Bottled Understanding: the role of lab-work in ecology

Adrian Currie

Penultimate version, forthcoming in the British Journal for Philosophy of Science.

Acknowledgements:

This paper has been presented at the 2016 PSA, the philosophy department at Calgary University and the Cambridge History & Philosophy of Science Department. I'm grateful to the comments I received, and to Angela Potochnik, Soazig Le Bihan and two anonymous referees for this journal for comments on drafts. I'd also like to thank Amy Payne for introducing me to Yoshida et al's work.

Abstract

It is often thought that the vindication of experimental work lies in its capacity to be revelatory of natural systems. I challenge this idea by examining laboratory experiments in ecology. A central task of community ecology involves combining mathematical models and observational data to identify trophic interactions in natural systems. But many ecologists are also lab scientists: constructing microcosm or 'bottle' experiments, physically realizing the idealized circumstances described in mathematical models. What vindicates such ecological experiments? I argue that 'extrapolationism', the view that ecological lab work is valuable because it generates truths about natural systems, does not exhaust the epistemic value of such practices. Instead, bottle experiments also generate 'understanding' of both ecological dynamics and empirical tools. Some lab-work, then, aids theoretical understanding, as well as targeting hypotheses about nature.

1. Introduction

Scientific investigations are often expected to yield truths about natural systems, albeit through torturous, ingenious routes. One such route involves the investigation of surrogate systems departing dramatically from nature: mathematical models, analogous evidence, lab-based ‘bottle’ or ‘microcosm’ experiments, and simulations. The expectation that science primarily delivers truths about natural systems underlies what I’ll call *extrapolationism* about surrogates. On this view, the vindication of scientific practices turns on their generating such truths. I’ll argue that extrapolation is too narrow. Focusing on ecological surrogates – bottle experiments in particular—we’ll see that the generation of understanding is a complementary epistemic good alongside obtaining propositional truths about the natural world.

Bottle experimentation involves lab-raised, easily managed critters in highly artificial environments[1] (Jessup et al [2004]). I’ll argue that although in some circumstances bottle experiments have only limited capacities to tell us about the natural world, bottle experiments have features which aid in our grasping the nature of ecological dynamics and the assumptions underlying the empirical tools ecologists employ. Specifically, they enable an ability: the capacity to navigate ecological possibility, a kind of modal understanding.

I’ll begin by sketching the investigations bottle experiments are often taken to inform: the relationship between the formal models which community and population ecologists take to explain ecological systems, and those systems themselves. I’ll then move to an illustrative case study. With this in place, I articulate extrapolation and argue that the heterogeneity of ecological systems potentially imperils extrapolationist defences. I’ll then turn to a positive defence of bottle experiments and their capacity to produce understanding. Note that my argument that bottle experiments provide understanding is independent of my reservations about extrapolationism; that negative argument matters for articulating extrapolationism and its

---

constraints, thus motivating the latter positive argument, but the positive argument doesn't depend upon it.

Three caveats. First, by isolating the practice of bottle experimentation, I do not imply (and do not think) that it is in epistemic competition or conflict with other ecological practices (theoretical modelling, field experiments, etc…). On my view, ecology is strong when it is opportunistic and pluralistic. Second, I don't argue that *all* experimentation serves to generate modal understanding. The epistemic power and value of a scientific practice turns crucially on the context at hand. As such, extending my argument to other practices requires explicitly examining them. Having said this, I suspect that experimental evolution—the use of bottle experiments to probe macro-evolutionary patterns—would turn out similarly. Third, I don't take myself to be making substantive claims about the nature of 'experiments' (be they natural or otherwise) 'models' or 'surrogates' generally. Rather, I am examining a particular ecological practice with an eye to identifying the epistemic dividends they bring. My appeals to notions such as 'experimentation' are pragmatic and I think could happily accommodate different substantive accounts.

2. Trophic Interactions & Observational Techniques

Here's a partial sketch of community ecology which emphasizes the role of theoretical models in identifying causal interactions in natural systems. I don't pretend this captures all there is to community ecology, but nonetheless it is a helpful starting point both for section three's case study and my account of extrapolationism.

Community ecologists are in the business of accounting for relatively short-term patterns of population abundance by positing interactions between populations. On this approach, a group of critters' fate is intertwined with the fate of who eats them, who they eat, and the strength of those interactions. Positing trophic interactions—predator and prey, parasite and host—

therefore plays a crucial explanatory role. These interactions are typically represented using relatively simple coupled differential equations, most infamously the Lokta-Volterra predator-prey model. Such models don't merely generate curves resembling the population-dynamics of actual groups of animals—they also posit causal information. Ecological models oscillate as they do in virtue of the links between the equations. In a predator-prey model, the prey number at one moment is in part determined by predator number at an earlier moment. For such models to carry explanatory oomph, the actual cause of the interactions in the wild should mirror the causes posited in the model.

Ecologists, then, must measure the strength of trophic interactions in natural systems. If trophic interactions are weak, if the dependencies between populations of predators and prey are *de*coupled, then the applicability of the relevant mathematical models are undermined. Strong trophic interactions in natural systems vindicate such models[2].

A common method for establishing trophic interactions involves measuring changes in species abundance in natural systems over time (Wooton [2005]). Observations of abundance-patterns generates data which is fed into a model designed to measure interaction strength. Let's call such methods *observational techniques*. The observed changes in populations are taken to be revelatory of trophic interactions. If the population of prey regularly dips after an increase in predators, this is taken to indicate a trophic interaction of the sort mathematical systems model.

Thus, ecologists often combine investigations of two kinds of systems: mathematical and natural. These together provide a means of establishing the trophic interactions which are used

---

[2] Sterelny ([2006]) & Odenbaugh ([2007]) make similar points, tying such vindication to the *existence or otherwise* of ecological communities. I'm interested in the tools ecologists use to get empirical traction on their targets, so leave such ontological questions for another day.

to explain patterns in population abundance. But ecologists also examine a third kind of system—*experimental systems*. How can lab work bear on observational techniques?

3. Cryptic Dynamics in Bottle Experiments

It is *prima-facie* plausible that ecological patterns underdetermine causes; several different mechanisms could generate the same population dynamics. This initial thought makes trouble for observational techniques which infer from observed patterns in natural systems to trophic interactions between populations. Could other interactions generate these patterns?

Imagine observing the following in an ecosystem: a predator population cycles, but prey remains constant. In this circumstance, observational techniques infer a weak trophic interaction between the two populations. As there are no changes in prey population, the cycling of predators must be independent; moreover, the prey population appears insensitive to predator number. However, imagine that the prey population contains two morphs, *tough* and *soft*. Toughs are resistant to predators, and so do well in predator-rich environments. However, they are outcompeted by softs when predation pressure is low. In those circumstances, it is possible that total prey number remains roughly even, while the population oscillates between high tough/low soft (when predator population is high), and low tough/high soft (when it is low). There are strong trophic links between predator and prey, but observational techniques predict weak interactions.

This is an example of a *cryptic population dynamic*: the strength of the relevant trophic interactions are masked by rapid evolution in the prey population[3]. Yoshida et al ([2007]) probe the phenomenon in experimental systems:

---

[3] This is only one example of cryptic dynamics. Another is 'apparent competition'. Here, unknown density-dependent predation makes it appear that interspecific competition is occurring.

> … we have discovered that ecologically relevant conditions exist under which there is little or no relationship between the abundance of a predator and the population growth rate of its prey, and vice versa, despite the existence of what is known to be a tightly coupled interaction. Under those conditions, a strong interaction would be completely missed in an analysis based on observed changes in species abundance (1868).

They describe three sets of experiments implicating cryptic dynamics, and explore them using mathematical modeling. They suggest their results undermine observational techniques. Although Yoshida et al's case is not representative of all work on bottle experiments, it is a useful example for our purposes here. Yoshida et al provide an extrapolationist defense, and they explicitly link their work to how community ecologists explain natural systems. As I'll discuss in section 4.2, other bottle experiments are less plausibly geared towards the natural systems (see the cases described in Griesemer [1988], Griesemer & Wade [1988]). Indeed, it may be that other examples of the strategy could be used to undermine the position I develop in sections 4 and 5, but such arguments must be made[4].

The first of Yoshida et al's experimental systems involves rotifers preying on algae. Experiments establish that the trophic interactions between these are strong: the out-of-phase oscillations which are distinctive of predator-prey cycles can be produced, and are disrupted as models predict via intervention. This is true even of populations containing several types of genotypically distinct algae of the same species.

Consider the following surprising cases.

First, a small number of rotifers are added to a mixed algae population. We should expect an initial increase in rotifer number, and a plunge in algae, before the populations settle into steady cycling. Instead, the algae population plunges, and stays low, while the rotifer population

---

[4] I do not generally think that case studies in the philosophy of science ought to be interpreted as bases for inductive inferences, see Currie ([2015]) and references therein.

oscillates. That is, the prey population remains steady while the predator population behaves as if in a regular predator-prey relationship (figure 1).

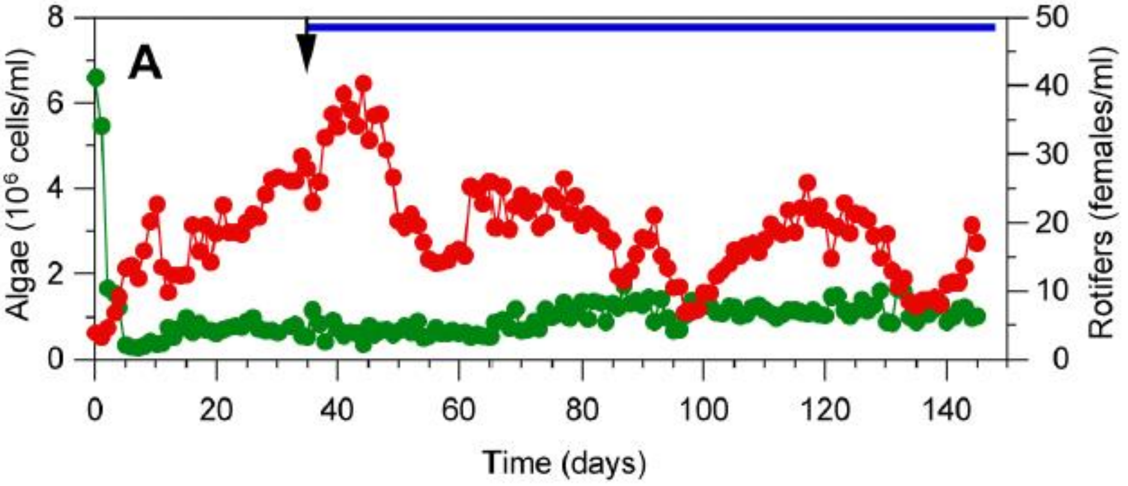

Figure 1: cryptic cycles in algae-rotifer system (detail from Figure 1, Yoshida et al [2007], open licence).

We know that rotifers and algae have strong trophic interactions, but here they don't act like it:

> Even though we know that the rotifers and algae are bound in a tight predator-prey relationship, a plot of their densities in a predator-prey phase plane suggests instead that rotifers and algal populations are completely decoupled. (1869)

Yoshida et al suggest that cryptic dynamics due to hidden population structure could explain their results. They suggest that there are two kinds of algae, mapping onto tough and soft. The algae population remains stable because of compensation between the two morphs. If these dynamics occurred in natural systems, observational techniques would mistakenly infer weak trophic links.

Yoshida et al explore their hypothesis using a simple mathematical model, effectively a predator/prey equation where prey are split into two groups. That is, total prey abundance at a time is the sum of the value of two morphs. These differ along two variables: their ‘palatability’, or the probability of being preyed upon, and their birth rate. We can represent toughs and softs as follows: toughs have low palatability and a correspondingly low birthrate, softs high palatability and high birthrate. The model indeed generates the kinds of cryptic cycles observed in the rotifer-algae system, driven by density compensation.

> [In the model] Prey evolution is driven by occasional predator outbreaks, but prey density remains nearly constant because the consumption of the vulnerable prey is almost exactly balanced by growth of the better-defended prey when they are released from competition with the vulnerable type (1871).

Now consider a second experiment: a mixed population of algae competed without rotifer disruption for a significant period (12 days). When rotifers were introduced, the two populations briefly switched to standard predator/prey oscillations, before settling to a low, constant population of algae with cycling rotifers (figure 2).

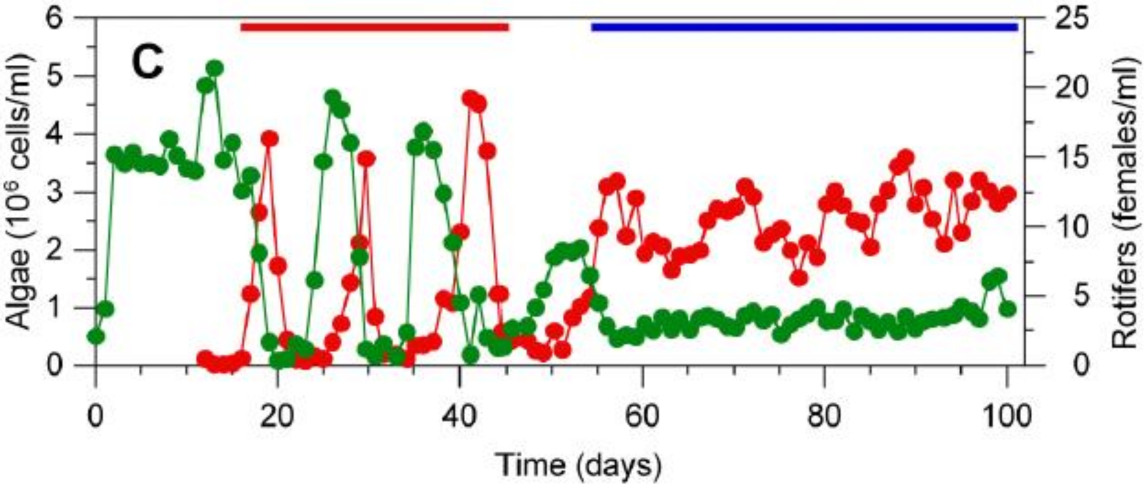

Figure 2: predator/prey cycle followed by cryptic cycle. (detail from figure 2, Yoshida et al 2007, open license).

Yoshida et al again explain these results in reference to rapid evolution between morphs. In circumstances without predation, *softs* would predominate. Upon the introduction of rotifers, the relatively homogeneous prey population creates a standard predator-prey cycle. Once the proportion of toughs increases sufficiently, the cryptic rapid-evolution cycle commences. Again, this is borne out in the mathematical system.

So, Yoshida et al posit an explanation for the behavior of the rotifer-algae system: that it is due to rapid evolution oscillating between tough and soft algae. Further, "… mathematical modeling demonstrates this is a plausible explanation" (1872). However, can we be sure that rapid evolution in fact occurs? Here, Yoshida et al turn to a different experimental system.

This system consists of bacteria and phage (viruses which target bacteria). In this setup, two different bacterial clones are used, one susceptible to phage and another resistant. The latter group carries a marker which enables tracking the relative density of genotypes in the bacteria population. In an experiment which mimics the original rotifer-algae study,

> The fraction of the susceptible genotype in the total bacteria population clearly showed evolutionary cycles in concert with cycles in bacteriophage density, as our model predicts (1872).

To summarize, we observe a phenomenon in the algae-rotifer system: apparently decoupled population dynamics in two populations we know to have strong trophic interactions. We then consider whether rapid evolution is to blame, constructing a simple model to represent and explore this hypothesis. The model predicts rapid evolution under highly simple circumstances. Finally, we turn to another experimental surrogate—bacteria and phage—and confirm the same dynamics, but this time directly observing cycles in the cryptic prey population.

This, it seems to me, provides persuasive evidence that rapid evolution occurred in the algae-rotifer system. First, we know there is tight trophic coupling in the system; second, we know the prey population is mixed; third, plausibility is granted by two surrogates (the mathematical model and the bacteria-phage experiments)[5]. On this picture, Yoshida et al infer from one bottle to

---

[5] On my view, the model's success in aping both the first and (more impressively) the second rotifer-algae experiments provide some epistemic grounds for the cryptic-dynamic explanation's plausibility. For

another; triangulating between two experimental systems and a mathematical one. However, they don't frame their work as explaining an odd result in a bottle experiment. Instead, they tie their study to the observational techniques used to establish trophic interactions in natural systems:

> The potential for the kind of cryptic dynamics that we have documented here means, however, that there are circumstances in which the absence of such statistical relationships cannot be reliably taken as evidence for the absence of important biological interactions (1874).

Further:

> … efforts to establish the nature and strength of interactions in ecological communities that fail to consider the potential for evolution (which is to say virtually all efforts to date) run a risk of being incorrect. Because essentially all natural populations have heritable variation for ecologically important traits… ignoring the potential for evolution to affect measurements of species interactions becomes untenable (1874).

According to Yoshida et al, then, their experimental results matter because they lead us to question the application of observational techniques to natural systems—they are extrapolationists. In the next section, I first expand upon and then question this kind of view.

4. Extrapolationism

An extrapolationist defense of a scientific investigation turns on its role in generating propositional knowledge—truths—pertaining to natural systems. That is, for the extrapolationist the value of an investigation is primarily due to its confirmatory prowess: it provides grounds for belief in some hypothesis pertaining to natural systems.

---

discussion of how under some conditions models can play confirmatory roles see Parke [2014], Currie [2018] chapter 9.

In this section I'll characterize extrapolationism in the context of ecological bottle experiments, and highlight a vulnerability in extrapolationist defenses of that practice. I'll begin by identifying the kind of support bottle experiments might provide *vis-à-vis* hypotheses concerning natural systems. I'll then argue that the heterogeneity of ecological systems potentially undermines the significance of that support. I take this position to motivate further consideration of the epistemic value of bottle experiments: something I turn to in section 5.

Surrogates, according to extrapolationism, target natural systems, and the resemblance between them facilitates extrapolating results from the former to the latter[6]. Our ecologist notices that, under certain conditions, the two populations in her model oscillate, the prey lagging by a quarter-cycle. These dynamics result from the links between the equations representing prey and predators. The resemblance between the natural pattern and those generated by the model provide the basis for extrapolation: natural ecological systems exhibit the patterns they do because of the kinds of interactions captured by the model.

Extrapolationism is, I think, a common view concerning bottle experiments. Jay Odenbaugh ([2006]), for instance, suggests something like this:

> … if ecologists can show that the differences between the artificial and natural systems are irrelevant and that the model of interest can accurately represent the dynamics of the former, then there is no reason to believe that it cannot do so with respect to the latter (728-729).

Bottle experiments are indeed informative of natural systems. However, I argue that extrapolationist defences of bottle experiments might not be sufficient alone. I should be explicit about the nature of my argument. We're here concerned with the *pursuitworthiness* of some

---

[6] See, for instance, views on modelling which emphasizes their role of representing target systems (Weisberg [2012]).

investigative strategy[7]. That is, given that ecologists could adopt a range of strategies, what are the epistemic dividends we should expect from bottle experimentation, and are they sufficient to justify the practice? To answer such questions, we should get a grip on the expected epistemic goods generated by an investigation. This can then be balanced against its cost and difficulty, and thereby with competing strategies. Arguing that, on extrapolationist grounds, the epistemic returns from bottle experiments are potentially quite weak doesn't involve arguing that they do not provide such goods. Far from it. Rather, the discussion ought to motivate a richer picture of the epistemic goods bottle experiments provide, a task I'll turn to in section 5.

4.1 Ecological Possibility & Actuality

On an extrapolationist reading, Yoshida et al's investigation matters due to informing us as to whether cryptic dynamics can occur in natural systems. To see how, distinguish between ecological *possibility* and *actuality*. The space of ecological possibility concerns the dynamics and interactions described by ecological theory. Such theory is diverse, but for our purposes we could consider it as the modal space described by predator-prey models and reasonable extensions such as Yoshia et al's[8]. Ecological actuality concerns that subset of ecological possibility space which we have reasonable grounds to expect to be instantiated. Yoshida et al's model demonstrates the ecological possibility of cryptic dynamics: nothing in ecological theory forbids their occurrence; under certain conditions they appear to be live ecological possibilities. The bottle experiment establishes ecological actuality insofar as under certain conditions a few physical systems can instantiate cryptic dynamics. But the extrapolationist is interested in whether cryptic dynamics are realized in *naturally occurring* ecological systems. For the

---

[7] Traditionally the context of pursuit concerns hypotheses (Franklin [1993], Laudan [1977]), but recently Rune Nyrup [forthcoming] has retooled the discussion towards investigations themselves.

[8] One might reasonably complain that it is difficult to ascertain which features and phenomena are endogenous or exogenous to ecology. I don't think there is a single correct answer to this question (Currie & Walsh [2018]) but also don't think this matters for our purposes here: we needn't specify an exact scope of ecological possibility in order to sensibly discuss it, or its relationship with ecological actuality. See Lange ([2005]) for more on ecological possibility.

extrapolationist, demonstrating actuality in an artificial system provides grounds for expecting actuality in natural systems. Let me explain further by returning to observational techniques and mathematical models.

As Yoshida et al point out, observational techniques assume "… that the impact of one species on another is revealed by patterns of covariation in their changes of abundance through time" (1868). Observational techniques shift from correlations between oscillations in natural populations to causal interactions between them. This differs epistemically from circumstances where our capacity to produce controlled and fine-tuned disruptions enables direct confirmation of trophic coupling—such as in the rotifer and algae system.

As I mentioned earlier, in principle patterns underdetermine causal mechanisms. Cryptic population dynamics are an example of such underdetermination in ecology. In the circumstances Yoshida et al are concerned with, a pattern involving oscillating predators and stable prey could be generated by (at least) either (1) decoupled populations or (2) cryptic dynamics between two prey morphs[9]. Articulating such alternative hypotheses matters for empirically distinguishing between them in natural systems. Indeed, Yoshida et al claim that their mathematical model provides the necessary conditions required for rapid evolution:

> Using a mathematical model, we have established the conditions necessary for these cryptic cycles: (1) predator-prey cycles would occur between the consumer and the more vulnerable genotype of the resource species if that genotype were the only one present, (2) the less vulnerable of the resource genotypes has an effective defense against the consumer, and (3) the cost of defense is fairly low in terms of ability to compete for limiting resources (1876).

---

[9] Of course, there may be *other* dynamics which could explain the observed pattern (perhaps involving unnoticed patterns of immigration, for instance). However, for Yoshida et al it is crucial to undermine the assumption that, in natural systems, we should expect observed patterns to track standard predator-prey models. As such, focusing on one plausible alternative explanation is sufficient.

This is undoubtedly an important theoretical result. The crucial question here concerns what the *experiments* bring to the table. Yoshida et al's mathematical system articulates the possible circumstances under which cryptic dynamics may occur, and thus provides hints as to what kinds of empirical considerations might guard against them—for instance, we should care about the homogeneity or otherwise of prey populations. We have, then, a demonstration of ecological possibility. However, they emphasize the importance of their experimental systems, and their connection to natural systems:

> Although our microcosms are extremely simple systems, they mimic the consumer-resource interactions in natural systems. Our rotifer-algal interaction is a herbivore consuming a primary producer… and our phage-bacterial interaction can be considered as either predator-prey or host-parasitoid (1876).

And so to our central question: if we expect different dynamics to generate the same patterns, and if our mathematical system expresses and explores these, what does actualizing them in a simplified experiment achieve?

The extrapolationist response is that demonstrating interactive dynamics like rapid evolution in experimental systems makes it plausible that such interactions are realized in natural systems. Bottle experiments are an empirical mid-way point between mathematical and natural systems. That is, demonstrating the actuality of the dynamics in a bottle experiment underwrites reasons for believing that they are actual in natural systems. The inference appears to be from the occurrence of dynamics in one area of ecological possibility space (artificial conditions) to their likely occurrence in a quite different part of that space (natural conditions). What might license such an inference? Although Yoshida et al emphasize the similarities between natural and artificial systems, their bottle experiment is specifically set up to mimic the theoretical dynamics that interest them. We might complain that it is unsurprising that the experiment behaved as it did: the circumstances were so constrained that the system had little choice (Crane & Molofsky [1999])

In typical experimental circumstances, we allow experimental subjects sufficient freedom to buck our expectations (Morgan [2005]). Bottle-experiments, however, are so highly controlled that their behaving other than models predict is unlikely. This undermines the effectiveness of extrapolationism, as the new empirical information we gain vis-à-vis natural systems appears weak. It is possible, of course, that the experiments turn out to be *unable* to behave as models predict. In such circumstances, this might give us reason to think that the relevant mathematics cannot be realized. However, what is significant about realizing dynamics in an ideal physical system, for claiming that those dynamics also occur (or possibly occur) in complex, natural systems?

The answer appears to be that experimental results prove ecological actuality in a broad sense, and this is an empirical stepping-stone for showing they are realizable in natural systems. If a scientist is unable to actualize the dynamics in question in a simple system then, the thinking goes, such dynamics are unlikely to occur in nature. If the scientist can, then, the hypothesis that the dynamics are possible passes an empirical test and is, *prima-facie*, more likely. Applied to the case at hand, Yoshida et al's experiments test whether it is possible to physically realize a system wherein trophic interactions are masked by fast-paced, cryptic dynamics. Because the two experimental systems (particularly the bacterial population, where morphs could be tracked) exhibit those dynamics, we have good reason to believe that cryptic evolution is ecologically actual—after all, we have an existence proof in the two experimental systems.

4.2 Ecological Heterogeneity

The extrapolationist defense of bottle experiments has vulnerabilities. Bottle experimentation is often hard: constructing and maintaining microcosm experiments, and getting them to behave as they do, is a significant achievement. But it is unclear whether demonstrating ecological actuality in an artificial system is epistemically significant on extrapolationist grounds. As we've seen, there is no direct extrapolation to be had from the artificial actuality to natural actuality. Rather, achieving artificial actuality gives some reason to think the dynamics are possible in natural systems. But the significance of even this can be questioned. The argument turns on the idea that a result's significance is sensitive to our priors concerning it. If we already expect a result, then the confirmatory dividends it brings are low. I want to suggest we have reason to think that a wide range of ecological dynamics are actual, and this to at least some extent undermines the epistemic value of demonstrating that particular dynamics are actual. Such an argument is not intended to show that extrapolationism is false—that there are *no* such dividends. Rather, I will demonstrate that its effectiveness as a defense of the pursuitworthiness of bottle experiments is open to question, and this should thus motivate us to ask whether there is more that might be said regarding their epistemic value. The argument turns on the heterogeneity of ecological systems.

Ecological systems range over many types of critters, trophic levels, scales, and interaction strengths. This grants some reason to expect ecological actuality to cover a fair portion of ecological possibility. In ecology, *constraints* on possibility are surprising—what is actual is often less so. John Matthewson ([2011]) has helpfully distinguished between 'complexity' and 'heterogeneity', and Alkistis Elliot-Graves ([2016]) adapts this conceptual machinery to an ecological context where the behavior of natural systems matters crucially: invasion biology. In what follows, I'll draw on this work to explain why under some conditions our priors regarding the instantiation of many dynamics in natural ecological systems should be high.

Briefly, Matthewson's distinction turns on the number of parts and interactions in a system, and the types of parts within that system. The former is *complexity*: a system may be composed of a large number of differently interacting objects which are nonetheless homogenous. The latter is *heterogeneity*: a system may be composed of few parts, but they may be of different types. Matthewson's suggestion is that heterogeneity, not complexity, plays an important role in explaining the differing successes of disciplines such as physics and chemistry on the one hand, and ecology on the other. Both target complex systems—there are many components and they interact in various ways. But ecological systems are also *heterogeneous*. In physics, although there may be many electrons, and their interaction may be somewhat bewildering, by and large one electron can be a model for all[10]. Whereas in ecology, different predators, say, predate in different ways—and these differences matter. Elliott-Graves focuses on heterogeneity *between systems*: different tokens of system types may contain quite different components. Such heterogeneity undermines our capacity to take one system as a guide to others.

Across heterogeneous systems, we should expect a wide range of dynamics to be actual. A set of dynamics—those described by a predator-prey equation, for instance—are abstract. They will only apply across heterogeneous systems very coarsely. To describe how such systems will in fact behave, they must be supplemented by further detail[11]. Which is to say, a *different*, more specific set of dynamics must be utilized. And so, we can move rather easily (perhaps even analytically) from a set of systems being heterogeneous, to their instantiating a wide range of dynamics. At the extreme end of the scale, we should expect that for any theoretically possible set of dynamics, those dynamics are ecologically realizable. This is surely a step too far: there are likely many constraints on what possible ecosystem dynamics can be actualized. However, the more causally heterogeneous we think a set of systems are, the more diverse we should expect the dynamics instantiated across those systems to be.

It is plausible that the significance of a result should turn on how surprising that result is, given our background knowledge. That is, if I already expect a result (for good reasons!), demonstrating the result is less significant than if the result was not expected. Significance is not

---

[10] It's worth sounding a note of caution here, physical systems characterised as being 'simpler' than biological systems often turn out to be much more tricky than we imagined (see Havstad [forthcoming] on H2o, for instance).

[11] Elliott-Graves argues that this should be done via 'constrained integration', models of causal dynamics are filled in on a case-by-base basis.

determined by this feature alone: the domain of the result might matter very deeply to us (as it might in cases of high inductive risk), or it might be deeply embedded in a range of other theories. But regardless, significance and epistemic expectations are tightly knit.

If these points hold water, then we have the following argument. First, the heterogeneity of ecological systems gives good reason to expect many ecologically possible dynamics to be actual. Second, the significance of demonstrating actuality is sensitive to our expectations about actuality; if we already expect a set of systems to exhibit a range of dynamics, demonstrating that a particular dynamic is actual is not very significant. Therefore—third—demonstrating the artificial ecological actuality of a dynamic is not so significant. However, fourth, extrapolationism argues that the pursuitworthiness of bottle experiments lies in their demonstrating the ecological actuality of some dynamic in an artificial system. Insofar as bottle experimentation is often a difficult, costly practice, then, we can conclude that fifth, if extrapolationism is our only option, then bottle experimentation's pursuitworthiness is in peril.

Theoretical modeling demonstrates the theoretical possibility of an ecological dynamic. We want to know whether this theoretical possibility is likely to be realized in a natural ecological system. The extrapolationist points to an artificial case, and argues this provides empirical grounds for expecting the same in natural systems. However, there is another route between ecological possibility and actuality: heterogeneity. Because we should already expect a wide variety of ecological dynamics to be realizable, the epistemic gap between possibility and natural actuality is not so wide. It doesn't follow from this that we will always be right: it could be that some areas of possibility are in fact very difficult to instantiate, and some apparent complexities common. However, it does suggest we shouldn't be particularly surprised if a particular set of dynamics are instantiated. This does not deny that demonstrations of ecological actuality in bottle experiments are epistemic achievements, however, such achievements are insignificant given our well-grounded expectations about natural ecological systems[12].

This is by no means the end of the story vis-a-vis extrapolationism, as the success of the argument above turns on a major assumption. For heterogeneity to lead us to expect significant overlap between ecological possibility and actuality, we must have an idea of the scope of ecological possibility. If this is sufficiently large, then it is perfectly possible for sections to be unactualized even across heterogeneous systems—under those conditions the kind of actuality-proofs generated by bottle experiments may be significant. I suspect the scope of ecological possibility is itself ambiguous (see footnote 8), but even so there is something of an impasse here: the significance of the practice turns in part on our beliefs about ecological possibility. I think this is sufficient to motivate considering options beyond extrapolationism.

I have based my analysis on Yoshida et al's study, and this might not be representative of bottle experiments. Both Jim Greisemer ([1988]) and Jay Odenbaugh ([2006]) discuss cases which appear to diverge. Griesemer discusses bottle experiments which attempt to examine interspecies competition in a kind of abstract, 'pure' form. These were not in the business of representing natural systems, and they did not involve the kind of purposeful interventions characteristic of Yoshida et al. These cases are not presented in extrapolationist ways, and—I think—are nicely captured by the account I'll provide in the next section (indeed, I'll return to Griesemer's view later). Odenbaugh discusses Costantino et al's ([1997]) demonstration that chaotic population dynamics can be instantiated in simple ecological systems. Although this had

---

[12] There may be subtler forms of extrapolationism. Take a view inspired by Strevens ([2008]), for instance. On this view, the experimental system works by isolating some crucial difference makers occurring in the natural system. In doing so, the artificial system generates a true explanation of the natural system insofar as they two share difference-makers. I don't think this view works in this context. Given the heterogeneity of ecological systems, and the theoretical results at hand, the epistemic benefit of the experimental is still minimal. The theoretical model itself plausibly isolates difference-makers (albeit more abstractly) and seems to be a source of explanations just as the experiment is. As in standard extrapolationism, the cost of the experimental practice doesn't appear justified by the payoff.

been demonstrated theoretically, according to Odenbaugh their occurrence in artificial systems was a genuine surprise. If Odenbaugh is correct, then in this instance our priors in possible results being actual are low, and so the results are significant. This should not be surprising: the applicability of my argument turns on assumptions about ecological possibility and its relationship with ecological actuality. Reasonably, some cases might turn be like Costantino's surprising results, while others might be like Yoshida et al's (I claim) unsurprising results.

Regardless, this argument puts pressure on extrapolationism about bottle experiments. That view relies on the value of taking artificial actuality as providing epistemic grounds for expecting natural actuality. But if we think the ecological world is highly diverse and heterogenous, in some circumstances we might see the step from ecological possibility to actuality as a short one. Under those circumstances, extrapolationism is weak. Happily, the epistemic benefits of bottle experiments are not exhausted by extrapolationism.

5. Understanding

Bottle-experiments are part of a wide range of investigative strategies: field observations and experiments, natural microcosm experiments, simulations, and so forth (Odenbaugh [2005], Love & Travisano [2013]); working across different scales is integral to ecological practice. (I don't think this attention to scales is independent of the considerations I'll consider). I've provided reason to worry about the pursuitability of bottle experiments on extrapolationist grounds alone. One option, mentioned above (and in Odenbaugh [2006]) is to construe theorizing in ecology as *about* what happens in bottles. Let's pause and consider this option.

Nancy Cartwright ([1999]), emphasizing the highly artificial, controlled nature of experimental physics, argues that physical laws are about the dispositions of objects to behave under controlled and simplified laboratory conditions, rather than natural ones. Similarly, Ken Waters ([2007]) argues that lab work enables the study of actual difference makers: the merely potential causal components of theories are actualized and interrogated in controlled circumstances. This captures something right about Yoshida et al's study. The original purpose of their theoretical model was to express a hypothesis about the cryptic dynamics of the rotifer-algae system. This, in combination with further studies of that system, and the algae-phage system, provided reason to think the hypothesis was true of that rotifer-algae system. Moreover, such a position saves something like extrapolationism. Scientific theorizing isn't about generating truths of natural systems, but generating truths about experimental systems. However, this is unsatisfying: if scientific truths are just about artificial circumstances which we construct, their significance to ecology is a puzzle. I earlier mentioned the interpretation by which Yoshida et al extrapolate from cryptic dynamics in one artificial system (composed of bacteria and phage) to another (composed of rotifers and algae). I think this is a plausible reading (especially considering my suspicion of extrapolationism in this instance), but what does such knowledge get us?

The answer, I think, is to see that although truths are generated about experimental systems, their epistemic value doesn't lie in such truths *per se.* And although the results from bottle experiments do demonstrate ecological actuality, their epistemic value is not exhausted by this. Instead, I want to suggest that much of the epistemic value of experimentation like Yoshida et al's lie in their capacity to produce *understanding.* Moving forwards, then, I have three tasks. First, I'll sketch an account of understanding as an epistemic good. Second, I'll argue that bottle experiments are a genuinely good tool for generating understanding. Third, I'll discuss how understanding may be exported from the laboratory contexts in which it is generated.

5.1 The Epistemic Good of Understanding.

The thought underlying extrapolationism is that science is in the business of producing truth, or something like it—and truth understood along broadly correspondence terms. However, it is clear that science doesn't aim for that kind of truth alone. Indeed, it is plausible that there is a wide range of epistemic goods which science produces (Potochnik 2017, 2013, Chang 2012). One set of plausible options is the production of *understanding*.

Philosophers disagree about the nature of understanding and its relationship to other epistemic goods, such as explanation and truth[13]. At base, there is some relationship between knowing something is the case, and grasping that it is. A common theme, and one which suits me here, is to distinguish between understanding, explanation and truth by designating understanding as a kind of *know-how*. That is, where explanations and truths are primarily propositional, representing *know-that*; understanding is primarily a kind of *ability.* Knowing that cryptic dynamics can reproduce decoupled cycles despite strong trophic links is one thing, but knowing *how* they do so is another: the latter involves the capacity to be able to *do something* with that knowledge; the former with our capacity to represent it. Undoubtedly these kinds of knowledge are tightly-knit, and determining the relationship between them is tricky, but for my purposes it is irrelevant whether one ultimately reduces to the other, or what their precise relationship is (see, for instance, Strevens [2013]). I'm interested in how limited human agents gain knowledge from particular scientific practices, and on my view the know/how distinction is crucial in that context, regardless of what their fundamental relationship is. What matters here is that bottle-experiments are perfect for the generation of a certain kind of know-how. But what know-how might bottle-experiments teach us? I think they grant us a kind of *modal ability*, specifically: the capacity to navigate modal space.

Soazig Le Bihan ([2016]) has developed such an account of understanding, and I'll sketch it here. Her starting point is the oft-noticed tendency of scientific representations to misrepresent in profound ways. And this leads to an oft-tackled puzzle: many scientific goods, such as explanation, have close ties to the truth, and so how can they be provisioned via

---

[13] For general discussion of the role and nature of scientific understanding, see the papers collected in De Regt, Leonelli & Eigner [2008].

misrepresentation? One common response appeals to the heuristic value of such theories and models: they make it *more likely* that we'll get to the relevant truths (Odenbaugh [2005] suggests this in an ecological context). But such a response is unsatisfying: scientists don't appear to treat misrepresentations as a mere means; a stepping stone on the way to truth, rather, scientists often happily embrace a wide range of idealized, abstract apparent misrepresentations (Angela Potochnik [2017] makes this point particularly forcefully). Le Bihan's solution is to claim that such representations generate a kind of knowledge, specifically: the capacity to navigate modal space.

Different phenomena occur across different sets of possible worlds. There are conditions (necessary, sufficient, probabilistic) which the production of phenomena depend upon; and conditions under which phenomena transform or disappear. But beyond knowing *that* certain phenomena have such dependencies; knowing *how to navigate* such spaces is a crucial part of scientific knowledge. Consider the relationship between cryptic dynamics and trophic links: under certain conditions, hidden population-structures can generate the appearance of decoupling between predators and prey despite their strong trophic relationship. In applying observational techniques, the ecologist must understand how various features of ecological systems might lead them to exhibit standard or cryptic dynamics. This does not merely rely on propositional knowledge about nature of systems and ecological theory, but requires the application of particular kinds of know-how pertaining to both observational techniques and the modal properties of ecological systems.

On Le Bihan's schema, we can analyze this modal understanding in terms of a possibility space S generated by a phenomenon's domain, *P*, and the *dependency structures* which determine subsets of *P* to be the case within S. For the case at hand, let's take our phenomenon domain to be circumstances involving tight trophic coupling. In some subset of that domain, coupling will correlate with the kinds of dynamics which observation techniques detect: well behaved predator-prey cycles, for instance. In others, coupling will correlate with cryptic cycles;

where observation techniques will fail to detect trophic coupling. There are a series of dependency relationships governing under what circumstances trophic interactions will generate different cycles. The modal space, then, consists of:

S, a possibility space;

P, the domain of *S* in which tropic interaction takes place.

P_1, the subset of *P* where trophic interaction generates well-behaved dynamics;

P_2, the subset of P where trophic interaction generates cryptic dynamics;

A set of *dependency structures* concerning *S,* which determine whether *P* holds, or whether P_1 or P_2 hold[14].

What, then is it to have modal understanding of *P* for Le Bihan?

Constitutive of modal understanding is the possession of "… 'navigating skills': the ability to 'navigate' (some of) the possibility space for some domain of phenomena *P*" (117). Le Bihan identifies various levels of know-how relating to navigation: first, how dependency structures generate *P* (that is, modal knowledge of the dependency structures), second, how different dependency relations interrelate (that is, modal knowledge of *P*), and third, understanding global constraints on *S* (see figure 3). In combination, these different modal abilities allow us to navigate within particular dependency structures, consider how they'll interact in different contexts, and grasp the general features necessary for *P.* We can clothe this abstract schema by considering Yoshida et al's work.

As we saw above, Yoshida et al's theoretical model captures various modal dependency structures. For instance, they suggest that one P_2 dependency requires defense costs (being a

---

[14] An alternative reading might index the phenomenon to different kinds of population dynamics, instead of to trophic interactions—this doesn't make any important difference to the result.

tough) to be sufficiently low, given competition with softs. They present this as (something like) a necessary (but insufficient) condition for P_2 to hold. This partly constitutes their understanding of how to navigate the relevant modal spaces. Another level of understanding involves how different population structures in ecology potentially generate very different subsets of phenomenon. For instance, homogeneous prey populations are likely to generate the well-behaved dynamics of P_1, while heterogeneous populations are more likely to generate P_2. If the argument in section 4 holds true, and ecological systems are likely to be heterogeneous (both in Matthewson's intra-system sense, and Elliott-Graves' inter-system sense), then we should expect much of the ecological world to not be amenable to observation techniques.

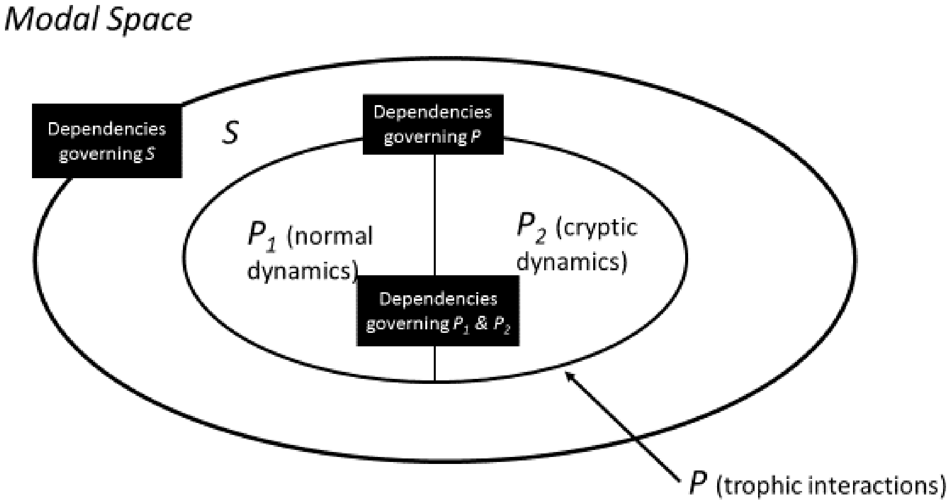

On such a view, then, some scientific investigations have epistemic (rather than merely heuristic) value in virtue of providing understanding, here understood as modal navigating capabilities. In the next section, I'll argue that bottle experiments are particularly well-placed to "… provide us with some understanding of the phenomena by providing us with some knowledge of how to navigate the possibility space" (122).

5.2 Bottle Experiments as Understanding-Generators.

Let’s take stock. I’ve thus far provided a case-study to ground our discussion of the merits of bottle-experiments. There, Yoshida et al triangulated between two experimental systems and a simple model to explore cryptic dynamics in ecology. They took their results to cause trouble for observational techniques: the practice of inferring trophic interactions from observed population dynamics. The view I’ve called ‘extrapolationism’ takes the epistemic value of bottle experiments to lie in demonstrating ecological actuality in an artificial system. That is, by generating cryptic dynamics in a controlled experiment, the hypothesis that they might be instantiated in a natural system passes a test. However, I have argued that the heterogeneity of ecological systems should lead us to expect ecological actuality to be broad and diverse, and thus we should expect significant overlap between ecological possibility and ecological actuality. Although that argument requires significant assumptions, I think it shows that a purely extrapolationist defense of bottle experimentation is limited.

I’ve just zeroed in on an epistemic good that I’ll argue bottle experiments are apt for generating: the know-how involved with modal-navigation. It’s worth noting that given the heterogeneity of ecological systems, such know-how is both particularly hard-won, and extremely valuable. Making sense of how various ecological dynamics might be instantiated in a range of circumstances—what dependency structures there are and how they might be interrelated—is precisely the kind of knowledge we need to explain the behavior of these tricky systems. So, why might we think that bottle-experiments are particularly suitable for this task? To see this, let’s briefly consider how we gain know-how of the sort concerning us here.

Learning how to navigate modal spaces involves *concretizing* those spaces. By ‘concretizing’, I do not necessarily mean making physical, but rather *making explicit*. We start from a set of rather abstract ideas and assumptions. For instance, echoing earlier in the discussion, we might think that it is plausible that different trophic interactions and population structures might underdetermine patterns in ecological systems. But seeing how these ideas and assumptions play out in a domain requires that we specify them in more detail. We need a way of representing, specifying—*concretizing*—those assumptions. For instance, a coupled differential equation is one way of representing a set of assumptions about the causal interactions between trophic levels. Positing such an equation sets a particular modal domain, and allows us to explore it. That is, it underwrites our learning how to navigate the domain specified by the equations. We see the same in bottle experiments. But—crucially—the mediums are different, the way we specify our assumptions are different, and both the modal domains, and the way we explore and navigate these domains are different: thus, I argue, different understandings are generated.

Another way of putting this point is to see that the construction and running of bottle experiments involve a different set of *constraints* than observing natural phenomena and running field experiments on the one hand, and using theoretical models and simulations on the other. Bottle experiments provide a set of constraints which make them well-suited to generating understanding of ecological dynamics. To demonstrate this, I'll briefly contrast bottle experiments with studies of naturally occurring systems and with mathematical models. Briefly, whereas bottle experiments are both simpler and more manipulable than natural systems, they are (in some relevant sense) more constrained than mathematical models. And these divergences facilitate the generation of rich knowledge about the dependencies involved in ecological dynamics.

Recall that I'm drawing on a notion of understanding where it consists in having the capacity to navigate a possibility space pertaining to some domain on the basis of understanding various dependencies within those spaces. The possibility spaces pertaining to ecology are wide indeed, and understanding them involves various approaches and strategies. Bottle experimentation is one such strategy. They are not the only strategy – indeed, the exploration of natural systems and theoretical systems are also required. But when we compare these, the 'sweet spot' bottle experiments occupy becomes clear.

Bottled systems are simpler than natural systems. Where natural systems instantiate a wide range—a chaos—of causal interactions which often require long periods of data collection to begin to unpick, bottle experiments are highly simplified, easily graspable systems. This allows two things. First, the relevant dependency relations are laid bare in clear, accessible ways. Consider the relationship between population structures in prey and cryptic cycles. In the bacteria-phage bottle system, because it is possible to mark the phage-resistant bacteria, we can track the compensation between toughs and softs within the bacterial population, as well as the relationship between it and the phage population. Doing the same in a wild population would be extremely tricky, both in terms of scale and causal heterogeneity: small scale allows the dynamics to unfold before our eyes, causal homogeneity allows us to identify much more precisely the dependencies governing the system's behavior.

Second, these dependencies are manipulable. As Yoshida et al could vary when to introduce predators to the experiment, they could run tests of their theory on the algae-rotifer system which asked after the dependencies posited in their hypothesis. As we saw above, cryptic population dynamics predict that an unpredated prey population should become dominated by softs, and upon the introduction of predators should switch to normal predator-prey oscillations, before exhibiting cryptic dynamics once tough numbers have increased sufficiently. Although some control is possible in natural systems, this is much amped up in the lab. The combination of tractability and manipubility make bottle experiments far preferable to natural systems for learning how to navigate ecological space. Where nature's heterogeneity will swamp and hide the relevant dependency relationships, bottle experiments lay them out.

These features—simplicity and manipubility—are shared between bottle experiments and formal models—particularly simulations studies – and so, why use bottle experiments as opposed to computer-run models? The answer is that the two systems complement one another.

In contrast with theoretical systems, bottle experiments are physical, experimental, systems and, as such, require difficult work in their design, development and ongoing maintenance. The chemostats that house Yoshida et al's rotifer-algae systems are complex apparatus requiring ongoing maintenance, as does keeping the critters within them alive. Raising, and experimenting with, such critters involves intimate knowledge of their ways. Undoubtedly, simulation work is also challenging and involves intimate knowledge of systems. But crucially, this involves different knowledge of different kinds of systems. These differences aid in our understanding of ecological dynamics in two ways.

First, one can consider the bottle experiment as building on the robustness of Yoshida et al's result[15]. As a particular model is just one way of concretizing a modal space, we want to know whether dynamics are similar across different models which pick out similar spaces in different ways. Concretizing the dynamics by using a differential equation as well as in the experimental system suggests that the results are not simply due to a quirk in the particular system – a quirk in that particular concretization exercise – but is a common features of ecological spaces. The ways in which bottle experiment behavior and theoretical model behavior converge, then, inform us about the resilience of the dependencies across the relevant modal spaces.

Second, the simple physicality of the bottle experiment leads to the exploration of a subtly different area of possibility space and a different set of dependencies. Different access, different properties, and different mediums guide the scientists in exploring different modal spaces in differing ways. Although switching from simple differential equations to, say, an agent-based simulation is a sizable departure, switching to a physical system involving living critters is much more radical. In short, concretizing a modal space in multiple ways is crucial for understanding,

---

[15] For discussion of robustness, see Calcott [2011], Matthewson & Weisberg [2009].

and the physicality of bottle experiments provide a quite different way of concretizing when compared to theoretical systems.

None of this says that the knowledge bottle experiments provide are unique *in principle*, but regardless I think such considerations make it plausible that they are an excellent way of gaining the know-how required to navigate ecological space. It also paints a picture of ecology in an abstract, non-extrapolationist key: ecological knowledge is of ecological dynamics generally speaking (rather than merely their instantiation in natural systems) and of the assumptions underlying observational techniques and other empirical tools. This suggests that my construal of ecology in section 2—as concerned narrowly with explaining patterns in natural ecosystems by positing connections between populations—was far too quick. Ecologists are also interested in theoretical questions about ecological dynamics which need not bear on natural systems.

To clarify the position, let's pause and contrast it with Griesemer & Wade's (1988) view on 'experimental evolution', which is readily adapted to ecological cases. They are primarily interested in how species selection can be understood generally by being realized in microcosms.

Griesemer & Wade use a version of the '*vera causa* ideal' to frame their discussion (for more recent investigation, see Novick [2016], Novick & Scholl [forthcoming]). The principle is a set of conditions governing when causal explanations are complete or successful. For a putative cause to be *vera causa*, we must establish (1) the 'competence' of the cause: if it were to occur, then the explanandum would follow; (2) the 'reality' of the cause—we should find evidence of it occurring in contexts beyond the explanandum case. Yoshida et al's work can be adapted fairly easily: they are interested in whether cryptic population structure might explain some patterns of population abundance. By realizing these dynamics in an experimental system they show that such structures may produce such dynamics (competence), and do so independently of the examples which motivated their investigation (reality). So far so good: this is a sophisticated way of saying that lab results demonstrate ecological actuality. However, Griesemer & Wade argue that the value of such demonstration is their role in *cause detection.*

A 'cause detector' is an instrument for training our expectations about the results of interactions in a set of systems:

> After experimenters have used a laboratory system in this way, they have a set of expectations about the relation between agency, the system's structure, and their effects which they can use to formulate causal hypotheses about other [natural] systems. (84).

Applied to Yoshida et al, the purpose of the experiments are to provision understanding of populations in ecology such that hypotheses about, say, under what conditions ecological systems will conspire to undermine observation techniques, may be generated.

My account of the value of bottle experiments is broader and, I think, encapsulates Griesemer & Wade's analysis. For them, to an extent experiments are still ultimately geared towards natural systems—after all, such experiments underwrite hypotheses about *natural causes.* In contrast, I have argued that the scope of ecology is broader, in two senses. First, the realm of ecological concern extends beyond the natural: in my view ecology is a deeply theoretical, modal science, concerned with much more than natural systems[16]. Second, ecological understanding concerns more than causes, including a wider set of modal dependency relationships[17]. Third, the generation of modal understanding explains *how* vera causa reasoning succeeds in bottle experimentation: by generating competent and real causes in controlled circumstances ecologists are able to learn how various properties of a system affect the dynamics it may instantiate. Cause detection, and other tasks, are thus enabled.

---

[16] Such a view resonates strongly with Marc Lange's view on ecological laws ([2005]).

[17] I'm here assuming that causation is a kind of modal relationship. I think this is a very plausible metaphysical view, but nothing crucial to my argument turns on this.

My argument that bottle experiments generate modal understanding, then, involves claiming that (1) the range of ecological interest is beyond naturally occurring systems; (2) ecological understanding involves a wide range of modal dependence relations pertaining to ecological dynamics; and (3) this modal know-how itself underwrites diverse downstream epistemic goods, only some of which pertain to natural systems. At base, the purpose and value of bottle experiments is in generating know-how about modal navigation[18].

Work focusing on understanding in science has taken its cue from the theoretical, not the experimental. Potochnik, Le Bihan and others start from the misrepresentative nature of scientific representations. Should we, then, see bottle experiments as an instance of the modeler's strategy (Godfrey-Smith [2006], Weisberg [2012])? Odenbaugh ([2005]) and Parke ([2014]) consider such activities to be experiments, while Levy and Currie ([2014]) argue that, properly speaking, they are physical models. I think my defense of the pursuitworthiness of bottle experiments turns on both their informing us about natural systems and the generation of understanding. This suggests that bottle experiments occupy a middle ground between the two. Indeed, if we delineate the experimental and the theoretical in terms of their epistemic outputs, and some practices have multiple such outputs, then the value of that way of distinguishing between modeling and experiment might be limited when applied to those practices (however, we shouldn't take the existence of indeterminate cases as alone undermining such distinctions!).

Experimental systems aid in nature-directed and more theoretical pursuits in virtue of occupying a sweet-spot in the generation of understanding. As physical systems, bottle experiments can guide and inspire theorizing, and generate empirical data, which mathematical models cannot. Because they are simple and controlled, they are more casually transparent and

---

[18] This suggests that bottle experiments may fail due to insufficiently capturing the intended dynamics, as opposed to not representing the world. Jim Griesemer's discussion ([1988]) provides such an example of an experiment being insufficiently 'pure', as does the objection considered in Griesemer & Wade ([1988]). I take it that examples of scientists objecting to bottle experiments on these kinds of grounds are grist to my mill.

amenable to study than natural systems. This makes them superb tools for exploring the nature of ecological dynamics and the assumptions underlying ecological theory. Even when truths about natural systems are not in the offing, grasping these other aspects of ecology pays significant epistemic dividends.

5.3 How Understanding Travels.

An objection: the dissemination of scientific knowledge is not produced via running experiments, but by publishing experimental results. If the epistemic good at hand is primarily an ability, how can the mere reporting of experimental procedures and their outcomes generate know-how? Happily, Sabina Leonelli ([2016]) has recently discussed a similar question, and I'll sketch and co-opt her answer.

Leonelli asks how the bountiful data generated by experimental model organism work are disseminated via databases; that is, how data travels from the context of its generation to other domains. On her view, database practices involve (1) using categorization to decontextualize data from the idiosyncrasies of origin; and then (2) facilitating recontextualization for use in new circumstances. For Leonelli, two things are necessary for recontextualization to succeed. First, meta-data must record information about the data's origin and subsequent travel, such that it can be made sense of and critically appraised by scientists. And second, those scientists must possess the relevant embodied knowledge in virtue of having, for instance, run similar studies to those from which the data emerged. On Leonelli's view, data journeys require both good recording practices, and the audience having the relevant expertise. In the good case, because the audience—the scientists utilizing the data base—have themselves had the relevant experience necessary to understand, criticize, and utilize the data, they are able to successfully reuse the data in new contexts.

Although the high-throughput biological cases Leonelli analyses are significantly more developed in terms of recording practices—a single publication is no data-base!—the same basic ideas are applicable here. Even if it is in virtue of being involved with the design, set-up and running-of experiments that Yoshida et al have generated the epistemic good, it doesn't follow from this that those goods cannot be transported via publication. As Leonelli makes explicit, the capacity to travel depends in part on the audience's knowledge. Ecologists undergo quite different types of training: some focus more on mathematical modeling, others on fieldwork, and others on the experimental work that concerns us here. But regardless, there is significant overlap. This overlap enables them to make sense of, and to learn the lessons of, Yoshida et al's work. Moreover, the diversity of the different training allows ecologists to draw lessons from the paper which Yoshida et al themselves might not be able to see: the plurality of embodied knowledge throughout ecology allows ecologists to draw different aspects from a single study.

6. Conclusion

I've argued that the vindication of ecological bottle experiments lies in their capacity to generate understanding of ecological dynamics as well as their capacity to generate truths about natural systems. Generally speaking understanding and truth are not independent, nor mutually exclusive. A defense of the pursuitability of bottle experiments, then, may draw on both.

I have focused on how bottle-experimentation in ecology promotes ecological understanding. But this vindication might be particularly strong when the experiments play dual roles: acting as empirical bridges between theory and the natural world as well as generating modal understanding. Indeed, given the heterogeneity of ecological systems, good expectations about the modal nature of the ecological world—good modal navigation skills—are likely a prerequisite for making meaningful empirical contact with them. We musn't relieve ourselves of both baby and bathwater by decoupling experiment and nature too quickly: natural phenomena still matter. Indeed, it is plausible that an understanding of ecological dynamics will themselves help us in our empirical investigations of complex and heterogeneous natural systems. Yoshida et al posit a set of dependencies governing cryptic dynamics generally speaking—these could potentially form the basis of more direct investigations involving ecological systems in play.

Even in circumstances where experimentation's capacity inform us about natural ecological systems is limited, such studies are often well placed to provide other epistemic goods—as I have suggested, to develop the capacities required to understand ecological dynamics and the assumptions underlying their epistemic tools. Bottle experiments generate understanding.

Bibliography

Calcott, B. (2011). Wimsatt and the robustness family: Review of Wimsatt's Re-engineering Philosophy for Limited Beings. Biology & Philosophy, 26(2), 281-293.

Cartwright, N. (1999). The dappled world: A study of the boundaries of science. Cambridge University Press.

Chang, H. (2012). Is water H2O?: Evidence, realism and pluralism (Vol. 293). Springer Science & Business Media.

Constantino, R. F., R. A. Desharnais, J. M. Cushing, and B. Dennis (1997). ChaoticDynamics in an Insect Population. Science 275: 389–391.

Currie, A., & Walsh, K. (2018). Newton on Islandworld: Ontic-Driven Explanations of Scientific Method. Perspectives on Science, 26(1), 119-156.

Currie, A. (2018). Rock, Bone, and Ruin: An Optimist's Guide to the Historical Sciences. MIT Press.

Currie, A. (2015). Philosophy of Science and the Curse of the Case Study. In The Palgrave Handbook of Philosophical Methods (pp. 553-572). Palgrave Macmillan, London.

De Regt, H. W., Leonelli, S., & Eigner, K. (Eds.). (2009). Scientific understanding: Philosophical perspectives. University of Pittsburgh Press.

Elliott-Graves, A. (2016). The problem of prediction in invasion biology. Biology & Philosophy, 31(3), 373-393.

Franklin, A. (1993): "Discovery, Pursuit, and Justification" Perspectives on Science 1: 252–284.

Godfrey-Smith, P. (2006). The strategy of model-based science. Biology and philosophy, 21(5), 725-740.

Griesemer, J. R. (1988). Causal explanation in laboratory ecology: the case of competitive indeterminacy. In PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association (Vol. 1988, No. 1, pp. 337-344). Philosophy of Science Association.

Griesemer, J. R., & Wade, M. J. (1988). Laboratory models, causal explanation and group selection. Biology and Philosophy, 3(1), 67-96.

Havstad, J. (forthcoming). Messy Chemical Kinds. British Journal for Philosophy of Science.

Jessup, C. M., Kassen, R., Forde, S. E., Kerr, B., Buckling, A., Rainey, P. B., & Bohannan, B. J. (2004). Big questions, small worlds: microbial model systems in ecology. Trends in Ecology & Evolution, 19(4), 189-197.

Lange, M. (2005). Ecological laws: what would they be and why would they matter?. Oikos, 110(2), 394-403.

Laudan, L. (1977): Progress and Its Problems. London: Routledge

Le Bihan, S. (2016). Enlightening Falsehoods. Explaining Understanding: New Perspectives from Epistemology and Philosophy of Science. Grimm, S. R., Baumberger, C., & Ammon, S. (Eds.). Routledge, pp 111 – 135

Leonelli, S. (2016). Data-centric biology: a philosophical study. University of Chicago Press.

Levins, R. (1966), "The Strategy of Model Building in Population Biology," American Scientist 54: 421-431.

Levy, A., & Currie, A. (2014). Model organisms are not (theoretical) models. The British Journal for the Philosophy of Science, 66(2), 327-348.

Love, A. C., & Travisano, M. (2013). Microbes modeling ontogeny. Biology & Philosophy, 28(2), 161-188.

Matthewson J (2011) Trade-offs in model-building: a more target-oriented approach. Stud Hist Philos Sci Part A 42(2):324–333

Matthewson, J., & Weisberg, M. (2009). The structure of tradeoffs in model building. Synthese, 170(1), 169-190.

Morgan, M (2005). Experiments versus models: New phenomena, inference and surprise. Journal of Economic Methodology 12 (2):317-329.

Novick, A & Scholl, R. (forthcoming). Presume it not: True causes in the search for the basis of heredity. British Journal for the Philosophy of Science

Novick, Aaron. (2016). Metaphysics and the Vera Causa Ideal: The Nun's Priest's Tale. Erkenntnis: 1-16.

Nyrup, R. (forthcoming). Of Water Drops and Atomic Nuclei: Analogies and Pursuit Worthiness in Science. British Journal for Philosophy of Science.

Odenbaugh, J. (2006). Message in the bottle: The constraints of experimentation on model building. Philosophy of Science, 73(5), 720-729.

Odenbaugh, J. (2007). Seeing the forest and the trees: Realism about communities and ecosystems. Philosophy of Science, 74(5), 628-641.

Odenbaugh, J. (2005). Idealized, inaccurate but successful: A pragmatic approach to evaluating models in theoretical ecology. Biology and Philosophy, 20(2-3), 231-255.

Parke, E. C. (2014). Experiments, simulations, and epistemic privilege. Philosophy of Science, 81(4), 516-536.

Potochnik, Angela. (in print). Idealization and the Aims of Science, The University of Chicago Press, forthcoming.

Potochnik, A. (2015). The diverse aims of science. Studies in History and Philosophy of Science Part A, 53, 71-80.

Sterelny, K. (2006). Local Ecological Communities*. Philosophy of Science, 73(2), 215-231.

Strevens, M. (2013). No understanding without explanation. Studies in history and philosophy of science Part A, 44(3), 510-515.

Waters, C. K. (2007). Causes that make a difference. The Journal of Philosophy, 104(11), 551-579.

Weisberg, M. (2012). Simulation and similarity: Using models to understand the world. Oxford University Press.

Wootton, J. T., & Emmerson, M. (2005). Measurement of interaction strength in nature. Annual Review of Ecology, Evolution, and Systematics, 419-444.

Yoshida, Takehito, Stephen P. Ellner, Laura E. Jones, Brendan JM Bohannan, Richard E. Lenski, and Nelson G. Hairston Jr. "Cryptic population dynamics: rapid evolution masks trophic interactions." PLoS Biol 5, no. 9 (2007): e235.