

Resource Abundance Promotes the Evolution of Public Goods Cooperation

Brian D. Connelly
Michigan State University
East Lansing, MI 48824, USA
connel42@cse.msu.edu

Benjamin E. Beckmann
Michigan State University
East Lansing, MI 48824, USA
beckma24@msu.edu

Philip K. McKinley
Michigan State University
East Lansing, MI 48824, USA
mckinley@cse.msu.edu

ABSTRACT

Understanding the evolution of cooperation as part of an evolutionary stable strategy (ESS) is a difficult problem that has been the focus of much work. The associated costs of cooperation may lower the fitness of an organism below that of its non-cooperating counterpart, allowing the more fit organism to persist and outcompete the cooperator. Insight into these behaviors can help provide a better understanding of many aspects of the natural world, as well as provide future avenues for fighting disease.

In this study, we use digital evolution to examine how the abundance of a required resource affects the cooperative production of a public good in an adverse environment. Evolutionary computation is an excellent tool for examining these problems, as it offers researchers complete access to organisms and total control over their environment. We find that stable cooperation can occur in otherwise competitive environments at discrete levels corresponding to the availability of a required resource. When resource levels are low, organisms focus solely on competitive behaviors. However, once resource levels cross a critical threshold, cooperation persists in populations. Further, this cooperation occurs in patches, where it is most likely to benefit relatives. Finally, we find that in some cases this cooperative behavior allows organisms to increase their competitive abilities as well.

Categories and Subject Descriptors

H.4 [Information Systems Applications]: Miscellaneous;
F.1.1 [Computation by Abstract Devices]: Models of Computation—*Self-modifying machines*

General Terms

Experimentation, Theory

Keywords

Artificial life, cooperation, public goods, digital evolution, virtual biofilms, structured populations

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO'10, July 7–11, 2010, Portland, Oregon, USA.
Copyright 2010 ACM 978-1-4503-0072-8/10/07 ...\$10.00.

1. INTRODUCTION

Insight into the evolution of cooperative strategies is critical in gaining an understanding of many fundamental aspects of life, including the emergence of multicellularity [25] and the evolution of pathogenesis and antibiotic resistance [7]. Cooperative acts incur some cost on the actor while providing some benefit to a recipient. In public goods cooperation, this benefit is placed into the environment, where anyone may take advantage of it. Upon first glance, the traditional “survival of the fittest” view of evolution would select against such behaviors, as other organisms that do not pay this cost would have higher individual fitness.

One mechanism allowing selection to favor these costly cooperative acts is *kin selection* [10], whereby organisms seek to maximize their *inclusive fitness*, which considers both the fitness of that particular organism and the fitness of its kin, which share either common ancestry or behavior. Because these organisms have some genetic similarity, the reproductive successes of an organism’s kin benefit the proliferation of that organism’s genes as well. This mechanism is captured in Hamilton’s equation, shown in Equation 1, which states that cooperation will be selected for when the costs (c) of the act are less than the benefits (b) conferred upon the recipient multiplied by the genetic relatedness (r) between the pair [10]. Therefore, cooperation is more likely to occur in environments with spatial structure, where organisms tend to interact primarily with close relatives.

$$c < r * b \quad (1)$$

Due to the difficulties in quantifying the effects of interactions in higher organisms, as well as their use of additional mechanisms such as reciprocity and punishment, cooperation at this fundamental level is increasingly studied using populations of microorganisms such as bacteria and yeast [24]. The short generation times of these organisms make evolutionary-scale studies feasible. Computational models of these systems are also frequently used, providing researchers complete environmental control and access to each organism [6, 8, 19].

One area in which cooperative behaviors are critical is in the formation and maintenance of biofilms [15], which are aggregates of microorganisms connected within a matrix of extracellular polymeric substance (EPS) [7]. This “stickiness” allows organisms in biofilms to cohabit in larger densities and resist being flushed out of their environment. In addition, biofilms play a large role in antibiotic resistance, and can render bacteria 1,000 times less likely to be affected by antibiotics [22]. Because this self-produced EPS resource

is a public good in this system, selection can produce organisms that take advantage of its benefits without producing it themselves (*cheaters*). The presence of such cheaters has been observed to significantly weaken biofilms [20].

Indeed, this observation is the basis for current research in “anti-infective” treatments of bacterial infections. In one recent study, mutants that do not produce or respond to a cooperative signaling molecule were introduced into infections in mice, resulting in a significant decrease in virulence and a corresponding increase in host survival [21]. Similarly, the use of benign or more manageable strains that can outcompete pathogenic strains has been considered [4]. As a further step, researchers have also proposed employing “Trojan Horse” techniques, whereby a competitive mutant strain containing beneficial behaviors is introduced [4]. In order for such treatments to be effective (and not harmful to the host), an intimate knowledge of the social interactions occurring in bacterial communities is needed.

These interactions depend not only on the organisms and communities themselves, but also on the environment in which they live. Often, the costs and benefits are driven by environmental factors such as disturbance, spatial limitations, and the availability of resource, which is examined in this work. Environmental disturbance, such as grazing by a higher trophic level species, can both promote and hinder cooperation. It has been observed that cooperation is maximized at intermediate levels of disturbance, where the balance between costs and benefits is positive [2]. This is because at low disturbance, there is less need for cooperative acts, so their associated benefits are small. At high disturbance, organisms are forced to focus on self-preserving behaviors, making the costs of cooperation prohibitively high.

In addition, the density of populations supported by the environment may drive selection for cooperation. Specifically, as groups grow to carrying capacity, increased competition among kin may hinder the evolution of cooperation through increased costs [12, 18]. In facilitation (interspecific cooperation), it is hypothesized that cooperation may allow organisms to increase their “realized niche”, thus enabling larger and more diverse communities [5]. Here, cooperative acts allow a species to survive in environments that would otherwise be uninhabitable, either environmentally or through competition.

The availability of a required resource often plays a large role in determining evolutionary outcomes. As resource becomes more abundant, the relative costs of the cooperative act decrease, allowing organisms to perform cooperative acts in addition to their self-preserving and competitive activities. When that resource becomes more scarce, however, organisms may be forced to focus solely on their own interests. Wet bench experiments have shown this relationship [3] as well as a link between resource availability and diversification [9].

The study presented here examines this relationship between available environmental resource and cooperation through the production of a public good. Populations of digital organisms evolved in digital environments with different resource availabilities, and we examined their ability to persist in an adverse environment through cooperation. In agreement with previous studies of natural organisms [3], we find that the frequency of cooperative organisms increased with available resource. Further, our study revealed two interesting features of this cooperation not reported previously.

First, that populations cooperate at discrete levels—those in rich environments cooperate with similar levels of effort regardless of resource abundance, while those in poor environments cooperate consistently less. Secondly, we find that success within these two distinct groups remains closely tied to competition. These findings provide additional insight into the environmental conditions under which cooperation occurs and may be used to help drive the development of treatments that target similar behaviors in disease.

2. METHODS

For this study, we used the Avida software platform to study evolution in populations of self-replicating digital organisms [17]. As depicted in Figure 1, each organism consisted of a circular list of assembly-like instructions (its “genome”), which were executed sequentially on virtual hardware allocated to that organism. This virtual hardware comprised three 32-bit registers, two stacks, and a CPU, which had four heads that controlled the execution flow of the organism’s genome and aided in self-replication. The instruction set available to organisms was Turing complete, so organisms could feasibly perform any possible calculation. Each organism resided independently in its own *cell* in a 100x100-cell bounded lattice, where each cell had eight neighbors. When an organism replicated, it created a copy of itself and placed this offspring into a random cell in the neighborhood, replacing any organism residing in that cell. By choosing neighboring cells, we limited dispersal and increased the likelihood that interactions would be among kin. Because an organism can be killed by new offspring, these environments presented a pressure to replicate quickly. Typically in Avida, organisms can replicate more often by shortening their genomes through mutations or by executing instructions more quickly.

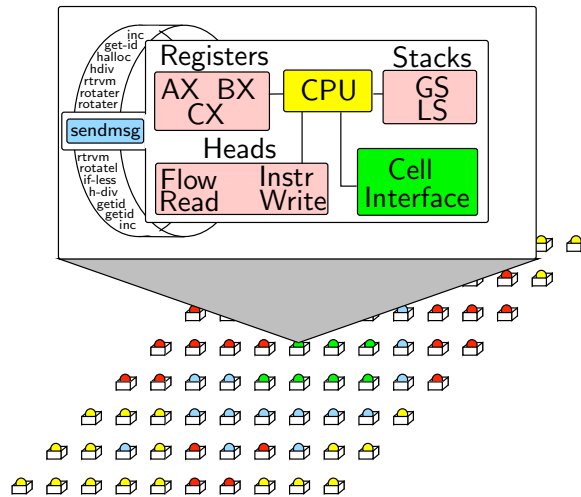


Figure 1: Representation of a digital organism (top) in an Avida population (bottom) [11]

As in other studies using Avida [14], organisms competed in this environment by completing logic *tasks*, which conferred a *merit* bonus. An organism with more merit received proportionally more CPU cycles, enabling it to execute its

Table 1: Logic tasks that could be completed by organisms. Upon completion, an organism was given the corresponding merit reward. The completion of OR NOT was not rewarded, but resulted in the production of beneficial public good at the organism’s cell.

Task	Input	Output	Merit Bonus
NOT	A	$\neg A$	2
NAND	A, B	$\neg(A \wedge B)$	2
AND	A, B	$A \wedge B$	4
OR	A, B	$A \vee B$	8
AND NOT	A, B	$A \wedge \neg B, \neg A \wedge B$	8
NOR	A, B	$\neg(A \vee B)$	16
XOR	A, B	$(A \wedge \neg B) \vee (\neg A \wedge B)$	16
EQU	A, B	$(A \wedge B) \vee (\neg A \wedge \neg B)$	32
OR NOT	A, B	$A \vee \neg B, \neg A \vee B$	0

genome faster and thus spread in the environment. These rewards were given at most once per task completed. In order to complete a task, an organism first had to execute the IO instruction, which placed randomly-generated numbers into its registers. The organism was then required to execute the instructions necessary to complete the task, place the result in the correct output register, and issue an additional IO instruction. A full list of the tasks available and their rewards is listed in Table 1.

In addition, an organism required a sufficient amount of a resource in order to complete each task. This resource flowed evenly into each cell in the environment and dissolved at a rate of 1% per *update* (Avida’s unit of time, further explained later). In this study, the inflow rates were varied, ranging from extremely limiting (equilibrium of 1 unit per cell) to effectively unlimited (equilibrium of 100 units per cell), where organisms were unable to complete tasks quickly enough to deplete the resource. For this work, each task consumed 1 unit of this resource. However, tasks had varying metabolic costs, or the number of instructions required for their completion, as is shown in Figure 2.

(a) IO	(b) IO	(c) IO	(d) IO
IO	IO	IO	IO
nop-c	nop-c	nop-c	nop-c
nand	nand	push	push
IO	nand	nand	nand
	IO	swap	swap
		nand	nand
		swap	swap
		nop-a	nop-a
		pop	pop
		nand	nand
		swap	swap
		nop-c	nop-c
		nand	nand
		IO	push
			pop
			nop-c
			nand
			IO

Figure 2: Sample genome segments required to complete several tasks given adequate resources. (a) The NAND task can be completed with 5 instructions, while (b) OR NOT requires 6. More difficult tasks such as (c) XOR and (d) EQU require 15 and 19 instructions, respectively.

An organism replicated by first allocating additional space in their genome to create a copy of itself. It then copied instructions from its genome into this new area line by line. Once this process was completed, the genome was cleaved, and the new organism was placed into a neighboring cell. As an organism copied its genome, instructions were modified with a probability of 0.0075, and as the organism divided, instructions were added or deleted with a probability of 0.05. Populations were seeded with a single organism capable only of replication. Any additional behaviors were evolved through mutations introduced during the replication process.

2.1 Virtual Biofilms

We adopted the concept of a virtual biofilm from previous work [6] as the model for our experiments. As in nature, an organism’s membership in a virtual biofilm provided it with some benefit that it could not easily obtain in a free or planktonic state. In our digital environment, being in a biofilm allowed organisms to survive a periodic “kill” event similar to grazing by higher organisms in nature.

To create a virtual biofilm, organisms completed the OR NOT task, which resulted in one unit of a public good being placed into the environment at their cell. This cooperative resource can be likened to EPS, enzymes, or other beneficial products commonly observed in biofilms. Unlike the other available tasks, OR NOT did not confer a merit reward (see Table 1); however, the product that it produced could enable that organism or its kin to avoid being killed, as discussed shortly. We therefore define cooperators as organisms who performed this task at some point during their lifetime, while cheaters focused solely on rewarded tasks.

A kill event periodically selected a focal cell at random, and each cell in a neighborhood comprising cells within a 5-cell radius (121 cells in total, or 1.21% of the environment) was examined. If the mean level of public good in these cells was below a threshold of 3 units, all organisms residing in that region were killed. This public good both decayed at 1% and diffused into neighboring cells at 1% per update. In this environment, cheaters could exploit both the diffusion of public good from neighboring cooperator cells as well as the kill event’s consideration of the mean public good level within a neighborhood to benefit from the production by neighbors while refraining from producing the public good.

Starting from a single seed organism, populations evolved in environments with different resource abundances for 50,000 updates. These abundances corresponded to equilibrium cell levels of 1, 2, 4, 8, 10, 20, 40, 60, 80, and 100 units per cell. During each update, each organism executed 30 instructions on average. Those organisms with higher merit executed proportionally more instructions per update than those with less. Between approximately 6,000 and 8,000 generations were observed in each of these runs. For each configuration, 30 independent populations were evolved. Each of these runs started with a different random seed, allowing each to explore different evolutionary paths. The kill event was executed every update starting with update 1,000, which then gave organisms a 1.21% chance of residing within a selected region.

3. RESULTS

As observed in studies of bacterial populations [3], the fraction of cooperators in our digital populations increased

with increasing resource abundance. In each of the different environments examined in this study, however, some populations were unable to escape the kill event in the early stages. This occurred when a population was small and had not had sufficient time to evolve public good production. The results reported here correspond with the populations that were able to run for the allotted time. For each data set shown in Figures contained in this section, the number of completed runs is indicated in each plot. Additionally, plots show data from a representative subset of the completed runs in order to improve readability.

Improved Response to the Kill Event—The fractions of organisms selected and killed by the kill event is shown in Figure 3. Here, organisms in high resource environments were able to produce enough public good so that approximately 70-80% of organisms successfully avoided being killed, while organisms in lower resource environments produced significantly less public good and were killed approximately 70% of the time. Interestingly, there is no clear relationship between the amount of resource available and the fraction of organisms that were able to avoid being killed. In fact, there is little difference in kill rates among environments with stable resource levels of 40 units per cell and above (in 3 out of 4 environments, however, these differences are significant; $p < 0.002$). There is also little difference among those in environments with levels of 10 units per cell and below (3 out of 5 runs were significantly different; $p < 10^{-6}$). These two qualitative groups, however, are significantly different ($p < 10^{-6}$) from each other, which indicates a clear difference in behaviors. This could be indicative of a critical resource level, below which cooperation is not favored, but above which cooperation occurs at relatively similar levels. This observation matches Hamilton’s equation (Equation 1), which indicates whether or not cooperation can occur under a given set of conditions. In this case, organisms have found it beneficial to cooperate in environments with sufficient resource; however, an increase in cooperative behavior has not proven to be more beneficial, so all populations maintain this behavior with the minimum required investment.

Increasing Levels of Public Good Produced—The observed survival in these adverse environments is due to above-threshold levels of public good existing in selected regions. The amount of this product created within these populations is shown in Figure 4. The populations that we observed to be most successful at preventing being killed were able to produce a mean per-cell level of resource well above the 3-unit threshold, while those in lower resource environments produced significantly less. As before, the differences within each of these two groups are minimal, indicating that there appears to be a level of required resource below which the cooperative act is too costly. Populations in the environment with a resource level of 20 units per cell were significantly different than these groups ($p < 10^{-6}$), indicating their level may be nearer to this critical value.

Increased Abundance of Cooperators—In populations that produce levels of public good well above the kill threshold, there is a large potential for cheaters to emerge that take advantage of this surplus. Figure 5 plots the fraction of cooperators present in the different populations. We define cooperators as organisms that perform the public-good-producing task at least once during their lifetime, while cheaters focused solely on rewarded tasks. We see that pop-




Figure 3: Fraction of organisms in different resource environments residing in the selected region that were killed by the kill event. As available resource increased, more organisms resided in cells with above-threshold public good levels, indicating that organisms were completing more cooperative tasks. Data are averaged over all completed runs (number is indicated in legend).




Figure 4: Per-cell level of public resource in runs with differing resource levels in environments with the neighborhood kill event. Organisms in high resource environments produced levels of public resource per cell that were above the kill threshold of 3.0. Data are averaged over all completed runs.

ulations are comprised of more cooperators when resource is abundant. This could be an indication of the decreased relative costs of cooperation in high resource environments: Organisms continue to compete through their completion of rewarded tasks, while cooperation remains as a common feature. Although a significant increase in the fraction of cooperators is observed in rich environments, cooperators remained less frequent than cheaters.




Figure 5: Fraction of cooperators in populations with different resource levels. Cooperators are defined as organisms that completed at least one OR NOT task, which produced public resource. Data are averaged over all completed runs.

Skewed Distribution of Public Goods—One explanation for both the abundance of cheaters and the production of public good well above the threshold is that cooperators are producing enough resource to support their entire populations. In this case, we would expect to see a fairly uniform distribution of this resource throughout the environment. Figures 6 and 7, however, show that this is not the case, and that both cooperators and public good abundance exist mostly in patches, which is common among cooperative systems with spatial structure [16]. Here at update 45,000, the mean per-cell level is above the kill threshold, yet 36% of cells are below this threshold due to the 58% cheaters in this population. These patches are continually attacked by cheaters, who reproduce more quickly and take over cells. However, once large clusters of cheaters form, they become susceptible to the kill event when the public resource decays and they can no longer take advantage of cooperating neighbors. These newly-vacated cells then become available to cooperators, and these oscillations continue. Figure 8 shows these dramatic oscillations between cooperators and cheaters within a selected 5x5-cell region during evolution in a typical high-resource environment. Although mean per-cell resource levels continue to rise in resource-rich environments, spatial constraints prohibit cheaters from invading and simply surviving from the resource produced earlier. Videos capturing these dynamics are available on our website¹.

¹<http://www.cse.msu.edu/thinktank>




Figure 6: Distribution of public resource in a typical high-resource environment at update 45,000. White indicates cells with above-threshold levels while black indicates cells with below-threshold levels. Videos showing this distribution over an entire run are available on our website¹.

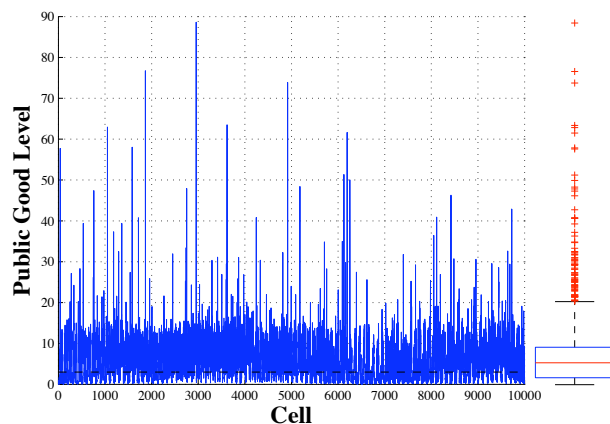


Figure 7: Distribution of public resource among cells in a typical high-resource environment at update 45,000. We see that a small number of cells have public resource levels much higher than others. The inclusion of such cells in any target area is likely to greatly assist in preventing the kill event from succeeding. In this 100x100-cell environment without periodic boundary conditions, cells in the first row were labeled 0-99, cells in the second 100-199, and so on.




Figure 8: Oscillating fractions of cooperators and cheaters present within a selected 5x5-cell region in one sample population in a high-level resource environment.

Overall Increase in Task Completion—In addition to avoiding the kill event, organisms also competed for space in these environments. Therefore, they faced pressure to complete rewarded tasks, enabling them to execute and reproduce more quickly. Figure 9 plots the number of rewarded tasks completed in different environments. Here we see that organisms in low-resource environments completed these tasks roughly in proportion to the level of resource available, indicating that this resource was limiting. Those organisms in high-resource environments, however, completed roughly the same number of tasks regardless of the level in their environments. In these environments, resource was not a limiting factor. Figure 10 shows the number of public-good-producing tasks completed. Here, the trends are similar among high-resource environments, and populations in low resource environments show a consistently-low cooperative effort. Overall, populations completed this cooperative task an order of magnitude less often than the rewarded tasks, indicating that competition remains the primary driving force in these environments.

Finally, Figure 11 plots the total number of tasks completed. We can see that in all populations, the competitive behaviors make up a much larger portion of the total number of tasks completed. However, those populations with intermediate resource availability and where cooperation occurs (e.g. equilibrium resource levels of 20 and 40) seemed to complete more tasks overall than would be predicted if the trends seen among lower-resource environments continued. This could indicate that cooperation enabled these populations to be more successful at competing in the task environment.

4. CONCLUSIONS

This study has demonstrated in digital populations a phenomenon reported in natural organisms: the relationship between the availability of a required resource and the fraction of cooperation observed in populations. We saw that




Figure 9: Number of rewarded tasks performed by organisms in environments with different resource levels. Organisms in low-resource environments were limited, while organisms in high-resource environments were able to continually improve. Data are averaged over all completed runs.




Figure 10: Number of cooperative, public-good-producing tasks performed by populations in environments with different resource levels. Data are averaged over all completed runs.




Figure 11: Total number of tasks, both rewarded and cooperative, performed by organisms in environments with different resource abundances. Data are averaged over all completed runs.

as available resource increased, populations produced more public good, enabling them to stave off being killed by a periodic event. Interestingly, however, we observed this relationship between resource and fraction of cooperation to be discrete. Above certain levels, organisms in all resource environments cooperated at a consistent level, while cooperation was consistently low below this critical level. We speculate that this critical level is where the costs of the cooperative behavior become less than the benefits provided. Further investigation of environments with resource levels near this dramatic change in behavior could yield insight into where this threshold exists and the true cost and benefit structures in these worlds. The presence of this critical resource level lends itself nicely to future investigation into anti-infective treatments. Following [21], this critical level could potentially be used to determine the amount of non-cooperative mutants to introduce to an infection in order to fall below the critical resource level and suppress cooperation and its negative effects.

This work also measured the competitive, rewarded tasks performed by organisms in addition to the cooperative ones, which is rarely done in computational models of these behaviors. We found that even in populations in which stable cooperative behaviors existed, organisms performed significantly more competitive tasks than cooperative ones. This result indicates that even in environments where competition dominates, cooperation can still emerge and may in fact enable a species or strain to become more competitive than a non-cooperating one. We observed this as cooperative populations in intermediate resource environments completed more tasks than would be expected based on other populations.

We also observed large surpluses in the level of public good in populations with significant fractions of cooperators, potentially supporting a large fraction of cheaters. We found, however, that the majority of this public good resided in clusters of cells which corresponded to clusters of coopera-

tors. Although this surplus made these cooperators susceptible to invasion by cheaters, we found that cooperation was still able to persist in the populations.

In considering the relationship between cooperation and resource availability, many additional questions remain. We have seen the effect at the population level, but how does the evolution of an individual's behavior change in different environments? Perhaps some level of restraint from cheating can be evolved at the genetic level as has been observed [23]. Additionally, plasticity [13] can play a large role in the investment an individual makes in cooperation based on the state of the environment and may significantly alter the percentage of cooperators that exist in populations, given a fluctuating environment.

Similarly, if the required resource is non-renewable and will eventually be depleted, it would be interesting to investigate the conditions under which populations evolve restraint in its consumption or can instead shift their dependence to another resource [1].

5. ACKNOWLEDGMENTS

We are grateful for the help and insight provided by Christopher Waters, Tracy Teal, and Luis Zaman. This work was supported in part by National Science Foundation grants CNS-0915885, CCF-0820220, and CNS-0751155; by U.S. Army Grant W911NF-08-1-0495; and by a Quality Fund Grant from Michigan State University.

6. REFERENCES

- [1] Z. Blount, C. Borland, and R. Lenski. Historical contingency and the evolution of a key innovation in an experimental population of *Escherichia coli*. *Proceedings of the National Academy of Sciences*, 105(23):7899–7906, 2008.
- [2] M. Brockhurst, A. Buckling, and A. Gardner. Cooperation peaks at intermediate disturbance. *Current Biology*, 17(9):761–765, 2007.
- [3] M. Brockhurst, A. Buckling, D. Racey, and A. Gardner. Resource supply and the evolution of public-goods cooperation in bacteria. *BMC Biology*, 6(1):20, 2008.
- [4] S. Brown, S. West, S. Diggle, and A. Griffin. Social evolution in micro-organisms and a Trojan horse approach to medical intervention strategies. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1533):3157–3168, 2009.
- [5] J. Bruno, J. Stachowicz, and M. Bertness. Inclusion of facilitation into ecological theory. *Trends in Ecology & Evolution*, 18(3):119–125, 2003.
- [6] B. D. Connelly and P. K. McKinley. Evolving social behavior in adverse environments. In *Proceedings of 10th European Conference on Artificial Life*, 2009.
- [7] J. Costerton, P. Stewart, and E. Greenberg. Bacterial biofilms: a common cause of persistent infections. *Science*, 284(5418):1318–1322, 1999.
- [8] T. Czárán and R. F. Hoekstra. Microbial communication, cooperation and cheating: Quorum sensing drives the evolution of cooperation in bacteria. *PLoS ONE*, 4(8):e6655, 2009.
- [9] A. R. Hall and N. Colegrave. How does resource supply affect evolutionary diversification? *Proceedings of the Royal Society B*, 274(1606):73–78, 2007.
- [10] W. D. Hamilton. The genetical evolution of social behaviour. I,II. *Journal of Theoretical Biology*, 7(1):1–52, 1964.
- [11] D. B. Knoester and P. K. McKinley. Cooperative network construction using digital germlines. In *Proceedings of the*

- Genetic and Evolutionary Computation Conference*, pages 217–224, 2008.
- [12] R. Kümmerli, A. Gardner, S. West, and A. Griffin. Limited dispersal, budding dispersal, and cooperation: an experimental study. *Evolution*, 63(4):939–949, 2009.
- [13] R. Kümmerli, N. Jiricny, L. Clarke, S. West, and A. Griffin. Phenotypic plasticity of a cooperative behaviour in bacteria. *Journal of Evolutionary Biology*, 22(3):589–598, 2009.
- [14] R. E. Lenski, C. Ofria, R. T. Pennock, and C. Adami. The evolutionary origin of complex features. *Nature*, 423:139–144, 2003.
- [15] C. Nadell, B. Bassler, and S. Levin. Observing bacteria through the lens of social evolution. *Journal of Biology*, 7(7):27, 2008.
- [16] M. Nowak and R. May. Evolutionary games and spatial chaos. *Nature*, 359(6398):826–829, 1992.
- [17] C. Ofria and C. Wilke. Avida: A software platform for research in computational evolutionary biology. *Artificial Life*, 10(2):191–229, 2004.
- [18] T. G. Platt and J. D. Bever. Kin competition and the evolution of cooperation. *Trends in Ecology & Evolution*, 24(7):370–377, 2009.
- [19] F. Prado and B. Kerr. The evolution of restraint in bacterial biofilms under nontransitive competition. *Evolution; international journal of organic evolution*, 62(3):538–548, 2008.
- [20] P. B. Rainey and K. Rainey. Evolution of cooperation and conflict in experimental bacterial populations. *Nature*, 425:72–74, 2003.
- [21] K. P. Rumbaugh, S. P. Diggle, C. M. Watters, A. Ross-Gillespie, A. S. Griffin, and S. A. West. Quorum sensing and the social evolution of bacterial virulence. *Current Biology*, 19(4):341–345, 2009.
- [22] P. S. Stewart and J. W. Costerton. Antibiotic resistance of bacteria in biofilms. *The Lancet*, 358(9276):135–138, 2001.
- [23] G. Velicer. Evolution of Cooperation: Does Selfishness Restraint Lie within? *Current Biology*, 15(5):173–175, 2005.
- [24] S. West, A. Griffin, A. Gardner, and S. Diggle. Social evolution theory for microorganisms. *Nature Reviews Microbiology*, 4(8):597–607, 2006.
- [25] J. B. Xavier, E. MartinezGarcia, and K. R. Foster. Social evolution of spatial patterns in bacterial biofilms: When conflict drives disorder. *The American Naturalist*, 174(1):1–12, 2009.