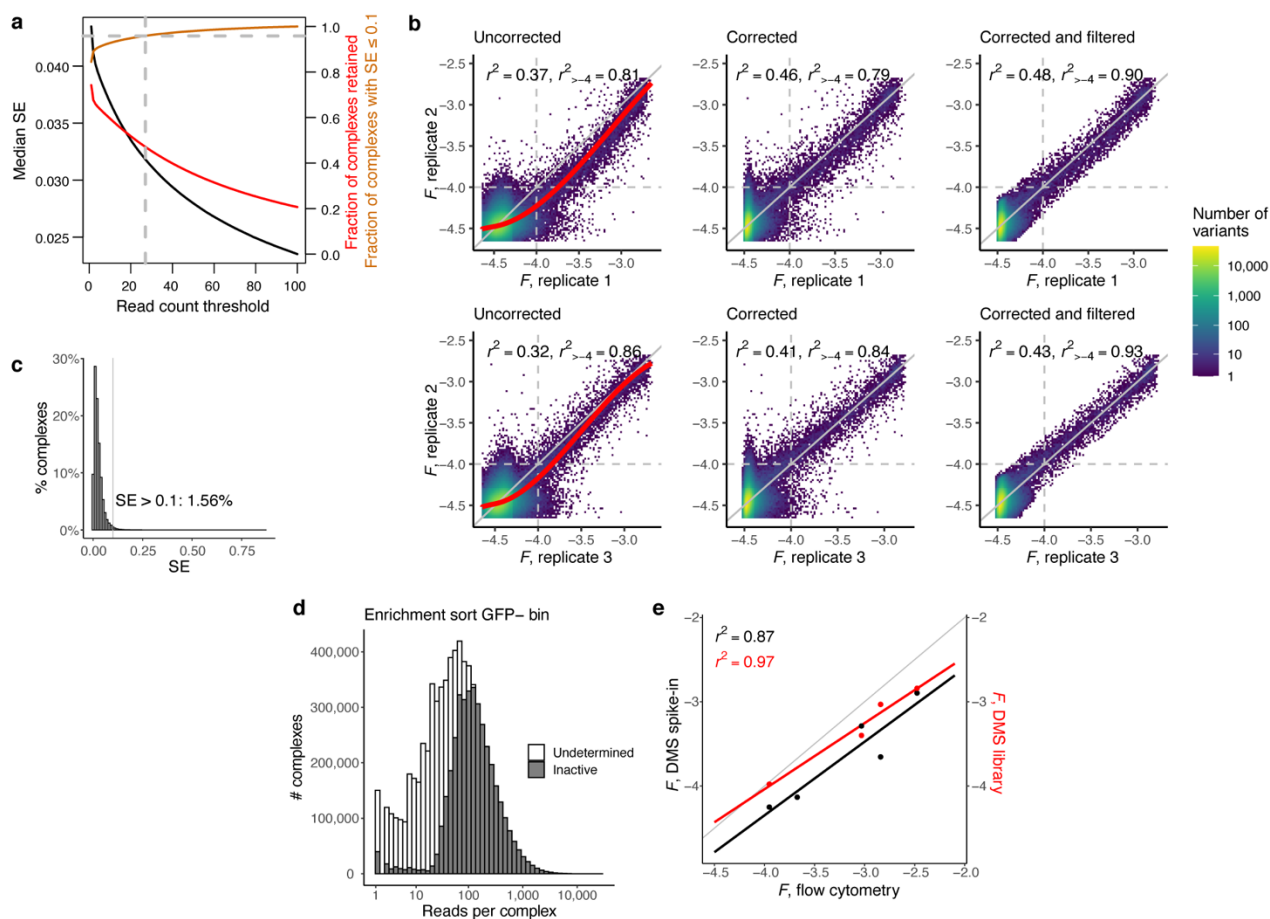
1045 **Additional information:** Supplementary information is available for this paper. Correspondence
1046 and material requests should be addressed to J.W.T (joet1@uchicago.edu).



1047

1048 **Extended Data Figure 1 | DBD library construction and sorting.** **a**, Design of the DBD
 1049 expression vector used for DMS. The SR DBD is fused to an N-terminal *S. cerevisiae* Gal4
 1050 Activation Domain. Its expression is under control of a bidirectional pGAL1/10 promoter, which
 1051 simultaneously drives mCherry expression to select cells that maintain the plasmid during
 1052 sorting. **b**, Design of DBD library oligos. NNS codons (pink) were used to generate all possible
 1053 combinations of amino acid mutations at the four RH scanning sites (marked as X in the amino
 1054 acid sequence). For each background (AncSR1, left; AncSR2, right), we synthesized 16 libraries,
 1055 each with a unique set of synonymous barcode mutations at five codons (purple, Supplementary
 1056 Table 1), which allows each to be associated with one RE strain. BsaI sites (orange) were used
 1057 for Golden Gate assembly into the pDBD2.1 backbone. **c–e**, Validation of the RE reporter
 1058 strains. GFP fluorescence was measured by flow cytometry in each strain in the presence
 1059 (+DBD) or absence (–DBD) of a universally high-affinity DBD variant (AncSR1+GGKA+11P,
 1060 ²⁹). In each row, the left peak corresponds to autofluorescence from cells that do not express

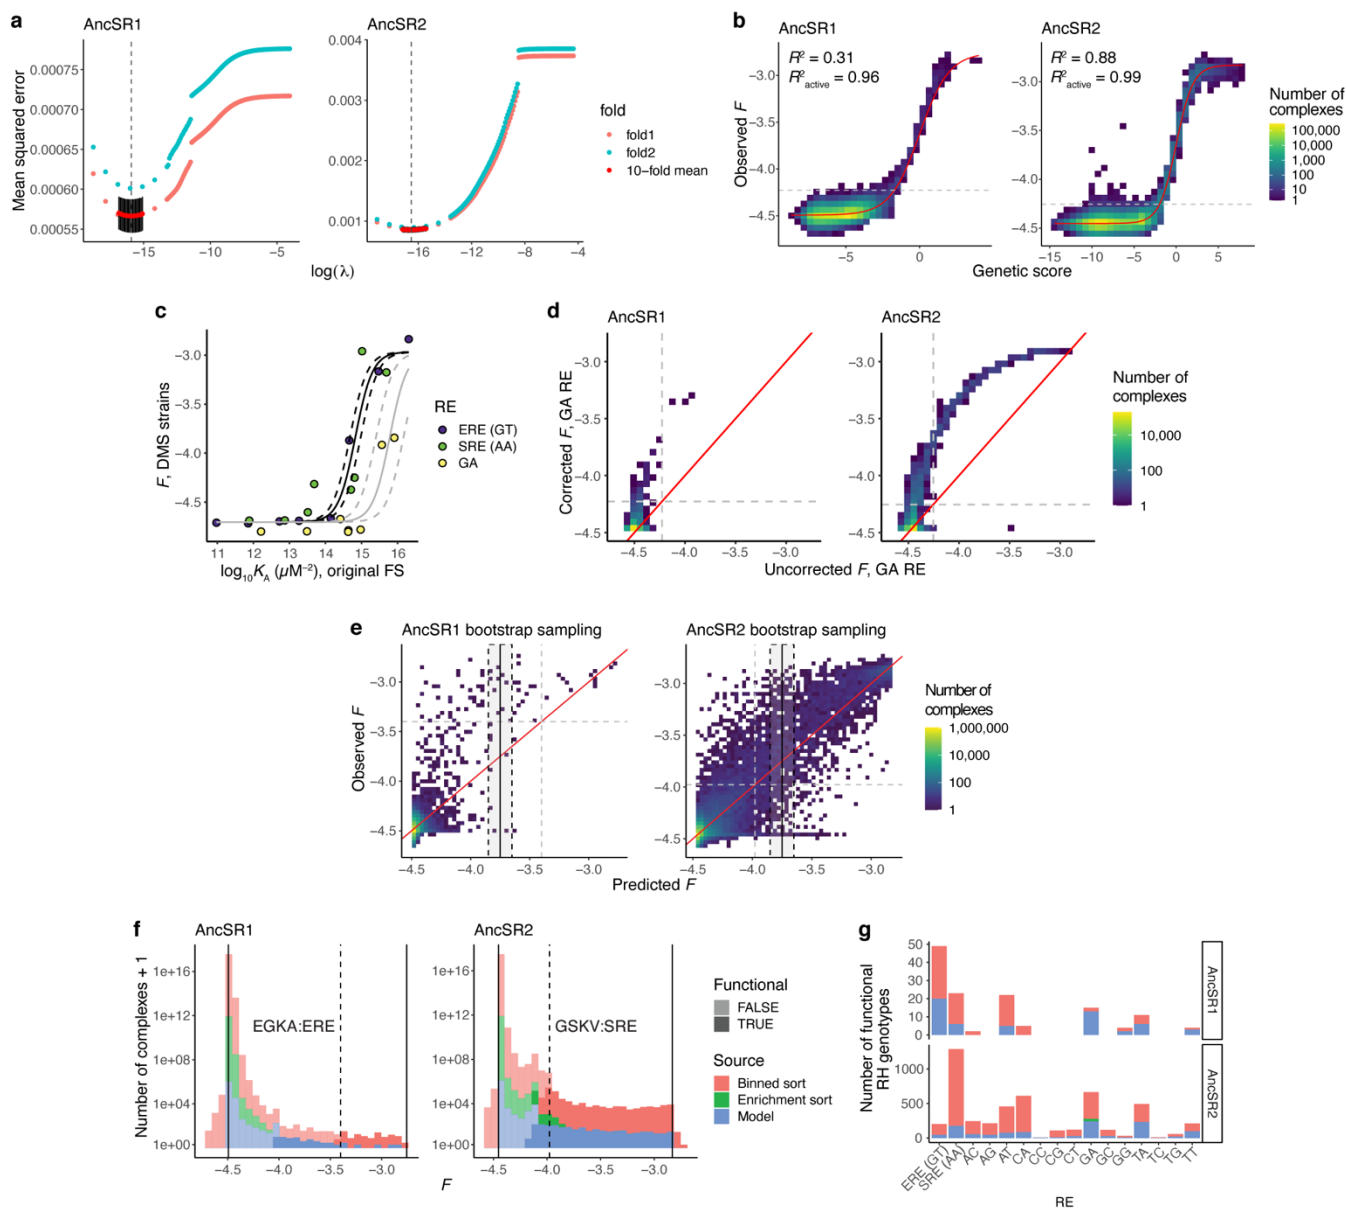
1061 GFP, either due to lack of DBD binding or loss of the DBD expression plasmid; the right peak
1062 corresponds to cells that are expressing GFP in response to DBD-RE binding. “FS mut” denotes
1063 strains with mutations in the flank/spacer regions of the RE that correct anomalous expression
1064 patterns shown in red (see Supplementary Methods). Red strains were not used in the final DMS
1065 experiment. Experiments were conducted on the same day within each panel. **c**, Fluorescence in
1066 the presence of high-affinity DBD. **d**, Fluorescence in the absence of DBD expression plasmid.
1067 **e**, Fluorescence in the CC and GA FS mut strains, with the ERE strain included as a negative
1068 control. **f–g**, Sorting gates used for DMS. **f**, Enrichment sort gates. Homogeneous single cells
1069 were first selected by gating on FSC-A vs. SSC-A and FSC-A vs. FSC-H (top). Plasmid
1070 retention was then selected for by gating on mCherry expression (PE-Texas Red-A, bottom left).
1071 Finally, cells were sorted into GFP+ (P4) and GFP– (P5) populations (bottom right). The
1072 boundary between the GFP+ and GFP– gates was drawn to have a slope of 1.5 on a log-FSC-A
1073 vs. log-GFP (FITC) scale so that populations were sorted by GFP expression relative to cell
1074 volume. **g**, Binned sort gates. Gates P1–P3 were drawn as in **C**. Cells were then sorted into four
1075 GFP bins, which were drawn to have roughly equal heights (P5–P7). The boundaries between
1076 GFP gates were again drawn to have a log-log slope of 1.5.
1077



1078

1079 **Extended Data Figure 2 | DMS data cleaning.** **a**, Curves show characteristics of the binned
 1080 sort dataset as a function of the read count threshold used to retain protein-RE complexes for
 1081 further analysis (x-axis). Black, standard error of F (SE, left axis); red, complexes retained,
 1082 expressed as a fraction of the number of complexes in the binned sort (right axis); gold, fraction
 1083 of complexes retained that have SE ≤ 0.1 (right axis). We used a read count threshold of 27
 1084 (vertical dashed line), at which $\geq 95\%$ of complexes have SE ≤ 0.1 (horizontal dashed line). **b**,
 1085 Correcting and filtering estimates of F from the binned sort. Left, correlation in F between
 1086 replicates before correction. Pearson's r^2 is shown for all complexes, and for the subset of
 1087 complexes with $F > -4$ in both replicates, which roughly corresponds to the boundary between
 1088 active and inactive complexes (gray dotted lines). Red curves, I-splines fit using complexes with
 1089 SE of $F < 0.1$. Center, correlation in F between replicates after correcting using the I-spline
 1090 transforms. Right, correlation in F between replicates after filtering corrected variants for SE \leq
 1091 0.1. **c**, Distribution of SE across all complexes in the binned sort after the I-spline correction.
 1092 Complexes with SE > 0.1 were discarded. **d**, Read count distribution for complexes sequenced in
 1093 the enrichment sort GFP- bin. Complexes were inferred to be inactive (gray) if they were not
 1094 observed in the binned sort, but had high enough inferred cell count in the enrichment sort to
 1095 have been detectable in the binned sort had they been at least minimally fluorescent (see
 1096 Supplementary Methods). **e**, Correlations between estimates of F from flow cytometry (x-axis)
 1097 and DMS (y-axes). Left y-axis (black points) shows estimates from isogenic strains that were
 1098 spiked into the DMS libraries prior to the binned sort. Right y-axis (red points) shows estimates
 1099 from complexes that were encoded in the DMS libraries.


1100



1101

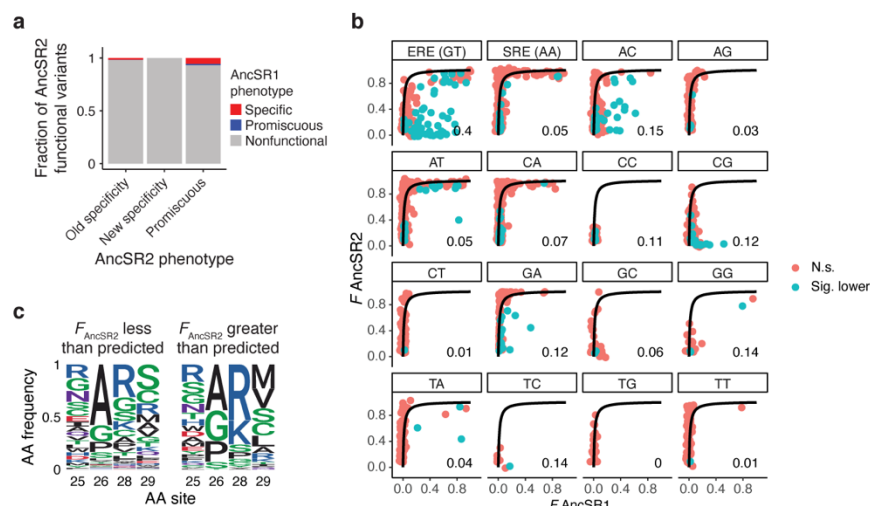
1102 **Extended Data Figure 3 | Fluorescence imputation, GA fluorescence correction, and**
 1103 **functional genotype classification.** **a, b,** A generalized linear model that predicts the
 1104 fluorescence of each protein-RE complex from its sequence was fit to the data for each
 1105 background, using L2 regularization to address overfitting. **a,** Ten-fold cross-validation (CV)
 1106 was used to identify the optimal L2 penalty parameter (λ). Red and black, mean and SE of the
 1107 out-of-sample mean squared error (MSE) across the 10 folds. Initial range finding was performed
 1108 using two folds (pink and cyan). Vertical line, λ that minimizes mean MSE. **b,** Genetic score
 1109 versus observed F for the regularized RFA models. Red line, best-fit nonspecific epistasis
 1110 function. For display, the distribution was discretized; colors show the number of variants in the
 1111 interval defined by each square. Coefficient of determination (R^2) is reported for all complexes
 1112 and for the subset of active complexes (above the gray line). **c, d,** Fluorescence correction for the
 1113 GA strain. **c,** Affinity (K_A) versus F for a panel of DBD variants measured on ERE, SRE, and
 1114 GA. Affinities, measured by fluorescence anisotropy on the three REs, all with the original
 1115 flank/spacer sequence, were previously reported^{12,29}. F was measured by flow cytometry in the

1116 RE strains that were used for DMS, of which the ERE and SRE strains had the original
1117 flank/spacer sequence, and the GA strain had a mutated flank/spacer sequence (see
1118 Supplementary Methods, Extended Data Fig. 1c–e). Curves, best-fit sigmoidal function. The
1119 same midpoint parameter was used for ERE and SRE (black); that for GA was independently
1120 estimated (gray). Dashed lines, sigmoidal functions using 95% confidence intervals on the
1121 midpoints. **d**, GA fluorescence correction based on the affinity effect estimated in **c**. Plots show
1122 F before and after the correction. Dashed gray lines, mean boundary between active and null
1123 variants. Red line, $y = x$. **e**, Bootstrap sampling strategy for classifying functional complexes with
1124 model-inferred fluorescence. Plots show concatenated out-of-sample predictions versus observed
1125 F across all 10 CV models. Bootstrap-sampled residuals from the interval within ± 0.1 units of a
1126 complex's predicted F were used to test whether a variant with model-inferred F was not
1127 significantly worse than the wild-type complex (dashed gray lines). An example for a complex
1128 with inferred $F = -3.75$ (solid black line) is shown, with the bootstrap interval shown as a shaded
1129 rectangle. Solid red line, $y = x$. **f**, Distribution of F across all 2,560,000 complexes in each DBD
1130 background. Solid vertical lines, upper and lower bounds of fluorescence inferred from the RFA
1131 models; dashed vertical lines, fluorescence of wild type complex (EGKA: ERE for AncSR1 and
1132 GSKV:SRE for AncSR2). Colors indicate the source from which F was estimated. Darker colors
1133 show functional variants, lighter colors nonfunctional. All “enrichment sort” complexes were
1134 assigned to the lower bound of fluorescence, except for GA RE variants whose fluorescence was
1135 corrected upward (**d**). Some model-predicted variants in the AncSR1 background have predicted
1136 F below the reference but are classified as functional, because the bootstrap test accounts for the
1137 AncSR1 RFA model's tendency to under-predict fluorescence (**e**, left). **g**, Bars show the number
1138 of functional RH variants per RE per DBD background, colored by source of F estimate as in **f**.
1139



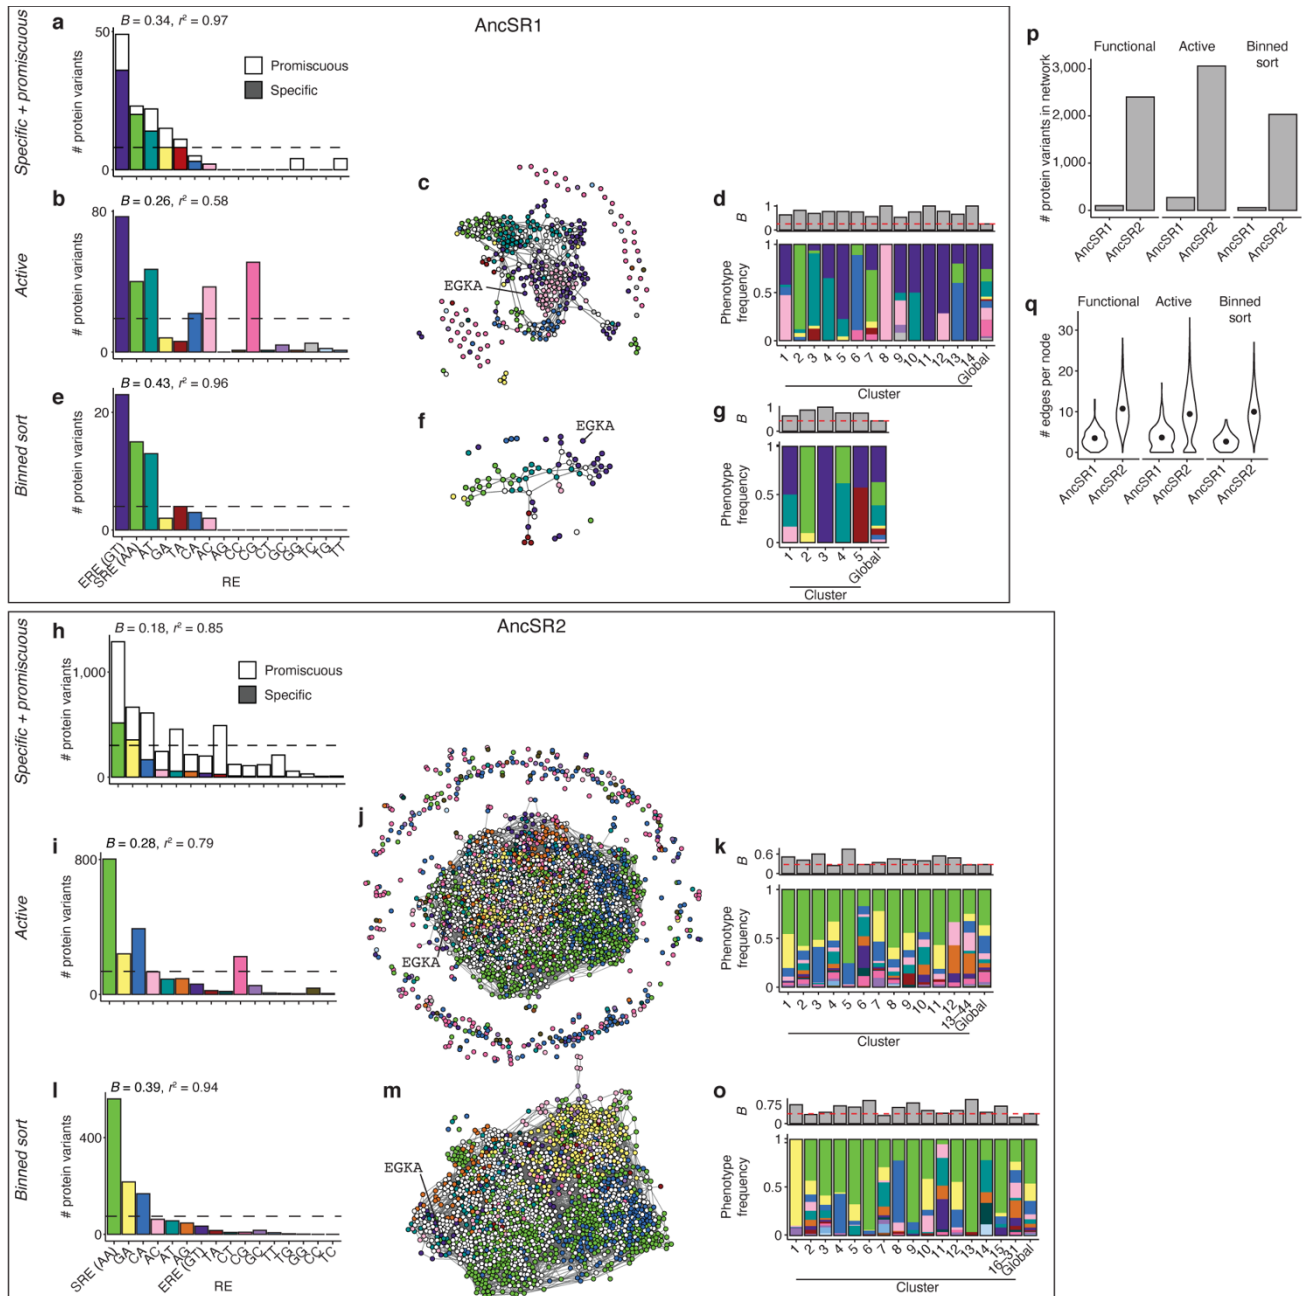
1140

1141 **Extended Data Figure 4 | Accessible new phenotypes after 3 substitution steps in the**
1142 **AncSR1 network.** Bars show the distribution over every starting genotype in the AncSR1 main
1143 component. Dashed line, mean.
1144



1145

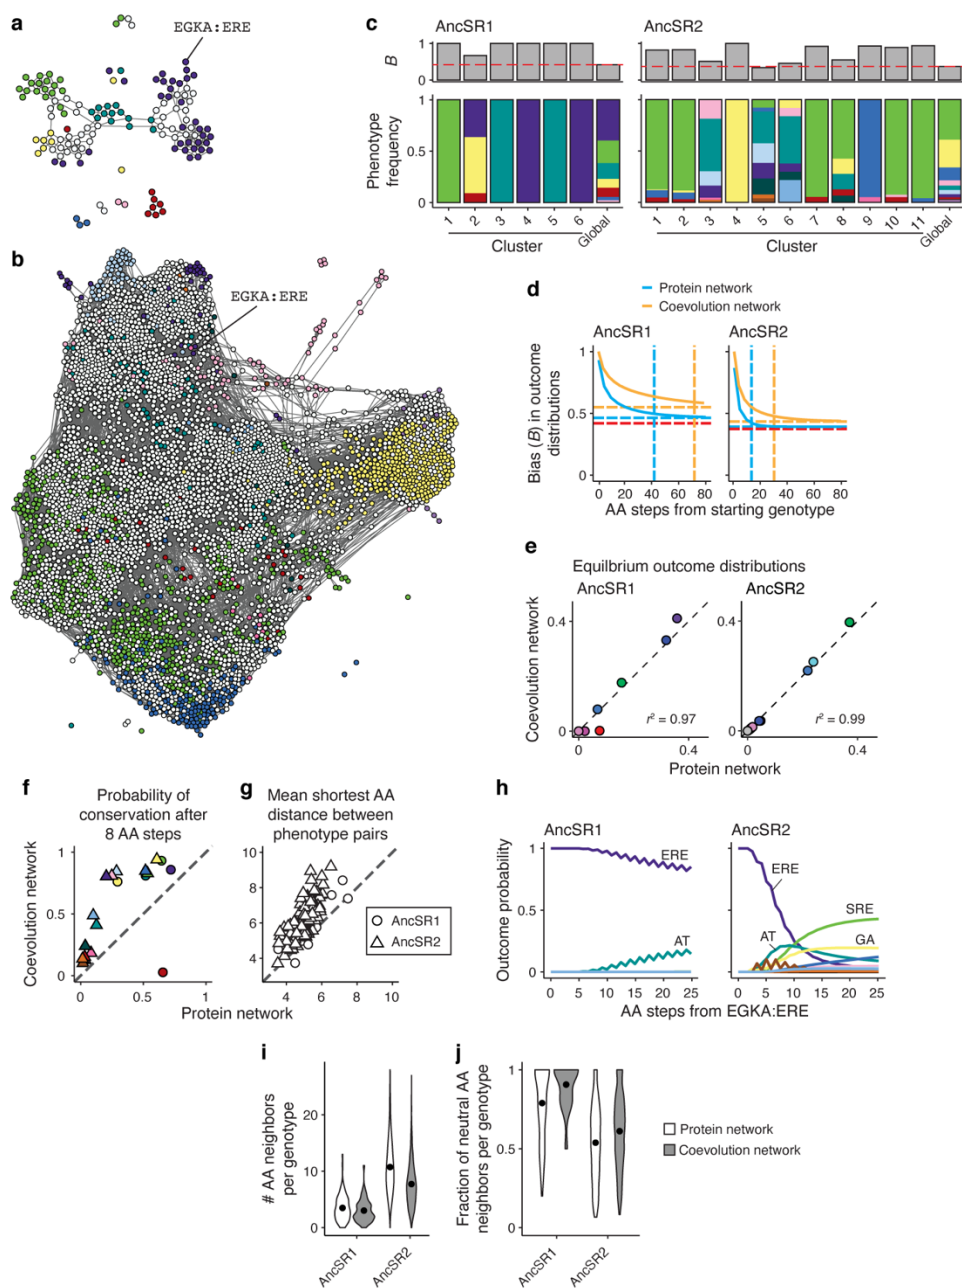
1146 **Extended Data Figure 5 | Additional analyses for effects of background substitutions on**
 1147 **DBD-RE affinity.** **a**, Changes in phenotype across the AncSR1-to-AncSR2 transition. Bars
 1148 represent the set of protein variants in AncSR2 that have different classes of phenotypes:
 1149 specificity phenotypes that were encoded in the AncSR1 map (old specificity), specificity
 1150 phenotypes not encoded in the AncSR1 map (new specificity), or promiscuous in AncSR2.
 1151 Colored sections show the fraction of variants in each class whose functional category in the
 1152 AncSR1 background was specific, promiscuous, or nonfunctional. **b**, Plots are the same as in
 1153 Fig. 6A, but split into panels by RE. Blue points, protein-DNA complexes with significantly
 1154 lower fluorescence in the AncSR2 background than predicted by the model; red, all other
 1155 variants. Numbers at the bottom-right of each panel show the fraction of plotted variants with
 1156 significantly lower than expected AncSR2 fluorescence. **c**, Amino acid frequencies at the RH
 1157 variable sites among all complexes that are significantly more (left) or less (right) fluorescent in
 1158 the AncSR2 background than predicted by the ERE-specific model in Fig. 6e. To test for
 1159 significance in **b** and **c**, we tested whether their Bonferroni-corrected 95% CI of fluorescence
 1160 was outside of the 95% CI of the model in both the AncSR1 and AncSR2 backgrounds.
 1161



1162

1163 **Extended Data Figure 6 | Robustness to alternative phenotype assignment methods. a,**
 1164 Global production distribution in the AncSR1 background, counting variants that bind
 1165 specifically (colored bars) and promiscuously (white bars) to each RE. Dashed line shows the
 1166 expected frequencies if the production distribution were isotropic. The bias, B , of the distribution
 1167 and r^2 to the production distribution for specific variants (Fig. 2a) are reported. **b,** Same as in **a**,
 1168 with phenotypes calculated using data from variants with fluorescence significantly higher than
 1169 that of nonsense variants (active variants). **c,** Sequence space network for AncSR1 active
 1170 variants. **d,** Bottom: Frequencies of specificity phenotypes within each genotype cluster in the
 1171 AncSR1 active variant networks; the global production distribution is shown for comparison.
 1172 Top: strength of phenotype bias (B) in each cluster. Red line, B of global production distribution.
 1173 **e–g,** Same as in **b–d**, but with phenotypes calculated using only data from the binned sort
 1174 experiment; protein-DNA complexes without experimental fluorescence measurements were

1175 assumed to have null fluorescence. **h–o**, Same as in **a–g**, but for the AncSR2 background. Note
1176 that the active variant datasets are likely to be enriched for false positives due to
1177 misclassification of variants whose fluorescence is by chance slightly higher than the nonsense
1178 variant distribution. This may explain the high frequency of variants that do not share any
1179 mutational connections to other active variants. It may also explain the high frequency of CG-
1180 specific variants compared to the original classification scheme, since the CG yeast strain has a
1181 slightly higher null fluorescence level than most other strains (Extended Data Fig. 1c, d) and
1182 most CG-specific variants are unconnected in the active variant genotype networks. **p**, Number
1183 of protein variants in each network under different methods of phenotype assignment.
1184 “Functional” indicates the original method used in the main text; note that this yields the same
1185 number of protein variants as the “specific + promiscuous” method. **q**, Number of edges per
1186 node in each network, with the original phenotype classification method (functional) shown for
1187 comparison.
1188



1189

1190 **Extended Data Figure 7 | Robustness to model of evolution using joint protein-DNA**

1191 **networks.** **a**, AncSR1 protein-DNA coevolution network. Nodes represent functional protein-RE

1192 complexes, colored by the RE specificity of the protein genotype; colors are as in Fig. 2b and 4b.

1193 Promiscuous protein genotypes are represented by multiple nodes, one for each RE it binds.

1194 Edges connect complexes that can be interconverted by a single nucleotide change in the RE or

1195 the coding sequence of the protein. **b**, AncSR2 protein-DNA coevolution network. **c**, Bottom:

1196 Frequencies of specificity phenotypes within each genotype cluster in the AncSR1 (left) and

1197 AncSR2 (right) coevolution networks; the global production distribution (right-most column) is


1198 shown for comparison. Top: strength of phenotype bias (B) in each cluster. Red line, B of global

1199 production distribution. **d**, Bias (B) in evolutionary outcomes as a function of the length of

1200 evolutionary trajectories. Solid curves, mean B across starting genotypes in the protein (cyan) or

1201 coevolution (orange) networks. Dashed horizontal lines, B of the equilibrium distribution in each

1202 network; dashed horizontal red line, global bias. Vertical dashed lines show the number of
1203 substitutions required for mean B to reach within 0.05 units of the equilibrium value within each
1204 type of network. The equilibrium distributions are more biased in the coevolution networks, and
1205 require more amino acid substitutions to be reached, because changes in protein genotype must
1206 occur between variants that can bind to the same RE sequence. **e**, Comparison between
1207 equilibrium outcome distributions of the protein-only evolution and protein-DNA coevolution
1208 networks in each AncSR1 (left) and AncSR2 (right) backgrounds. Pearson's r^2 between the two
1209 distributions are shown. Dashed line, $y = x$. **f**, Probability of conservation of each phenotype after
1210 8 amino acid substitution steps in the protein vs. coevolution networks. **g**, Mean shortest amino
1211 acid distance between all possible pairs of phenotypes in the coevolution vs. protein networks,
1212 calculated as in Fig. 2g. Circles, AncSR1 networks, triangles, AncSR2 networks. Dashed line, y
1213 $= x$. **h**, Probability of evolving each specificity phenotype as a function of the number of amino
1214 acid substitutions away from EGKA:ERE in the AncSR1 (left) and AncSR2 (right) coevolution
1215 networks. In both backgrounds, conservation is more likely at short trajectory lengths than in the
1216 corresponding protein networks (Fig. 3g, 5f), but the relative likelihood of achieving each
1217 phenotypic outcome is similar. **i**, Distribution of the number of neighbors per genotype with
1218 distinct RH sequences in each type of network. Dots, means of distributions. **j**, Distribution of
1219 the fraction of neutral neighbors per node with distinct RH genotypes in each network. Dots,
1220 means of distributions.
1221



1222

1223 **Extended Data Figure 8 | Robustness of RH mutation effects to uncertainty in ancestral**
1224 **reconstruction.** Effects on ERE binding of all possible single amino acid mutations at the four
1225 variable RH sites in the background of the maximum *a posteriori* (MAP) wild type AncSR1
1226 protein (*x*-axis), and in the background of the AltAll wild type AncSR1 protein, which has the
1227 second-most likely amino acid state at all sites at which the posterior probability of the MAP
1228 state is less than 0.8 (*y*-axis)³³. Pearson's r^2 is shown.
1229

