

Toward Open-Ended Fraternal Transitions in Individuality

Matthew Andres Moreno and Charles Ofria

BEACON Center, Michigan State University, East Lansing, MI 48824
mmore500@msu.edu

Abstract

1 The emergence of new replicating entities from the union of
2 simpler entities represent some of the most profound events
3 in natural evolutionary history. Such transitions in individuality
4 are essential to the evolution of the most complex forms
5 of life. As such, understanding these transitions is critical
6 to building artificial systems capable of open-ended evolution.
7 Alas, these transitions are challenging to induce or detect,
8 even with computational organisms. Here, we introduce the
9 DISHTINY (DIStributed Hierarchical Transitions in Individuality)
10 platform, which provides simple cell-like organisms with the
11 ability and incentive to unite into new individuals in a manner
12 that can continue to scale to subsequent transitions. The system
13 is designed to encourage these transitions so that they can be
14 studied: organisms that coordinate spatiotemporally can
15 maximize the rate of resource harvest, which is closely linked
16 to their reproductive ability. We demonstrate the hierarchical
17 emergence of multiple levels of individuality among simple
18 cell-like organisms that evolve parameters for manually-designed
19 strategies. During evolution, we observe reproductive division
20 of labor and close cooperation among cells, including resource-
21 sharing, aggregation of resource endowments for propagules,
22 and emergence of an apoptosis response to somatic mutation.
23 Many replicate populations evolved to direct their resources
24 toward low-level groups (behaving like multi-cellular
25 individuals) and many others evolved to direct their resources
26 toward high-level groups (acting as larger-scale multi-cellular
27 individuals).
28

Introduction

29
30 Artificial life researchers design systems that exhibit properties
31 of biological life in order to better understand their dynamics
32 and, often, to apply these principles toward engineering
33 applications such as artificial intelligence (Bedau, 2003).
34 Studies of evolution have been of particular interest to the
35 community, especially in regard to how organisms are produced
36 with increasing sophistication and complexity (Goldsby et al.,
37 2017). This particular issue is often described as “open-ended
38 evolution.” Although precise definitions and measures of open-
39 ended evolution are still being established, this term is generally
40 understood to refer to evolving systems that exhibit the continued
41 production of novelty (Taylor et al., 2016). Evolutionary
42 transitions in

43 individuality, which are key to the complexification and
44 diversification of biological life (Smith and Szathmary, 1997),
45 have been highlighted as key research targets with respect to
46 the question of open-ended evolution (Ray, 1996; Banzhaf
47 et al., 2016). In an evolutionary transition of individuality,
48 a new, more complex replicating entity is derived from the
49 combination of cooperating replicating entities that have
50 irrevocably entwined their long-term fates (West et al., 2015).
51 In particular, we focus on fraternal transition in individuality,
52 events where closely-related kin come together or stay together
53 to form a higher-level organism (Queller, 1997). Eusocial
54 insect colonies and multicellular organisms exemplify this
55 phenomenon (Smith and Szathmary, 1997). Like the definition
56 of open-ended evolution, the notion of what constitutes an
57 evolving individual is not concretely established. Commonly
58 indicated features include: close coordination and cooperation,
59 reproductive division of labor, reproductive bottlenecks, and
60 loss of ability to replicate independently (Ereshefsky and
61 Pedroso, 2015; Bouchard, 2013).

62 Our appreciation of fraternal transitions in individuality
63 benefits from experimental work probing the origins of
64 multicellularity. In the biological domain, Ratcliff et al. have
65 demonstrated evolution of multicellularity in yeast, deriving
66 fraternal clusters of cells that cling together in order to
67 maximize their settling rate (Ratcliff et al., 2012). The
68 contributions of Goldsby and collaborators are particularly
69 notable among computational Artificial life work on the origins
70 of multicellularity. Their evolutionary experiments track a
71 population composed of demes, distinct spatial domains
72 inhabited by clonal colonies of cells. Two distinct types of
73 reproduction occur: (1) cells reproduce within demes and (2)
74 deme reproduction, where a target deme is sterilized then
75 re-inoculated with genetic material from the parent deme.
76 With such methods, Goldsby et al. have studied division
77 of labor (Goldsby et al., 2010, 2012), the origin of soma
78 (Goldsby et al., 2014), and the evolution of morphological
79 development (Goldsby et al., 2017). We aspire to complement
80 deme-based approaches with a framework where higher level
81 individuality unfolds via cellular reproductions within a single
82 unified space. In particular, we are interested

83 in the potential for such a system to undergo nested hierar- 135
 84 chical transitions. 136
 85 Major challenges in studying evolutionary transitions in 137
 86 individuality include (1) determining the environmental con- 138
 87 ditions that will promote such a transition and then (2) rec- 139
 88 ognizing that a transition has occurred. In order to begin 140
 89 exploring transitions in individuality, we must devise a sys- 141
 90 tem in which we expect such transitions to occur repeatably 142
 91 and in a detectable manner. Once we can consistently in- 143
 92 duce and observe evolutionary transitions in individuality, 144
 93 we may subsequently proceed to relax aspects of such a sys- 145
 94 tem to explore in greater detail what conditions are neces- 146
 95 sary to induce transitions and how transitions can be det- 147
 96 ected. For now, we will focus on these initial goals in the 148
 97 context of fraternal transitions in individuality. 149
 98 To this end, we introduce the DISHTINY (DIStributed 150
 99 Hierarchical Transitions in IndividualitY) platform, which 151
 100 seeks to achieve the evolution of transitions in individual- 152
 101 ity by explicitly registering organisms in cooperating groups 153
 102 that coordinate spatiotemporally to maximize the harvest of 154
 103 a resource. Detection of such a transition in DISHTINY 155
 104 is accomplished by identifying resource-sharing and repro- 156
 105 ductive division of labor among organisms registered to the 157
 106 same cooperating group. We designed this system such that 158
 107 hierarchal transitions across an arbitrary number of levels of 159
 108 individuality can be selected for and meaningfully detected. 160
 109 We have focused this system on a rigid form of major tran- 161
 110 sition using simple organisms, but the underlying principles 162
 111 can be applied to a wide range of artificial life systems. Fur- 163
 112 thermore, DISHTINY is decentralized and amenable to mas- 164
 113 sive parallelization via distributed computing. We believe 165
 114 that such scalability — with respect to both concept and im- 166
 115 plementation — is an essential consideration in the pursuit 167
 116 of artificial systems capable of generating complexity and 168
 117 novelty rivaling that of biological life via open-ended evolu- 169
 118 tion (Ackley and Cannon, 2011; Ackley, 2016). 170

119 Methods 171

120 In order to demonstrate that the DISHTINY platform selects 173
 121 for detectable hierarchical transitions in individuality, we 174
 122 performed experiments where cell-like organisms evolved 175
 123 parameters to control manually designed behaviors such as 176
 124 resource-sharing, reproductive decision-making, and apop- 177
 125 tosis. We will first cover the design of the DISHTINY plat- 178
 126 form and then describe the simple cell-like organisms we 179
 127 used to evaluate the platform. 180

128 DISHTINY 182

129 DISHTINY allows cell-like organisms to replicate across a 183
 130 toroidal grid. Over discrete timesteps (“updates”), the cells 184
 131 can collect a continuous-valued resource. Once sufficient 185
 132 resource has been accrued, cells may pay 8.0 resource to 186
 133 place a daughter cell on an adjoining tile of the toroidal grid 187
 134 (i.e., reproduce), replacing any existing cell already there. 188

As cells reproduce, they can choose to include offspring in the parent’s cooperating “signaling channel” group or force offspring to create a new cooperating “signaling channel” group.

As shown at the top of Figure 1, resources appear at a single point then spread outwards update-by-update in a diamond-shaped wave, disappearing when the expanding wave reaches a predefined limit. Cells must be in a costly “activated” state to collect resource as it passes. The cell at the starting position of a resource wave is automatically activated, and will send the activate signal to neighboring cells on the same signaling channel. The newly activated cells, in turn, activate their own neighbors registered to the same signaling channel. Neighbors registered to other signaling channels do not activate. Each cell, after sending the activation signal, enters a temporary quiescent state so as not to reactivate from the signal. In this manner, cells sharing a signaling channel activate in concert with the expanding resource wave. As shown Figure 1a, b, the rate of resource collection for a cell is determined by the size and shape of its same-channel signaling network; small or fragmented same-channel signaling networks will frequently miss out on resource as it passes by.

Each cell pays a resource cost when it activates. This cost is outweighed by the resource collected such that cells that activate in concert with a resource wave derive a net benefit. Recall, though, that resource waves have a limited extent. Cells that activate outside the extent of a resource wave or activate out of sync with the resource wave (due to an indirect path from the cell that originated the signal) pay the activation cost but collect no resource. Cells that frequently activate erroneously use up their resource and die. In our implementation, organisms that accrue a resource debt of -11 or greater are killed. This erroneous activation scenario is depicted in Figure 1c.

In this manner, “Goldilocks” — not too small and not too big — signaling networks are selected for. Based on a randomly chosen starting location, resource wave starting points (seeds) are tiled over the toroidal grid such that the extents of the resource waves touch, but do not overlap. All waves start and proceed synchronously; when they complete, the next resource waves are seeded. This process ensures that selection for “Goldilocks” same-channel signaling networks is uniformly distributed over the toroidal grid.

Cells control the size and shape of their same-channel signaling group through strategic reproduction. Three choices are afforded: whether to reproduce at all, where among the four adjoining tiles of the toroidal grid to place their offspring, and whether the offspring should be registered to the parent’s signaling channel or be given a random channel ID (in the range 1 to 2^{22}). No guarantees are made about the uniqueness of a newly-generated channel ID, but chance collisions are rare.

Hierarchical levels are introduced into the system through




Figure 1: **Activation signaling, and net resource collection for three different-sized same-channel networks during a resource wave event.** At the top, a resource wave is depicted propagating over three updates and then ceasing for four updates (left to right). In row *a*, a small two-cell channel-signaling group (far left, in green) is activated; tracking the resource wave (top) yields a small net resource harvest (far right). In row *b*, an intermediate-sized 13-cell channel-signaling group yields a high net resource harvest. Finally, in row *c*, a large 29-cell channel-signaling group incurs a net negative resource harvest. In rows *a*, *b*, and *c*, dark purple indicates the active state, light purple indicates the quiescent state, and white indicates the ready state.

189 multiple separate, but overlaid, instantiations of this re- 211
 190 source wave/channel-signaling scheme. We refer to each 212
 191 independent resource wave/channel-signaling system as a 213
 192 “level.” In our experiments, we allowed two resource 214
 193 wave/channel-signaling levels, identified here as level one 215
 194 and level two. On level one, resource waves extended a ra- 216
 195 dius of three toroidal tiles. On level two they extended a 217
 196 radius of 24 toroidal tiles. On both levels, activated cells 218
 197 netted +1.0 resource from a resource wave, but suffered an 219
 198 activation penalty of -5.0 if no resource was available. Due 220
 199 to the different radii of resource waves on different levels, 221
 200 level one selects for small same-channel signaling networks 222
 201 and level two selects for large same-channel signaling net- 223
 202 works. 224

203 Cells were marked with two separate channel IDs, one 225
 204 for level one and another for level two. We enforced hier- 226
 205 archical nesting of same-channel signaling networks during 227
 206 reproduction: daughter cells may inherit neither channel ID, 228
 207 just the level-two channel ID, or both channel IDs. Daugh- 229
 208 ter cells may not inherit only the level-one channel ID while 230
 209 having a different level-two channel ID. The distribution of 230
 210 IDs across the level-two and level-one channels can be envi- 231

211 sioned by analogy to political countries and territories. Each 212
 213 country (i.e., level-two channel network) may have one or 214
 215 many territories (i.e., level-one channel network). However, 216
 217 no territory spans more than one country. Figure 2 depicts 218
 219 hierarchically nested channel states at the end of three evo- 220
 221 lutionary runs. 222

223 Channel IDs enable straightforward detection of an evolu- 224
 225 tionary transition in individuality. Because common channel 226
 227 IDs may only arise systematically through inheritance, com- 228
 229 mon channel IDs indicate a close hereditary relationship in 229
 230 addition to a close cooperative relationship. Because new 230
 231 channel IDs arise first in a single cell, same-channel sig- 231
 232 naling networks are reproductively bottlenecked, ensuring 232
 233 meaningful reproductive lineages at the level of the same- 233
 234 channel signaling network. To recognize an evolutionary 234
 235 transition in individuality, we therefore evaluate 235

- 236 1. Do cells with the same channel ID choose to share re- 236
 237 sources (e.g., cooperate)? 237
- 238 2. Is there division of reproductive labor between members 238
 239 of the same channel (e.g., do cells at the interior of a net- 239
 240 work cede reproduction to those at the periphery)? 240

232 If these conditions are met among cells sharing the same 285
 233 level-one channel, a first-level transition in individuality 286
 234 may have occurred. Likewise, if these conditions are met 287
 235 among cells sharing the same level-two channel, a second- 288
 236 level transition in individuality may have occurred. In either 289
 237 case, observation of altruistic behavior, such as an apoptosis 290
 238 response to mutation, would further evidence a transition. 291

239 Organisms

240 We performed our experiments using cell-like organisms 294
 241 composed of 15 floating-point parameters, each controlling 295
 242 a specific strategy component pertinent to transitions in indi- 296
 243 viduality (i.e., reproductive division of labor, resource pool- 297
 244 ing, apoptosis, propagule generation, and propagule endow- 298
 245 ment). These particular cell-like organisms are in no way 299
 246 inherent to the DISHTINY platform, but were merely de- 300
 247 veloped to study transitions using as simple a model system 301
 248 as feasible. On reproduction, we applied mutation to each 302
 249 parameter independently with probability 0.00005. 303

250 The **aversion parameters** (A_1 and A_2) allow cells to 304
 251 avoid reproducing over neighbors sharing the same signal- 305
 252 ing channel. Specifically, they control the probability that a 306
 253 cell declines to supplant a neighbor sharing the same level- 307
 254 one (A_1) or level-two (A_2) channel ID. If a cell declines 308
 255 to place its offspring in all four adjoining tiles, it does not 309
 256 reproduce. Mutation is performed by a redraw from the uni- 310
 257 form distribution $U(-0.5, 1.5)$ clamped to the range $[0, 1]$. 311

258 The **resource allocation parameters** control the propor- 312
 259 tion of resources that go to the cell's stockpile (P_c), its level- 313
 260 one channel's resource pool (P_1), or its level-two channel's 314
 261 resource pool (P_2). These parameters are initialized by a 315
 262 draw from $U(-1.0, 2.0)$ clamped to the range $[0, 1]$ and mu- 316
 263 tated by addition of a normal value drawn from $N(0.0, 0.2)$ 317
 264 with the result clamped to the range $[0, 1]$. The set P_c, P_1, P_2 318
 265 is always normalized to sum to 1. 319

266 Channel resource pools are identical to an organism's 320
 267 individual stockpile, except that any deficit is distributed 321
 268 evenly among the individual organism's stockpile. On every 322
 269 update, cells can spend from their individual stockpile 323
 270 to reproduce or from a channel pool, with priority given to 324
 271 cells nearest to the centroid of that pool's members. As such, 325
 272 pool-funded reproduction fills in a same-channel signaling 326
 273 network from the inside out and help produce diamond- 327
 274 shaped same-channel signaling networks. (Distance is mea- 328
 275 sured using the taxicab metric.) 329

276 **Channel cap parameters** C_1 and C_2 regulate the size 330
 277 of same-channel signaling networks. When an organism re- 331
 278 produces, it checks the size of its level-one signaling net- 332
 279 work against C_1 and the size of its level-two signaling group 333
 280 against C_2 . If neither cap is met or exceeded, then the or- 334
 281 ganism will produce an offspring sharing both of its chan- 335
 282 nel IDs. If only the C_1 cap is exceeded, then the organism 336
 283 will produce an offspring with new level-one channel ID but 337
 284 identical level-two channel ID. Finally, if the C_2 cap is ex- 338

ceeded, then the organism will produce an offspring with 339
 new IDs for both channels. For level-one caps, these pa- 340
 rameters are initialized by a draw from $U(0.0, 16.0)$. For 341
 level-two caps, these parameters are initialized by a draw 342
 from $U(0.0, 128.0)$. Both are mutated by addition of a value 343
 drawn from $N(0.0, 24.0)$ with the result clamped to be non- 344
 negative. 345

The **endowment parameters** E_c , E_1 , and E_2 deter- 346
 mine the amount of resource provided to offspring. This 347
 endowment is paid as an additional cost by the cell stock- 348
 pile (or same-channel resource pool) funding a reproduc- 349
 tion. The full amount of the received endowment is di- 350
 vided between the daughter cell's stockpile, level-one same- 351
 channel resource pool, and level-two same-channel resource 352
 pool according to the offspring's resource allocation param- 353
 eters. E_c is the endowment amount paid to an offspring 354
 that shares both channel IDs of the parent; E_1 is the en- 355
 dowment paid to an offspring that shares just the level-two 356
 channel ID of the parent; and E_2 is the endowment paid to 357
 an offspring that shares neither the level-one nor the level- 358
 two channel ID of the parent. Endowed resources help new- 359
 channel propagules to rapidly grow their signaling network 360
 in order to begin collecting resource at a rate competitive 361
 to other well-established same-channel signaling networks. 362
 In order that adequate resource remain to ensure parental 363
 stability, endowment was paid out only after twice the en- 364
 dowment amount had been accrued (leaving an amount of 365
 resource equal to the endowment remaining with the par- 366
 ent). Cell level endowments are initialized by a draw from 367
 $U(0.0, 5.0)$. Level-one endowments are initialized by a 368
 draw from $U(0.0, 80.0)$. Level-two endowments are initial- 369
 ized by a draw from $U(0.0, 405.0)$. All endowments are mu- 370
 tated by addition of a value drawn from $N(0.0, 10.0)$ with 371
 the result clamped to be non-negative 372

Parameters M_c , M_1 , and M_2 control the **apoptosis re-** 373
sponse to mutation. Each time that a mutation occurs 374
 during reproduction, the mutated offspring attempts suicide 375
 with probability M_c if it shares both channel IDs of its 376
 parent, probability M_1 if it shares just the level-two chan- 377
 nel ID of its parent, and probability M_2 if it shares nei- 378
 ther channel ID of the parent. The M_x value applied is 379
 from the offspring's genotype after mutation. Attempted 380
 suicide succeeds 90% of the time. This capacity enables 381
 first- or second-level individuals to combat somatic muta- 382
 tion. Initialization and mutation each of these parameters is 383
 performed by a redraw from the distribution $U(-0.5, 1.5)$ 384
 clamped to the range $[0, 1]$. 385

Finally, parameters S_1 and S_2 **fine-tune site choice for** 386
offspring placement. If an organism is placing an off- 387
 spring with identical channel IDs, with probability S_1 the 388
 four possible sites for offspring placement are considered in 389
 order of increasing distance from the centroid of the par- 390
 ent's level-one signaling network. If an organism is placing 391
 an offspring with identical level-two channel ID but differ- 392

339 ent level-one channel ID, with probability S_2 the four possi- 392
340 ble sites for offspring placement are considered in order of 393
341 increasing distance from the centroid of the parent's level- 394
342 two same-channel signaling network. Otherwise, the four 395
343 possible sites for offspring placement are considered in a 396
344 random order. Initialization and mutation are performed by 397
345 a draw from the distribution $U(-0.5, 1.5)$ clamped to the 398
346 range $[0, 1]$. 399

347 Treatments 401

348 Our standard treatment was designed to assess the evolu- 402
349 tionary trajectories of populations in DISHTINY. We seeded 403
350 each tile on the 120×120 toroidal grid with a randomized 404
351 organism and ran the simulation for 20 million updates. In 405
352 order to facilitate turnover, we culled the population inter- 406
353 mittently. Starting at update 500,000, and every 50,000 up- 407
354 dates thereafter, we randomly selected second-level channel 408
355 IDs and killed all cells with that channel ID, continuing un- 409
356 til at least 5% of grid tiles were empty. We performed 50 410
357 replicates within this treatment. On average, each cellular 411
358 generation took just over 500 updates. Across all succes- 412
359 sive 10,000 update segments of all replicates, the mean num- 413
360 ber of cellular generations elapsed per 10,000 updates was 414
361 19.2 with a standard deviation of 2.7 cellular generations per 415
362 10,000 updates. 416

363 In order to detangle the impact of same-channel signaling 417
364 networks with respect to kin recognition versus cooperation 418
365 to increase resource collection rate, we performed control 419
366 evolutionary trials where same-channel signaling networks 420
367 did not affect cellular resource collection rate. Under control 421
368 conditions, same-channel signaling networks just helped 422
369 cells recognize other related cells. In our implementation, 423
370 this treatment corresponded to a constant per-update inflow 424
371 of 0.02 resource units into all cells. All cells were activated 425
372 (in order to take up the resource) at all updates and no cost 426
373 for activation was assessed. We chose this resource inflow 427
374 rate in order to approximately match the cellular generation 428
375 rate of the control treatment to that of the standard treat- 429
376 ment. In control runs, each cellular generation took around 430
377 450 updates. Across all successive 10,000 update segments 431
378 of all replicates, the mean number of cellular generations 432
379 elapsed per 10,000 updates was 22.0 with a standard devia- 433
380 tion of 2.0 cellular generations per 10,000 updates. Due to 434
381 checkpoint-restart failures on our compute cluster, control 435
382 experiments were curtailed at 3 million updates. All other 436
383 aspects of control runs, including culling and the function- 437
384 ality of all lifestyle parameters, were otherwise identical to 438
385 standard conditions. We performed 50 replications of the 439
386 control treatment. 440

387 In standard evolutionary runs, we observed a spectrum of 441
388 evolved resource-caching strategies. To assess the relative 442
389 fitness of these evolved organisms, we ran competitions be- 443
390 tween the most common genotype from three standard evo-
391 lutionary runs. The first genotype allocated resource ex-

clusively to its first-level same-channel resource pool (i.e.,
 $P_1 = 1.0$), the second split resource evenly between its first-
level and second-level resource pool (i.e., $P_1 = P_2 = 0.5$),
and the third allocated resource primarily to the second-level
resource pool (i.e., $P_2 > P_1$). (No most-common genotypes
allocated resource exclusively to the second-level resource
pool.) We seeded each competition with three copies of each
genotype, uniformly spaced over the 120×120 toroidal grid
with random arrangement. Each competition lasted 2 mil-
lion updates. We performed 50 runs in this experiment.

402 Implementation 406

403 We performed our computational experiments at the Michi-
404 gan State University High Performance Computing Center. Each
405 replicate of standard evolutionary experiments required approxi-
406 mately six days of compute time to reach 20 million updates.
407 Each replicate of control evolutionary experiments expended
408 approximately two days of compute time to reach 3 million
409 updates. Control runs were somewhat slower than standard
410 runs, perhaps due to increased computational overhead associ-
411 ated with bookkeeping for the larger same-channel groups
412 that evolved under the control conditions. Each replicate of
413 competition experiments consumed approximately ten hours of
414 compute time. For standard evolutionary experiments, data
415 processing required approximately four hours of compute time
416 per run. Other data processing was computationally negligible.
417

418 We implemented our experimental system using the
419 Empirical library for scientific software development in
420 C++, available at [https://github.com/devosoft/
421 Empirical](https://github.com/devosoft/Empirical). The code used to perform and analyze our
422 experiments, our figures, data from our experiments, and a
423 live in-browser demo of our system is available via the Open
424 Science Framework at <https://osf.io/ewvg8/>.


Results and Discussion

Standard Evolutionary Experiments


A spectrum of resource allocation strategies ranging from
purely allocation to level-one same-channel resource pools
to primarily allocation to level-two same-channel resource
pools were observed at the conclusion of different runs
of our evolutionary simulation (mean cellular generation
37,168 with standard deviation 4,684). We interpret these
outcomes as ranging between individuality at the level of
first-level same-channel groups to individuality at the level
of second-level same-channel groups. Figure 2 shows the
level-one and level-two signaling networks at the end of
runs where first-, split-, and second-level resource alloca-
tion evolved, respectively. First-level allocators form some-
what irregular level-two amalgamations of diverse level-
one networks. Second-level allocators form highly regu-
lar diamond-shaped level-two signaling networks. Split-
allocation individuals exhibit a level-two phenotype of inter-
mediate regularity. Figure 3 shows a time series of signal-

	Competitors			Mean Dominant ($\pm S.D.$)			Pop Mean ($\pm S.D.$)		Control Pop Mean ($\pm S.D.$)
	$P_1 = 1.0$	$P_2 = P_1$	$P_2 > P_1$	$P_1 = 1.0$	$1.0 > P_1 > P_2$	$P_2 \geq P_1$	<i>all</i>	<i>all</i>	<i>all</i>
	29920	33852	47507	30841 \pm 3183	35346 \pm 3444	39315 \pm 3346	6670 \pm 729	6069 \pm 672	6626 \pm 377
<i>Cell Gen.</i>	20M	20M	20M	20M	20M	20M	3.3M	3M	3M
<i>Upd. n</i>	1	1	1	9	7	34	50	50	50
A_1	0.00	0.00	0.89	0.23 \pm 0.35	0.50 \pm 0.47	0.57 \pm 0.46	0.53 \pm 0.37	0.53 \pm 0.35	0.56 \pm 0.34
A_2	1.00	1.00	1.00	1.00 \pm 0.00	1.00 \pm 0.00	1.00 \pm 0.00	1.00 \pm 0.00	1.00 \pm 0.00	0.99 \pm 0.01
P_c	0.00	0.00	0.00	0.00 \pm 0.00	0.00 \pm 0.00	0.03 \pm 0.05	0.02 \pm 0.03	0.02 \pm 0.02	0.00 \pm 0.00
P_1	1.00	0.50	0.00	1.00 \pm 0.00	0.60 \pm 0.07	0.28 \pm 0.16	0.42 \pm 0.25	0.42 \pm 0.24	0.56 \pm 0.37
P_2	0.00	0.50	1.00	0.00 \pm 0.00	0.40 \pm 0.07	0.69 \pm 0.14	0.56 \pm 0.24	0.56 \pm 0.24	0.44 \pm 0.37
C_1	3.13	3.45	2.04	3.90 \pm 0.60	3.38 \pm 0.33	3.03 \pm 0.69	3.21 \pm 0.63	3.21 \pm 0.60	28.6 \pm 21.7
C_2	233.2	238.6	290.2	230.6 \pm 71.1	192.7 \pm 45.3	271.6 \pm 73.6	201.5 \pm 58.1	195.8 \pm 55.3	484.0 \pm 123.5
E_c	0.87	0.14	4.20	0.29 \pm 0.37	0.44 \pm 0.59	0.21 \pm 0.75	1.14 \pm 1.07	1.21 \pm 1.05	1.50 \pm 1.08
E_1	33.4	11.7	4.80	47.2 \pm 21.7	21.3 \pm 12.0	4.62 \pm 7.05	18.1 \pm 16.2	19.2 \pm 15.9	28.9 \pm 22.2
E_2	341.4	397.4	321.1	231.2 \pm 94.3	283.1 \pm 57.0	325.4 \pm 68.9	303.0 \pm 66.5	302.7 \pm 65.2	317.3 \pm 66.3
M_c	0.11	1.00	0.66	0.33 \pm 0.41	0.74 \pm 0.31	0.67 \pm 0.35	0.39 \pm 0.32	0.39 \pm 0.31	0.18 \pm 0.23
M_1	0.00	1.00	0.40	0.52 \pm 0.41	0.65 \pm 0.46	0.68 \pm 0.38	0.52 \pm 0.37	0.51 \pm 0.35	0.48 \pm 0.33
M_2	0.00	0.44	1.00	0.45 \pm 0.39	0.52 \pm 0.37	0.50 \pm 0.42	0.47 \pm 0.33	0.47 \pm 0.32	0.53 \pm 0.36
S_1	0.00	1.00	1.00	0.65 \pm 0.38	0.55 \pm 0.40	0.47 \pm 0.42	0.39 \pm 0.36	0.40 \pm 0.34	0.47 \pm 0.34
S_2	0.00	0.01	0.46	0.51 \pm 0.43	0.35 \pm 0.39	0.45 \pm 0.39	0.47 \pm 0.34	0.46 \pm 0.34	0.55 \pm 0.34


Table 1: The leftmost two table segments enumerate genotypes used as seeds for competition experiments (“Competitors”) and the mean values of the most abundant genotype at the end of evolutionary runs (“Mean Dominant”), both partitioned by resource-caching strategy. The rightmost table segments enumerate the population mean genotype values for standard evolutionary trials (“Pop Mean”) and control treatments (“Control Pop Mean”), matched at both absolute update count and (approximately) elapsed cellular generations.



(a) Mean $P_c=0.03$, $P_1=0.75$, $P_2=0.23$; cell gen. 29920



(b) Mean $P_c=0.03$, $P_1=0.51$, $P_2=0.49$; cell gen. 33852



(c) Mean $P_c=0.08$, $P_1=0.01$, $P_2=0.90$; cell gen. 47507

Figure 2: End state of same-channel signaling networks in replicates where resource was exclusively allocated to first-level channel pools (2a), was split evenly between first- and second-level channel pools (2b), and was primarily allocated to second-level channel pools (2c). Level-one channels are coded by color saturation and level-two channels are coded by color hue. A single cell-like organism occupies each grid tile except for black tiles, which are empty.




Figure 3: Progression of of same-channel level-one and level-two signaling networks states in an evolutionary run where level-two resource sharing evolved. Level-one channels are coded by color saturation and level-two channels are coded by color hue. A single cell-like organism occupies each grid tile except for black tiles, which are empty.

ing network snapshots in an evolutionary run where second-level individuality evolved. Table 1 summarizes most-common genotypes observed at the end of our evolutionary simulations. In the standard treatment, all evolved genotypes had A_2 fixed at 1.0. So reproduction over cells sharing the same level-two channel was universally avoided; genotypes evolved so that cells declined to reproduce when they were located at the interior of level-two same-channel signaling networks. However, a variety of resource-caching strategies evolved. Most-abundant genotypes at the end of nine evolutionary runs exclusively cached resource in organisms' level-one signaling network's pool (i.e., $P_1 = 1.0$). We observed strategies where resource was primarily, but not entirely, cached in an organism's level-one signaling network pool (i.e., $1.0 > P_1 > P_2$) as the most-abundant genotype at the end of seven evolutionary runs. In one run, the most-abundant final genotype split resources evenly between an organism's level-one and level-two signaling network pool ($P_1 = P_2 = 0.5$). Finally, we observed strategies where resource was primarily, but not entirely, cached in an organism's level-two signaling network pool (i.e., $1.0 > P_2 > P_1$) as the most-abundant genotype at the end of 33 evolutionary runs.

We suspect that a trade-off between growth rate and long-term stability prompted the universal allocation of at least some resource to level-one pools and/or cell stockpiles. Cell- and level-one resource caching might function something like saving for a rainy day. Because reproduction over level-two channel-mates was universally avoided, cells and level-one same-channel networks situated at the interior of a larger level-two same-channel network do not expend their resource pools unless that larger level-two same-channel network is damaged, exposing them to directly-adjacent cells of a different level-two channel. Thus, resource accumulates in cell stockpiles and level-one pools until the level-two same-channel network comes under stress. Split allocation might also represent hedging against defecation of a second-level channel-mate by via somatic mutation.

Indeed, we did observe selection for apoptosis in the 41 replicates where the dominant genotype employed second-level resource caching. In these replicates, the average population mean value of M_c was 0.68 with standard deviation 0.33, significantly greater than the value $M_c = 0.5$ we would expect in the absence of a selective pressure on apoptosis response to mutation ($p < 0.001$, bootstrap test).


To assess whether heavy second-level resource allocators, which we characterize as higher-level individuals, were more likely to employ apoptosis to mitigate somatic mutation, we examined the relationship between first- and second-level resource pooling and cellular apoptosis at the conclusion of our 50 replicate evolutionary trials. We observed a significant negative correlation between dominant genotype P_1 and M_c ($p < 0.05$; bootstrap test; Figure 4a) and a significant positive correlation between dominant genotype P_2 and M_c ($p < 0.05$; bootstrap test; Figure 4b). This result suggests that second-level individuals, in particular, relied on apoptosis to mitigate somatic mutation.

We also assessed whether higher-level individuals provided larger resource endowments to their second-level propagules (offspring sharing neither the level-one nor the level-two channel ID with the parent). We examined the relationship between first and second-level resource pooling and dominant genotype second-level propagule endowment at the conclusion of our 50 replicate evolutionary trials. We observed a significant negative correlation between dominant genotype P_1 and E_2 ($p < 0.05$; bootstrap test) and a significant positive correlation between dominant genotype P_2 and E_2 ($p < 0.05$; bootstrap test). Second-level individuals might provide larger endowments to propagules simply due to a greater capacity to collect resource or perhaps because of stronger selection for well-endowed offspring when competing against other second-level individuals.

This result prompts the reverse question: do lower-level individuals provide larger resource endowments to first-level propagules (offspring that do not share level-one channel ID with the parent but may or may not share level-two channel ID with the parent)? Indeed, we observed a significant positive correlation between first-level resource sharing and first-level endowment ($p < 0.0001$; bootstrap test) and a significant negative correlation between second-level resource sharing and first-level endowment ($p < 0.0001$; bootstrap test). Cells that pool resource with their smaller level-one same-channel group tend to invest more heavily into the direct offshoots of their level-one same-channel group than cells that pool resource with their larger level-two same-channel group. This observation suggests that, although cells do not directly displace their level-one channel-mates, competitive dynamics between may be at play.

Competition Experiments

Next, we wanted to compare first-, second-, and split-level allocators to determine which genotype was the most fit. We ran competition experiments between dominant genotypes from evolutionary runs representative of each of these strategies. To prevent further evolution, we disabled mutation for these experiments. To represent first-level allocators, we selected randomly from the nine pure first-level allocator dominant genotypes we observed. To represent the split-level allocators, we selected the single dominant genotype where resource was partitioned exactly evenly between first- and second-level channel pools. To represent second-level allocators, we selected the dominant genotype with the largest second-level allocation proportion. Table 1 enumerates the three representative genotypes used. Figure 5 shows a time series of signaling network snapshots in a competition experiment run. Colonies of each genotype can be seen to grow from each seed and then clash, ultimately yielding a



(a) Correlation plot of dominant genotype P_1 and dominant genotype M_c . (b) Correlation plot of dominant genotype P_2 and dominant genotype M_c .

Figure 4: Plots of dominant resource caching strategies and dominant apoptosis strategies. A bootstrapped 95% confidence interval for the fit is shaded. Both correlations are statistically significant ($p < 0.05$; bootstrap test).

551 population dominated by second-level allocators.

552 Indeed, the second-level resource caching strategy be-
553 came most abundant in all 50 trials. Across the 50 replicates,
554 at update 1.5 million (cellular generation 3489 with stan-
555 dard deviation 40) the second-level resource caching strat-
556 egy constituted 90.2%, with standard deviation 3.8%, of the
557 competing population of cells. In the absence of mutation,
558 second-level allocators tend to exhibit greater fitness than
559 split- and first-level allocators ($p < 0.0001$; two-tailed exact
560 test).

561 In competition experiments, however, higher-level indi-
562 viduals likely benefited from elimination of somatic muta-
563 tion. To assess the relative fitness of first- and second-level
564 individuals without mutation disabled, we examined the rela-
565 tionship between first- and second-level resource pooling
566 and the rate of cellular reproduction at the end of each of
567 the 50 replicate evolutionary trials performed. We observed
568 a significant negative correlation between mean P_1 and cel-
569 lular reproduction rate ($p < 0.0001$; bootstrap test; Figure
570 6a) and a significant positive correlation between mean P_2
571 and cellular reproduction rate ($p < 0.0001$; bootstrap test;
572 Figure 6b). This result suggests that second-level allocators
573 tend to collect resource more effectively than split- and first-
574 level allocators.

575 Control Evolutionary Experiments

576 Under control conditions where resource was distributed
577 evenly to all cells regardless of same-channel group con-
578 figuration, split-level resource caching also evolved. Split-
579 level allocation was the most common strategy at update 3
580 million in all replicates. Strategies where resource was pri-
581 marily, but not entirely, cached in an organism's level-one

582 signaling network pool (i.e., $1.0 > P_1 > P_2$) were most-
583 abundant at the end of 33 evolutionary runs and strategies
584 where resource was primarily, but not entirely, cached in
585 an organism's level-two signaling network pool (i.e., $1.0 >$
586 $P_2 > P_1$) were most-abundant at the end of 17 evolution-
587 ary runs. As shown in Table 1, the average population mean
588 of P_1 is greater in the control treatment than in the standard
589 treatment at time-points matched by absolute elapsed update
590 count and approximate elapsed cellular generations, but this
591 difference is not statistically significant.

592 Consistent with the standard treatment, we observed
593 strong selection against direct reproductive competition be-
594 tween channel-mates at update 3 million in the control
595 treatment. Nearly all most-common genotypes completely
596 avoided reproducing over level-two channel-mates (i.e.,
597 $A_2 = 1.0$), except for a single most-common genotype
598 where a very slim probability of reproducing over level-two
599 channel-mates was allowed ($A_2 = 0.996$).

600 The emergence of resource-sharing and competition
601 avoidance under control conditions suggests kin recognition
602 alone can prompt some aspects of higher-level individual-
603 ity. However, we observed selection *against* the apoptosis
604 response to mutation, M_c , under control conditions. Across
605 50 replicates of the control treatment, the average population
606 mean value of M_c was 0.18 with standard deviation 0.23 —
607 significantly less than the value $M_c = 0.5$ expected with-
608 out selective pressure against apoptosis response to mutation
609 ($p < 0.0001$, two-tailed t test). Indeed, population mean M_c
610 for control runs was also significantly reduced compared to
611 the standard treatment at time-points matched by absolute
612 elapsed update count ($p < 0.001$; two-tailed t test) and by
613 approximate elapsed cellular generations ($p < 0.001$; two-





Figure 5: Progression of same-channel level-one and level-two signaling networks states in a competition run. We seeded the grid with three copies of each of three champion genotypes from evolutionary trials. Then, with mutation disabled to prevent further evolution, the genotypes competed. Level-one channels are coded by color saturation and level-two channels are coded by color hue. A single cell-like organism occupies each grid tile. <https://doi.org/10.7554/peerj.c.36007> for details. Access this article at <https://doi.org/10.7554/peerj.c.36007>



(a) Correlation plot of population mean P_1 and population net reproduction rate.

(b) Correlation plot of population mean P_2 and population net reproduction rate.

Figure 6: Mean resource caching strategies and net reproduction rate across populations. A bootstrapped 95% confidence interval for the fit is shaded. Both correlations are statistically significant ($p < 0.0001$; bootstrap test).




Figure 7: End state (update 3000000, cell gen. 6916) of same-channel signaling networks evolved under the control treatment. Level-one channels are coded by color saturation and level-two channels are coded by color hue. A single cell-like organism occupies each grid tile except for black tiles, which are empty.

614 tailed t test). Perhaps under control conditions, the apoptosis
 615 response to mutation is disfavored because kin groups stand
 616 to lose less from mutant members (i.e., the resource penalty
 617 for excessive same-channel network expansion is absent). It
 618 appears that, at least in our system, kin recognition alone
 619 does not suffice to prompt full-fledged fraternal transitions
 620 in individuality.

621 In the absence of resource penalties for erroneous acti-
 622 vation under control conditions, we also observed the evo-
 623 lution of larger same-channel groups. At update 3 million,
 624 most-common genotypes encoded a level-two same-channel
 625 cap C_2 of 484.0 cells with standard deviation of 123.5.
 626 Compared to the standard treatment, control runs exhibited
 627 larger mean level-two same-channel caps C_2 at time-points
 628 matched by absolute elapsed update count ($p < 0.0001$;
 629 two-tailed t test) and approximate elapsed cellular genera-
 630 tions ($p < 0.0001$; two-tailed t test). Even at 20 million up-
 631 dates, when evolution had elapsed around six times as many
 632 cellular generations in the standard treatment compared to
 633 the control treatment at update 3 million, mean level-two
 634 same-channel caps C_2 reached only 262.9 with standard de-
 635 viation 72.2 under the standard treatment. This is signifi-
 636 cantly smaller than mean C_2 under the control treatment at
 637 update 3 million ($p < 0.0001$; two-tailed t test). Figure 7 de-
 638 picts the comparatively large same-channel level two groups
 639 present at the end of a control run. Table 1 summarizes most-
 640 common genotypes observed under the control treatment.

641

Conclusion

642 Using simple organisms that evolve parameters for a set
 643 of manually-designed strategies, we have demonstrated that
 644 DISHTINY selects for genotypes that exhibit high-level in-
 645 dividuality. We observed a spectrum of first- and second-

646 level individuality among evolutionary outcomes. Specifi- 698
 647 cally, we observed 699

- 648 1. reproductive division of labor among members of the 700
 649 same channel (i.e., individuals enveloped in a same- 701
 650 channel signaling network ceded reproduction to those at 702
 651 the periphery), 703
 704
- 652 2. cooperation between members of the same channel (i.e., 705
 653 pooling of resource on same-channel signaling networks), 706
 707
- 654 3. reproductive bottlenecking (i.e., groups of cells sharing a 708
 655 channel ID descend from a single originator of that chan- 709
 656 nel ID), and 710
- 657 4. suppression of somatic mutation via apoptosis coincident 711
 658 with second-level individuality. 712

659 Competition experiments revealed that second-level indi- 713
 660 viduals usually outcompete lower-level individuals. The 714
 661 magnitude of resource endowment for propagules was also
 662 correlated with second-level individuality. 715

663 Although shifts in individuality to level-one and level- 716
 664 two signaling networks were both observed, the question of 717
 665 whether these transitions were truly hierarchical in nature is 718
 666 debatable. That is, it is not clear whether level-one individ- 719
 667 uality was to some extent preserved in or necessary for the 720
 668 emergence of level-two individuality. Given the nature of 721
 669 the manually-designed strategies for resource-pooling and 722
 670 reproductive division of labor, level-two resource pooling 723
 671 and division of labor could readily leapfrog over level-one 724
 672 resource pooling and division of labor and, in many ways, 725
 673 seemed to completely supersede those level-one efforts. 726
 674 We believe that this is a shortcoming of the manual design 727
 675 of behaviors for which simple cell-like organisms evolved 728
 676 parameters, not the DISHTINY platform itself. We have 729
 677 nevertheless demonstrated that DISHTINY ultimately se- 730
 678 lects for high-level individuality. We are eager to work 731
 679 with more sophisticated cell-like organisms capable of arbi- 732
 680 trary computation via genetic programming in order to pur- 733
 681 sue more open-ended evolutionary experiments. We will 734
 682 also test the implications of relaxing current arbitrary re- 735
 683 strictions that artificially promote transitions, such as the hi- 736
 684 erarchical nesting of same-channel signaling networks and 737
 685 the explicitly-defined signaling networks themselves, leav- 738
 686 ing these details to evolution to figure out. Further work 739
 687 will provide valuable insight into scientific questions relat- 740
 688 ing to major evolutionary transitions such as the role of pre- 741
 689 existing phenotypic plasticity (Clune et al., 2007; Lalejini 742
 690 and Ofria, 2016), pre-existing environmental interactions, 743
 691 pre-existing reproductive division of labor, and how transi- 744
 692 tions relate to increases in organizational (Goldsby et al., 745
 693 2012), structural, and functional (Goldsby et al., 2014) com- 746
 694 plexity. 747
 695 We believe that such an approach also provides a unique 748
 696 opportunity to fundamentally advance Artificial life with re- 749
 697 spect to open-ended evolution. Fundamental to this goal is

scale. The DISHTINY platform trivially scales to select for
 an arbitrary number of hierarchical levels of individuality
 (not just the two hierarchical levels explored in these exper-
 iments). Importantly, the platform is implemented in a de-
 centralized manner and can comfortably scale as additional
 computing resources are provided. Parallel computing is
 widely exploited in evolutionary computing, where subpop-
 ulations are farmed out for periods of isolated evolution or
 single genotypes are farmed out for fitness evaluation (Lin
 et al., 1994; Real et al., 2017). DISHTINY presents a more
 fundamental parallelization potential: principled paralleliza-
 tion of the evolving individual phenotype at arbitrary scale
 (i.e., a high-level individual as a large collection of individ-
 ual cells on the toroidal grid). Such parallelization will be
 key to realizing evolving computational systems with scale
 — and, perhaps, complexity — approaching biological sys-
 tems.

Acknowledgements

Thanks to members of the DEVOLAB, in particular Michael
 J. Wiser, for feedback on statistical methods and Heather
 Goldsby for concept and editing feedback. This research
 was supported in part by NSF grants DEB-1655715 and
 DBI-0939454, and by Michigan State University through
 the computational resources provided by the Institute for
 Cyber-Enabled Research. This material is based upon work
 supported by the National Science Foundation Graduate Re-
 search Fellowship under Grant No. DGE-1424871. Any
 opinions, findings, and conclusions or recommendations ex-
 pressed in this material are those of the author(s) and do not
 necessarily reflect the views of the National Science Founda-
 tion.

References

- Ackley, D. H. (2016). Indefinite scalability for living computation.
 In *Proceedings of the Thirtieth AAAI Conference on Artificial
 Intelligence*.
- Ackley, D. H. and Cannon, D. C. (2011). Pursue robust indefinite
 scalability. In *HotOS*.
- Banzhaf, W., Baumgaertner, B., Beslon, G., Doursat, R., Fos-
 ter, J. A., McMullin, B., De Melo, V. V., Miconi, T., Spec-
 tor, L., Stepney, S., et al. (2016). Defining and simulat-
 ing open-ended novelty: requirements, guidelines, and chal-
 lenges. *Theory in Biosciences*, 135(3):131–161.
- Bedau, M. A. (2003). Artificial life: organization, adaptation and
 complexity from the bottom up. *Trends in cognitive sciences*,
 7(11):505–512.
- Bouchard, F. (2013). What is a symbiotic superindividual and how
 do you measure its fitness. *From groups to individuals: evo-
 lution and emerging individuality*, 243.
- Clune, J., Ofria, C., and Pennock, R. T. (2007). Investigating the
 emergence of phenotypic plasticity in evolving digital organ-
 isms. In *European Conference on Artificial Life*, pages 74–83.
 Springer.

- 750 Ereshefsky, M. and Pedroso, M. (2015). Rethinking evolutionary
751 individuality. *Proceedings of the National Academy of Sci-*
752 *ences*, 112(33):10126–10132.
- 753 Goldsby, H. J., Dornhaus, A., Kerr, B., and Ofria, C. (2012). Task-
754 switching costs promote the evolution of division of labor and
755 shifts in individuality. *Proceedings of the National Academy*
756 *of Sciences*, 109(34):13686–13691.
- 757 Goldsby, H. J., Knoester, D. B., and Ofria, C. (2010). Evolution
758 of division of labor in genetically homogenous groups. In
759 *Proceedings of the 12th annual conference on Genetic and*
760 *evolutionary computation*, pages 135–142. ACM.
- 761 Goldsby, H. J., Knoester, D. B., Ofria, C., and Kerr, B. (2014).
762 The evolutionary origin of somatic cells under the dirty work
763 hypothesis. *PLoS biology*, 12(5):e1001858.
- 764 Goldsby, H. J., Young, R. L., Hofmann, H. A., and Hintze, A.
765 (2017). Increasing the complexity of solutions produced by
766 an evolutionary developmental system. In *Proceedings of the*
767 *Genetic and Evolutionary Computation Conference Compan-*
768 *ion*, pages 57–58. ACM.
- 769 Lalejini, A. and Ofria, C. (2016). The evolutionary origins of phe-
770 notypic plasticity. In *Proceedings of the Artificial Life Con-*
771 *ference*.
- 772 Lin, S.-C., Punch, W. F., and Goodman, E. D. (1994). Coarse-
773 grain parallel genetic algorithms: Categorization and new ap-
774 proach. In *Parallel and Distributed Processing, 1994. Pro-*
775 *ceedings. Sixth IEEE Symposium on*, pages 28–37. IEEE.
- 776 Queller, D. C. (1997). Cooperators since life began. *The Quarterly*
777 *Review of Biology*, 72(2):184–188.
- 778 Ratcliff, W. C., Denison, R. F., Borrello, M., and Travisano, M.
779 (2012). Experimental evolution of multicellularity. *Proceed-*
780 *ings of the National Academy of Sciences*, 109(5):1595–1600.
- 781 Ray, T. S. (1996). Evolving parallel computation. *Complex Sys-*
782 *tems*, 10:229–237.
- 783 Real, E., Moore, S., Selle, A., Saxena, S., Suematsu, Y. L., Tan, J.,
784 Le, Q. V., and Kurakin, A. (2017). Large-scale evolution of
785 image classifiers. In Precup, D. and Teh, Y. W., editors, *Pro-*
786 *ceedings of the 34th International Conference on Machine*
787 *Learning*, volume 70 of *Proceedings of Machine Learning*
788 *Research*, pages 2902–2911, International Convention Cen-
789 tre, Sydney, Australia. PMLR.
- 790 Smith, J. and Szathmary, E. (1997). *The Major Transitions in Evo-*
791 *lution*. OUP Oxford.
- 792 Taylor, T., Bedau, M., Channon, A., Ackley, D., Banzhaf, W.,
793 Beslon, G., Dolson, E., Froese, T., Hickinbotham, S.,
794 Ikegami, T., et al. (2016). Open-ended evolution: per-
795 spectives from the oee workshop in york. *Artificial life*,
796 22(3):408–423.
- 797 West, S. A., Fisher, R. M., Gardner, A., and Kiers, E. T. (2015).
798 Major evolutionary transitions in individuality. *Proceedings*
799 *of the National Academy of Sciences*, 112(33):10112–10119.