

Predator-Prey Dynamics and the Red Queen Hypothesis: Putting Limits on the Evolutionary Arms Race

Ted Carmichael, UNC Charlotte and TutorGen, Inc.
Mirsad Hadzikadic, UNC Charlotte

Abstract

Computer simulations of complex food-webs are an important tool for deepening our understanding of these systems. Yet most computer models assume, rather than generate, key system-level patterns, or use mathematical modeling approaches that make it difficult to account fully for non-linear dynamics. In this article we present a computer simulation model that addresses these concerns by focusing on assumptions of agent attributes rather than agent outcomes. Our model utilizes the techniques of Complex Adaptive Systems and Agent-Based Modeling so that system-level patterns of a general ecosystem emerge from the interactions of thousands of individual simulated agents. This methodology has been validated in previous work by using this general simulation model to replicate fundamental properties of an ecosystem, including: (1) the predator-prey oscillations found in Lotka-Volterra; (2) the “stepped pattern” of biomass accrual from resource enrichment; (3) the Paradox of Enrichment; and (4) Gause’s Law. In this work we explore further the fundamental properties of this generative model in the context of the Red Queen Hypothesis, also referred to as the “arms race” between antagonistic species, e.g. predators and prey. We find that improvements in the competitive landscape for a single entity in a predator species does not generally confer a benefit on the predator species as a whole, and may even be detrimental to the predator population. This non-intuitive result is shown through two methods of adjusting the predators effectiveness in consuming prey. We further explore this idea by explicitly accounting for individual entity’s energy requirements, and also allowing evolutionary adaptation for an effectiveness / energy trade-off.

1.0 Overview

The literature on marine and terrestrial ecosystems is long and varied, encompassing both theoretical models (e.g.: Grimm, 1999; DeAngelis & Mooij, 2005) and empirical surveys (e.g.: Christensen et al., 2003; Frank et al., 2005). Some significant differences between model results and real-world surveys have persisted for years, and it has been difficult identifying fundamental principles relative to the many complicating factors that can be found in existent ecosystems. For example, in the early 1980s Oksanen et al. examined multiple trophic levels in a predator-prey system using mathematical models, in order to determine whether species population (bio-mass) is fundamentally controlled by resources – as was the conventional wisdom at the time – or dominated by predation (Oksanen et al., 1981). In describing this work, Power states that these models produce “a stepped pattern of biomass accrual” (Power, 1992); Brett and Goldman further characterize the Oksanen et al. results, saying that “In food webs with an odd number of trophic levels, increases in primary production should lead to increased biomass for odd-numbered trophic levels and no change in biomass for even-numbered trophic levels. Conversely, in food webs with an even number of trophic levels, increases in primary production should lead to increased biomass for even-numbered trophic levels and no change in biomass for odd-numbered trophic levels” (Brett & Goldman, 1997).

More recently, researchers have begun to ask, not which process (bottom-up resources vs. top-down predation) dominates overall, but rather how these two forces interact under different spatial and temporal scenarios (Reid et al., 2000). Many scientists think that this new approach will help reveal “how resilient food webs are to the combined effect of resource-mediated and predator-mediated forces” (Casini et al., 2008). This suggestion leads naturally to the view that ecosystems can be productively studied as Complex Adaptive Systems; that is, systems that are characterized by their emergent properties, self-organization, and non-linear dynamics (Holland, 1992; Gell-Mann, 1994; Allesina & Pascual, 2007; Brose & Dunne, 2010; Valdovinos et al., 2010). The field of Complex Adaptive Systems (CAS) recognizes that these systems are generally robust and flexible, consisting of multiple negative feedbacks that produce one or more “basin(s) of attraction,” i.e., the emergence of resilient, system-level patterns. Indeed, it is the identification of the key feedbacks and their contributions to system resilience that is the ultimate goal of ecological research.

Advances in computing technology have allowed for more robust *in silico* simulations that can model these complex ecosystem dynamics in ways previously intractable. In order to address certain limitations of purely mathematical models (DeAngelis & Mooij, 2005), which generally simplify individual variation and spatial representation, and represent global properties in a top-down manner, we have created a general CAS model of a marine ecosystem. Central to a more flexible, Agent-Based Modeling (ABM) approach (sometimes called Individual-Based Modeling, or Agent-based Complex Systems) is that the outcomes of the simulation are

generated in a bottom-up design process, rather than via top-down constraints (Grimm et al., 2005). It is important to note that the phrase “bottom-up” design has a somewhat different meaning in the CAS literature than the similarly-worded “bottom-up” forces referred to in marine and terrestrial ecosystem literature. A CAS-based ABM implies that the system-level patterns, such as population growth, aggregate predation rates, etc., are *generated* from the bottom up, rather than assumed as a “top down” constraint on the modeled system (see e.g., Epstein, 1999). That is, the system-level patterns are emergent properties that arise from the interactions of the autonomous agents that comprise the simulated system. Thus, the method of computer simulation relies on assumptions of agent *attributes* rather than agent *outcomes*. (The similar phrasing in the ecosystem literature is unfortunate; one might come across, e.g., a discussion of the controlling forces in constraining a particular population: i.e., “bottom-up” forces of resource availability vs. “top-down” predation. This is an entirely different context than “bottom-up” emergence in an ABM.)

Here we present a general ABM marine ecosystem with a focus on key phenomena in population dynamics in the context of the Red Queen Hypothesis, also referred to as the “arms race” between antagonistic species, e.g. predators and prey. This model has been validated in previous work by replicating fundamental properties of an ecosystem, including: the predator-prey oscillations found in Lotka-Volterra; the “stepped pattern” of biomass accrual from resource enrichment found in Oksanen, et al.; the Paradox of Enrichment; and Gause’s Law (Carmichael & Hadzikadic, 2013). In this work we will extend our understanding of found in Oksanen, et al., by considering how these patterns of biomass accrual change when the predator population becomes more (or less) efficient at catching prey. We will also consider these changing patterns in the context of the Red Queen Hypothesis and how evolutionary pressures may be aligned with limits on the Red Queen effect, in ways that could prevent an escalating arms race between predators and prey.

2.0 Background

2.1 Previous Results

In previous work we validated the results of Oksanen, et al. (1981), by showing that our agent-based model (ABM) exhibits the “stepped pattern of biomass accrual” found in simple predator-prey dynamics. Briefly, in a three-trophic-level simulation (food, prey, predators), increasing the food available to the prey will cause the *predators* to increase in population but, surprisingly, does not change the population size of the prey themselves. That is, even though the prey are able to eat more – and thus reproduce faster – with an increase in food, the predators immediately match this change by consuming the extra prey that are now above the equilibrium level. Add food at an even faster rate and the predator population will grow to an even higher population level, and continue to consume all the extra prey. Thus, when resources are increased to the prey, the prey population can eat faster, and reproduce faster, but they are also consumed faster, while the predators alone increase in population size. This non-intuitive result matches the predictions found in Oksanen, et. al.

Another consequence of this dynamic is that the predators do not increase their equilibrium consumption rate at the *per capita* level. Predator consumption rate does change when they are out-of-equilibrium, but only until the new population level is reached; at that point the *per capita* consumption rate has returned to the predators' equilibrium level of consumption. Figure 2 illustrates the increased population size for the predators based on an increased resource level for the prey, as well as the concurrent reduction in the prey's average age.




Figure 1: Population growth of predators (purple) based on increased resources to the prey population (left) and the concurrent change in the average age of the prey (right). The simulation was run for 6000 time steps; resource rate was increased after 2000 steps (a) from 0.06 to 0.12, and again at step 4000 (b) from 0.12 to 0.18. Adapted from Carmichael & Hadzikadic, 2013.

2.2 Model Description

This generative ABM is purposefully kept as simple as possible, in order to determine baseline properties and consequences of the interacting populations with as few complicating factors as possible. In the ecological literature it is noted that the outcomes of this model are rarely if ever found in the real world. This includes Oksanen's results, Gause's Law, the Paradox of Enrichment, and even the Lotka-Volterra model of predator-prey

dynamics. This incongruity is acceptable, even expected, as the real world attributes of various species are much more complicated and nuanced than their simulated counterparts in our ABM. However, it is incredibly difficult to tease out which of these additional attributes are important and which are inconsequential. A generative model allows us to add complexity one layer at a time and determine, from a simulation standpoint, which are important, and might therefore give us a better understanding of the many different food webs that are found in nature.

Therefore, the assumptions encoded in this model are relatively straightforward. The food is grown on each patch based on a “slider” in the user interface; as the slider is increased, the amount of food grown per patch, per simulation time-step increases in a linear fashion. The prey perform three basic functions during each turn: they move at random, eat if there is food present, and increase their age. They also reproduce asexually as a linear function of how much food they have consumed. The threshold for reproduction can be controlled by a slider; however, this linear relationship can not be altered in this simulation. The predators are exactly the same: they move at random, eat if there is prey available, and increase in age during each turn. Note that both the prey and the predators have a slider labeled “turns per tick.” This allows for additional control and experimentation, in that the number of turns for each species during each simulation time-step can be altered, which changes the number of actions performed each “tick” relative to the other populations. The consequences of changing this controller is discussed below. As with the prey, the predators reproduce asexually as a function of the total amount of food (in this case, prey) that is consumed. In some experiments there is also a “top predator” that consumes the predator in the same method that the predator consumes the prey. This changes the dynamics of all the populations in various ways; however, only the simple three-level model is discussed here. All models were created using the NetLogo modeling environment (Wilensky, 1999).

2.3 The Red Queen Hypothesis

The Red Queen Hypothesis was first introduced in 1973 (Van Valen) and expresses the idea of an “arms race” between antagonistic species, such as predators and prey in a common ecosystem. Given the intuitive benefits of increased efficiency for a member of one species, it seems likely that an advantage in the phenotype would ensure that the related genotype would more likely survive and spread on evolutionary time scales throughout the population. However if, for example, the predators become better hunters, then there is assumed to be subsequent pressure on the prey to also adapt, in order to better survive. Once the prey adapt better survival techniques, then the predators adapt again, and so on.

Given this theory the question naturally arises: why do species not continually increase in efficiency? There must be some mechanism that reduces or even eliminates continued evolutionary advantage. The first and most intuitive answer is that there is a cost associated with efficiency increases and, at some point along a continuum this cost is greater than the additional benefit. Related to this idea is the natural diversity in abilities across the prey species; in particular, individual prey that are very young, very old, or sick might generally have less ability than those in their prime. If this is the case, then a predator individual might have to work much harder for only a slight increase in the number of prey who are susceptible to predation.

More subtly, there could also be certain predation strategies that work against continued spread of highly effective genes in the population. For example, many predator species share the results of the hunt, which raises the possibility that the so-called “free rider” problem is a limiting factor, by conferring a benefit to a diversity of predator individuals and not just the most effective hunters. This reduces the gains for effective evolutionary adaptations and therefore increases the cost-to-benefit ratio associated with such gains.

In the present work the model described in section 2.2 is expanded to consider two alternative ways to increase (or decrease) the efficiency of the predator population. The outcomes of these two changes are then compared and contrasted in terms of the effects on the population as a whole, for both predators and prey. The second set of experiments expands the model even further, to include an explicit “energy” requirement for individual predators to move and hunt, and also a mechanism for evolving efficiency, but also with a concurrent change in the energy expenditure required, so that an increase in hunting effectiveness requires a greater expenditure in energy.

3.0. Experimental Design – First Experiment

Here we consider two different methods for increasing (or decreasing) the efficiency of the predator population. As noted in section 1.2, the number of “turns per tick” for the predators can be changed, which has consequences for both predator and prey populations. If this number is reduced, then the predator population as a whole has fewer actions relative to the prey for a given period of time. Another control for changing the attributes of the predator population has been added to the model called the “predator success rate.” In the baseline model if a predator finds prey on its current patch it will eat one hundred percent of the time. With this new controller there is now a chance that the predator will “miss” the prey, or that the prey will escape.

Intuitively these two different methods of controlling the predator population would seem to have very similar effects; both act to reduce the effectiveness of the predator population. If each predator has, for example, half as many actions per time-step, that would seem to be similar in effect as if each predator misses its prey half the time. However, as shown below, this turns out not to be the case.

3.1. Results – First Experiment

In the first case, where the number of turns per tick is reduced by one-third, the new equilibrium result is the predator population is higher, also by approximately one-third (Figure 2). Interestingly, none of the other monitored outcomes change: not the prey population size, nor the prey consumption rate or average age, nor the predators' relative consumption rate and average age. Note that for the predators, changing the “turns per tick” in reality does change their average age and consumption rate when compared to the prey. However, the average predator still consumes the same number of prey over its own lifetime; reducing the “turns per tick” by one-third simply increases the predator lifetime by one-third, because a unit of “lifetime” is advanced during each turn, not during a simulation time-step.

In the second case, we reduce the effectiveness of each predator, so that approximately one-third of the time the predator will “miss” the local prey. In this situation, however, we find quite a different result. The predator population does not decrease at all for being less effective hunters, but this does allow the prey population to increase, by approximately one-third, and the prey average age also increases by a similar amount (Figure 3). The predator consumption rate is unchanged, but the prey consume food approximately one-third less often which, perhaps paradoxically, reduces the total available food in the system (but by approximately one-half).




Figure 2: Population and average age changes as predators' "turns per tick" is reduced by one-third.

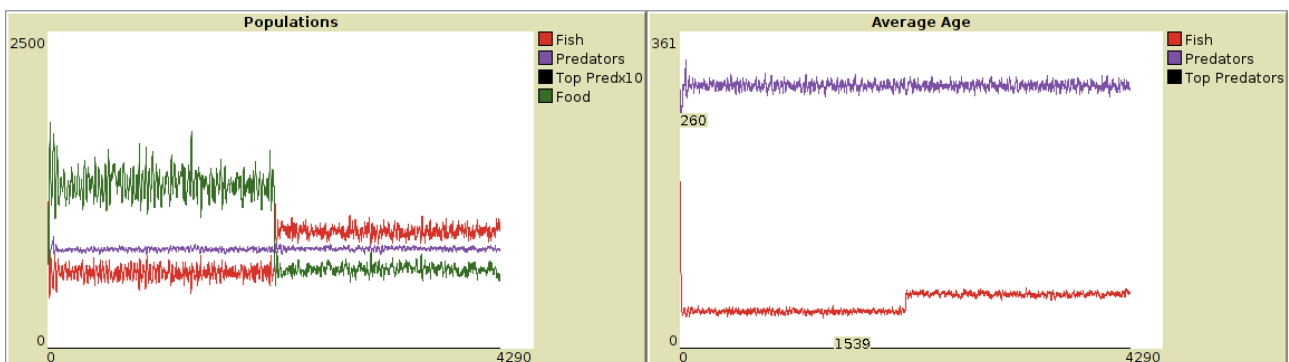


Figure 3: Population and average age changes as predators' "success rate" is reduced by one-third.

The first case above can be interpreted as a slowdown in the metabolism of all the individual predators; their lifetime is longer relative to the prey, but they don't consume any more prey across that lifetime. It is as if they have more resources available during a simulation time-step, and therefore the predator population numbers benefit from this effect. The second case shows that even if all the predators become less effective hunters, by missing their prey one-third of the time, their total population does not suffer in any way. Conversely, if the predators improve their hunting success, the prey suffer both individually and collectively, living a shorter lifespan and displaying a concurrent drop in total population size. If left unchecked continued predation improvements would conceivably drive the prey to extinction. And so the next set of experiments considers what mechanism is sufficient to prevent this outcome.

4.0. Experimental Design – Second Experiment

Using the same generative model, “energy” is added as an explicit constraint on the individual predators. Recall that in the baseline model each predator tracks how many prey have been consumed, and they will generate a new predator (asexual reproduction) once a threshold is reached. This could be considered an accounting of energy accumulation among the agents. Energy is therefore added to this model by imposing a small “movement cost” for the predators, reducing this accumulation by a small amount each turn.

Each predator is instantiated with a “success rate” of 0.50; however, now each predator has a random chance to mutate its own success rate in a way that is coupled to movement costs. If the success rate is mutated higher, then movement costs increase; if it mutates lower, then movement costs decrease. Specifically there is a 1/1000 chance for a positive mutation each turn, and a 1/1000 chance for a negative mutation. Consuming one fish adds one unit to a predator’s accumulated resources, and for the following experiments movement costs are set as 1/100th of a resource unit per turn, or 0.01 units.

4.1. Results – Second Experiment

A series of simulation runs was conducted with various values coupling success rate with movement costs. Figure 4 shows the results from five of these. The top image reflects the outcome if, for every mutation of a 1% increase in efficiency, the individual will have a 0.0001 increase in the movement cost (from a base rate of 0.01 – therefore, also 1/100th of a change). In the next image this trade-off was 0.0002 added to movement cost, for every 1% increase in efficiency. The third had a 0.00025 additional cost; the fourth, 0.0003; and the bottom is 0.0004.

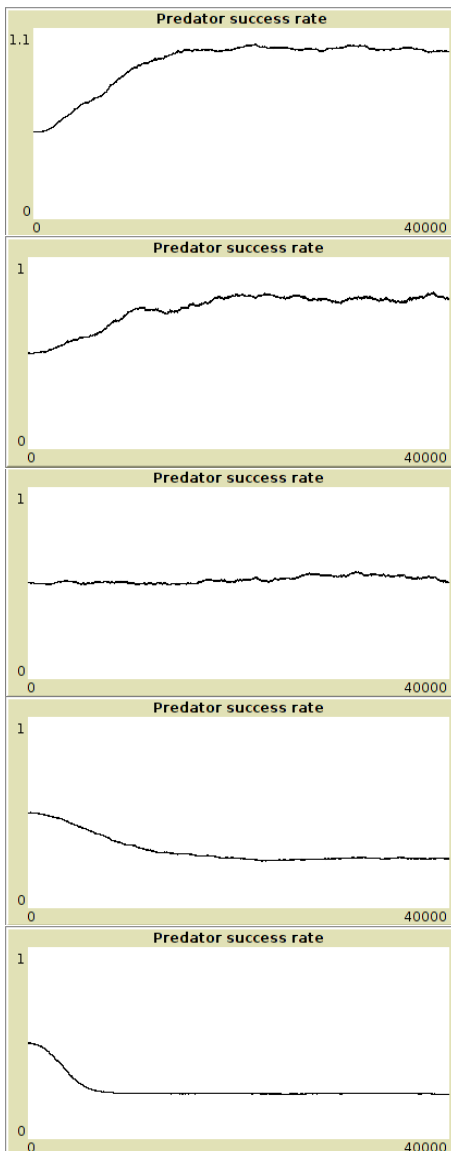


Figure 4: Evolved predator success rate, for five different levels of trade-offs between changes in effectiveness and changes in movement costs. The first and last (top and bottom) results fall outside of the model’s parameters; however, the middle three charts show the results of low, medium, and high cost trade-offs.

The top and bottom chart both produce results that are meaningless, under the definitions of this model. When the trade-off is cheap for additional hunting effectiveness (top chart, 0.0001), then the predators keep evolving until they reach 100% effectiveness, catching a prey each turn. Of course the model doesn’t allow more than 100% effectiveness, so this result is ignored. Similarly, when the trade-off is expensive (bottom chart, 0.0004), then the predators will continually evolve to be less and less effective hunters, in order to save energy. This evolution continues until the predators reach an average movement cost of approximately zero. This is an invalid result, since the model assumes all individuals will have a positive energy requirement for movement.

The three remaining charts, however, show that there is clearly an equilibrium rate for predator success, where the individual predators will, at some point, stop evolving their hunting effectiveness. In the second-from-top chart the equilibrium success rate is high, approximately 97% on average, with an average movement cost of ~0.015 (or about 50% more than at start). The middle chart the success rate is stable at 50%, with an average movement cost remaining very close to the 0.010 baseline (start) amount. And the second-from-bottom chart stabilizes at a success rate of ~25%, with movement cost reduced to ~0.003 (about 70% less than the start amount).

The most remarkable aspect of these experiments, adding a movement cost to the predators and allowing hunting effectiveness to evolve, is that the outcome for both populations is decidedly different than in the previous experiment. Recall that, if predator effectiveness was reduced from 100% to ~67%, then this would benefit the prey, increasing their equilibrium population size as well as their average age. However, now that energy is explicitly accounted for with a movement cost, and individual effectiveness can only be reduced by reducing this movement cost, the benefits of a similar change no longer accrue to the prey population. Figure 5 illustrates this result.

The bottom most row in Figure 5 reproduces the second-from-bottom chart in Figure 4, where the predators’ individual effectiveness is reduced from 50% to approximately 25%. Rather than benefiting, the prey population is reduced in size by a very small amount. Conversely,

the predator population increases quite dramatically: their final population size is double that at the start of the simulation, and their average age has increased as well. Clearly they have benefited from the reduction in movement costs, living longer while eating less, and not significantly reducing the overall size of their prey population. When the trade-off is high, predators will choose to reduce their efforts, both individually and as a

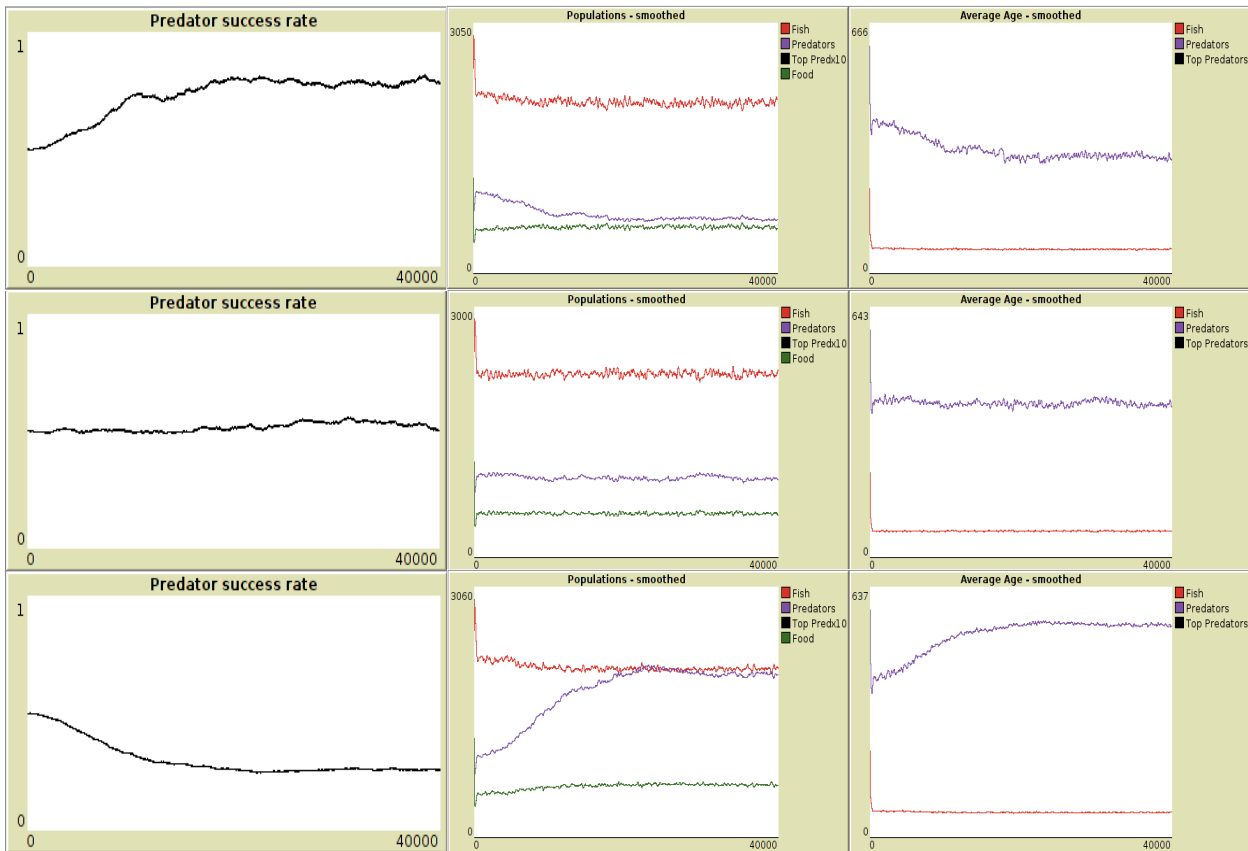


Figure 5: Low cost trade-off (top row), medium cost trade-off (middle row), and high cost trade-off (bottom row), between changes in predation effectiveness and changes in movement costs. The middle column displays the resulting changes in population levels for the predators (purple), prey (red), and food (green). The right-most column displays the resulting changes in average age for each population collective.

A moderate trade-off (0.0025, middle row) and a low trade-off (0.002, top row) both produce less of a change, especially the moderate trade-off. As we can see, the moderate trade-off results in a hunting effectiveness equilibrium that is almost unchanged, still at ~50%, and a movement cost that is nearly the same as well: 0.01 units per turn. As such, the population size and average age for both predators and prey is virtually the same at the start as at the end of the simulation run. Even though the individual predators could evolve to hunt better (or worse), at this trade-off they do not.

The low trade-off condition (0.002, top row) does have negative consequences for both populations. Recall from Figure 4 that the predators increase their hunting effectiveness, from 50% to very nearly perfect, ~97% on average. In the original model that doesn't account for energy costs, this would be harmful for the prey population, but not affect the predators at all, in either population size or average age. Here, things are different. The prey population experiences an almost imperceptible decrease in size, while the predators experience a ~40% decline in numbers, and a ~30% decline in average age. The lesson seems to be: when predator success comes too easy, everyone loses.

5.0. Discussion and Future Work.

In our previous results we showed how increasing the food to the prey population does not truly help the prey; rather, only the predator population increases in size. Because of this increase in predators, the prey are consumed faster than otherwise: they eat faster and reproduce faster, but since there is no change in the predators' hunting ability, the prey must also have a shorter lifespan when there are more predators. The Red Queen Hypothesis raises questions, however, about the predators' effectiveness: what happens when this is changed? And in

particular, can we discover the minimal conditions that put a limit on such change?

In the first set of experiments we adjusted this effectiveness directly, by simply dictating how often (stochastically) the predators would miss catching a prey. This potentially helps us understand the net effect on population levels that the Red Queen Hypothesis can have, but doesn't help explain the mechanism that will get us there. Absent a detailed study it may be assumed that both the predators and the prey should adapt endlessly, engendering an arms race in their antagonistic abilities. The second set of experiments, however, explicitly allows for a trade-off between the amount of energy expended by individuals vs. their success rate, so that they can improve, but only at a cost. Conversely, the predators can also reduce their effectiveness in order to conserve resources. As this simulation shows, when the trade-off is high this is exactly what they will do.

Throughout this paper we have assumed that improving or reducing predation effectiveness has been an evolutionary mechanism; i.e., through the mutation and spread of changes in the genotype. This is taken from the original formulation of the Red Queen Hypothesis. But these models are not limited to that interpretation. The changes are abstracted into a very simple measure of effectiveness, and this could also be interpreted to mean that the predators simply learn to become more (or less) effective hunters. Thus the applicability of this model is much broader than just changes on the evolutionary scale.

There is much more to be done in future work. The most obvious is to put similar energy constraints on the prey population as well as the predators. If they can also adapt their ability to escape predation, given similar resource trade-offs, how will they react, both individually and collectively? Another consideration is to instantiate a more complicated food-web. Due to the constraints of the Competitive Exclusion Principle there is not yet a simple way to add (and preserve) multiple species at a particular trophic level. However, we can easily add and experiment with a "top predator," in order to extend the food chain to a four-trophic-level system. As seen in previous results this addition will affect all population levels, in both size and average age. While we expect it is the size of the changes and not the direction that would be affected, it remains to be seen if the results presented here are robust to such an addition.

Conclusions

Even in a simple model of population dynamics with very basic assumptions we find many outcomes that are non-intuitive in nature. The power of this ABM, however, is that by understanding the fundamental properties in the simplest model first we can perhaps better understand how additional and more complicated factors affect dynamical food webs in the real world. Further, an ABM allows us to monitor many aspects of these simulated populations that are difficult or impossible to monitor for their real world counterparts. For example, average consumption by a predator population would be time- and resource-intensive to record in the field, while in a simulated environment it is elementary. Even more exciting however is the possibility that a properly calibrated ABM – one that is grounded to a specific, real-world food web – might also provide other simulated measures that can be collected in the field, and that make it possible to infer those that cannot be. For example, the average age of each population is a variable that has particular (and sometimes non-intuitive) characteristics that at least indicate what (for example) average consumption rates might be. Average age is much easier to record in a real-world environment, and so a realistic and grounded ABM has the potential to provide information that is difficult or even impossible to collect any other way.

References

- Allesina, S., & Pascual, M. (2007). Network structure, predator–prey modules, and stability in large food webs. *Theoretical Ecology*, *1*(1), 55-64.
- Brett, M. T., & Goldman, C. R. (1997). Consumer Versus Resource Control in Freshwater Pelagic Food Webs. *Science*, *275*(5298), 384-386.
- Brose, U., & Dunne, J. A. (2010). Modelling the dynamics of complex food webs. In H. A. Verhoef & P. J. Morin (Eds.), *Community Ecology Processes, Models, and Applications* (Vol. 1, p. 37–45). Oxford University Press.
- Carmichael, T., & Hadzikadic, M. (2013). Emergent Features In A General Food Web Simulation: Lotka–Volterra, Gause's Law, And The Paradox Of Enrichment. *Advances in Complex Systems*, *16*(08), 1350014.
- Casini, M., Lövgren, J., Hjelm, J., Cardinale, M., Molinero, J.-C., & Kornilovs, G. (2008). Multi-level trophic cascades in a heavily exploited open marine ecosystem. *Proceedings. Biological sciences / The Royal Society*, *275*(1644), 1793-801.
- Christensen, V., Guenette, S., Heymans, J. J., Walters, C. J., Watson, R., Zeller, D., et al. (2003). Hundred-year decline of North Atlantic predatory fishes. *Fish and Fisheries*, *4*(1), 1-24.
- DeAngelis, D. L., & Mooij, W. M. (2005). Individual-Based Modeling of Ecological and Evolutionary Processes. *Annual Review of Ecology, Evolution, and Systematics*, *36*(1), 147-168.

- Dunne, J. A. (2006). The network structure of food webs. In M. Pascual & J. A. Dunne (Eds.), *Ecological Networks: Linking Structure to Dynamics in Food Webs* (p. 27–86). Oxford University Press.
- Eichelberger, C. & Hadzikadic, M. (2006). Complex Adaptive Systems: Using a Free-Market Simulation to Estimate Attribute Relevance. In F. Esposito, Z. W. Raš, D. Malerba, & G. Semeraro (Eds.), *Foundations of Intelligent Systems* (Vol. 4203, pp. 671-680). Berlin, Heidelberg: Springer Berlin Heidelberg. doi: 10.1007/11875604_74.
- Epstein, J. M. (1999). Agent-based Models and Generative Social Science. *Complexity*, 4(5), 41-60.
- Findley, M. G. (2008). Agents and conflict: Adaptation and the dynamics of war. *Complexity*, 14(1), 22-35.
- Frank, K. T., Petrie, B., Choi, J. S., & Leggett, W. C. (2005). Trophic cascades in a formerly cod-dominated ecosystem. *Science (New York, N.Y.)*, 308(5728), 1621-3.
- Gell-Mann, M. (1994). Complex Adaptive Systems. In G. A. Cowan, D. Pines, & D. E. Meltzer (Eds.), *Complexity: Metaphors, Models, and Reality* (pp. 17-45). Addison-Wesley.
- Grimm, V. (1999). Ten years of individual-based modelling in ecology: what have we learned and what could we learn in the future? *Ecological Modelling*, 115(2-3), 129-148.
- Grimm, V., Revilla, E., Berger, U., Jeltsch, F., Mooij, W. M., Railsback, S. F., et al. (2005). Pattern-oriented modeling of agent-based complex systems: lessons from ecology. *Science*, 310(5750), 987-991.
- Holland, J. H. (1992). Complex Adaptive Systems. *Daedalus*, 121(1), 17–30.
- Khouja, M., Hadžikadić, M., Rajagopalan, H. K., & Tsay, L.-S. (2008). Application of complex adaptive systems to pricing of reproducible information goods. *Decision Support Systems*, 44(3), 725-739.
- Levin, S. A. (2003). Complex adaptive systems: Exploring the known, the unknown and the unknowable. *Bulletin-American Mathematical Society*, 40(1), 3–20. American Mathematical Society.
- Miller, J. H., & Page, S. E. (2007). *Complex Adaptive Systems: An Introduction to Computational Models of Social Life (Princeton Studies in Complexity)* (p. 284). Princeton University Press.
- Mitchell, M. (2009). *Complexity: A Guided Tour* (p. 368). Oxford University Press, USA.
- Oksanen, L., Fretwell, S. D., Arruda, J., & Niemelä, P. (1981). Exploitation ecosystems in gradients of primary productivity. *American Naturalist*, 118(2), 240–261.
- Power, M. E. (1992). Top-down and bottom-up forces in food webs: do plants have primacy. *Ecology*, 73(3), 733–746. Eco Soc America.
- Reid, P. C., Battle, E. J. V., Batten, S. D., & Brander, K. M. (2000). Impacts of fisheries on plankton community structure. *ICES Journal of Marine Science*, 57(3), 495-502.
- Rosenzweig, M. L. (1971). Paradox of Enrichment: Destabilization of Exploitation Ecosystems in Ecological Time. *Science*, 171(3969), 385-387.
- Roy, S., & Chattopadhyay, J. (2007). The stability of ecosystems: a brief overview of the paradox of enrichment. *Journal of biosciences*, 32(2), 421-8.
- Scheffer, M., & De Boer, R. J. (1995). Implications of spatial heterogeneity for the paradox of enrichment. *Ecology*, 76(7), 2270–2277.
- Valdovinos, F. S., Ramos-Jiliberto, R., Garay-Narváez, L., Urbani, P., & Dunne, J. A. (2010). Consequences of adaptive behaviour for the structure and dynamics of food webs. *Ecology Letters*, 13(12), 1546–1559.
- Van Valen, L., (1973). A new evolutionary law. *Evolutionary Theory*, 1: 1-30.
- Wilensky, U. (1999). NetLogo. <http://ccl.northwestern.edu/netlogo/>. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.