

Local Ecological Communities*

Kim Sterelny^{†‡}

A phenomenological community is an identifiable assemblage of organisms in a local habitat patch: a local wetland or mudflat are typical examples. Such communities are typically persistent: membership and abundance stay fairly constant over time. In this paper I discuss whether phenomenological communities are functionally structured, causal systems that play a role in determining the presence and abundance of organisms in a local habitat patch. I argue they are not, if individualist models of community assembly are vindicated; i.e., if the presence of one species is not typically explained by the presence or absence of specific other species. I discuss two alternatives to individualism, and conclude by arguing for a dimensional model of phenomenological communities. The causal salience of a phenomenological community depends on three factors: the extent to which it is internally regulated, the extent to which it has robust boundaries, and the extent to which it has emergent properties. I conclude by using this model to frame a natural research agenda for community ecology.

1. Phenomenological Communities. Organisms interact in many ways. Birds, for example, variously eat the seeds of trees, eat insects that infest them, disperse their seeds, and use them as perches and nest sites. These innumerable interactions collectively constitute the economy of nature: they make up the distinctive demographic and geographic patterns of ecology. One such pattern is the heterogeneity of typical landscapes. Most regions are patchy: they are composed of local habitats of different types: local wetlands, parks, farms, local bush reserves, river corridors and the like, each with their distinctive inhabitants. These assemblages of popu-

*Received August 2005; revised June 2006.

[†]To contact the author, please write to: Kim Sterelny, Philosophy Program, Victoria University of Wellington, PO Box 600 Wellington, New Zealand; e-mail: Kim.Sterelny@vuw.ac.nz, or to Philosophy Program, RISS, Australian National University, 0200 ACT, Australia; e-mail kimbo@coombs.anu.edu.au.

[‡]Thanks to Mark Colyvan, Michael Weisberg, Tim Lewens, Jay Odenbaugh, Peter Godfrey-Smith, John Odling-Smee, and three referees and the editor of this journal for very helpful comments on an earlier draft of this paper. The dimensional way of thinking about communities was suggested to me by Peter Godfrey-Smith's parallel treatment of Darwinian populations. Thanks also to the audience of HPS Cambridge for their helpful feedback on a presentation based on this material.

Philosophy of Science, 73 (April 2006) pp. 215–231. 0031-8248/2006/7302-0005\$10.00
Copyright 2006 by the Philosophy of Science Association. All rights reserved.

lations, each in their distinctive patch, are local ecological communities. The ecology of a region is composed of a patchwork of such local communities. These are *phenomenological communities*: persisting assemblages of organisms in spatial association within some identifiable chunk of habitat. Typically, these assemblages are reasonably stable in their composition year by year. That stability underwrites the common practices of producing local field guides, floras, natural history notes and environmental impact reports. One important dispute in ecology concerns the extent to which these phenomenological communities are causal systems; whether *communities themselves* are structured, functionally organized or integrated systems that play a role in determining the presence, abundance and fate of the populations out of which they are composed (Cooper 2003).

Suppose, for example, that a local population of banksias and another of eucalypts in Black Mountain, a bush reserve near the ANU, are associated only because both populations happen to tolerate the temperature, soils and rainfall characteristic of this location. Of course local populations do not live completely independently of one another. Life at Black Mountain is diffusely interdependent; the organisms that live there are not autonomous islands of life. Plants often depend on animals for pollination and seed dispersal, on symbiotic partners for crucial nutrients, and on detritivores for nutrient cycling. Consumer guilds—herbivores and carnivores—are obviously dependent on other organisms. But though organisms depend on their biological as well as their physical environment, it remains possible that the biological conditions of life are typically coarse grained, with species having broad banded biological conditions of existence. On this *individualist* view of communities, the presence of one population makes little difference to the abundance of the other. If most of the Black Mountain populations are largely indifferent to one another's presence and abundance, and the populations are associated mainly because they happen to tolerate similar physical conditions, then the Black Mountain community is not an organized, structured system. For while there might be a more or less deterministic explanation of why particular species are represented in Black Mountain—soils, rainfall and temperature make it hospitable to some members of the regional species pool but not others—these explanations do not appeal to the presence or absence of specific species. The community itself is not a causal system whose characteristics constrain membership or population size. Species in the regional pool are neither excluded nor incorporated into the Black Mountain assemblage as a result of strong interactions with other community members.

So the view that phenomenological communities are integrated causal systems seems committed to rejecting individualism. Populations in a local

community must matter to one another. Moreover, causal systems have boundaries and that also poses something of a challenge. Black Mountain consists of a hill and some surrounding area, covered in bush. In some moister gullies, there is a relatively closed canopy. In other areas, it is open forest. Near the edges there is open woodland with grasses, weedy annuals, shrubs and a scattering of trees. In some parts there has been an invasion of exotic weeds, but in most areas, the vegetation resembles that of pre-European settlement. But as one moves from place to place, these changes in character are gradual rather than sharp. So "Black Mountain" names a quite heterogeneous region of about 10 square kilometers with gentle variation from patch to patch. As a consequence of these gradual changes in character, the different populations might not be correlated. A local brushtail possum population may overlap with a local ringtail possum population, a local boobook owl population, a greater glider population, and a number of eucalyptus populations. For on Black Mountain, there are no sharp changes that matter to all of these species, keeping local populations congruent with one another.

So are Black Mountain and similar local patches merely phenomenological communities, just part of the descriptive phenomenology of biology? If so, while such communities have correct natural history descriptions, they have no objective bounds in space or time, and nor do they have explanatorily salient organization. This individualist perspective on phenomenological communities poses the problem of this paper: do communities themselves have causally salient, functional properties?

In answering this question, I need first to make clear the notion of function and organization in play. In evolutionary biology, functional traits depend on selection histories, and communities do not have selection histories. But Robert Cummins has shown that there is an alternative conception of function: a part of a system has a Cummins function when its activity makes a distinctive, stable contribution to the system as a whole (Cummins 1973; Godfrey-Smith 1993, 1994). Communities may be organized systems: as a whole they behave in complex and interesting ways, and their various components have Cummins functions. A keystone predator, for example, might maintain diversity amongst the consumers in a community by limiting the population size of a species which would otherwise exclude the others (Paine 1966). Starfish are not selected to maintain diversity by eating mussels, but within that community, diversity is a stable effect of this particular activity. So our question becomes: are phenomenological communities typically organized systems? Are they stable, bounded, and with enduring global features of biological importance to which particular components make a regular contribution?

In the rest of this paper I assess individualism and contrast it with two alternative proposals. One suggests that local communities are causal sys-

tems because they are internally regulated. The other suggests that local communities are causal systems because they have emergent properties—community ensembles have effects that cannot be explained by appealing to the biological properties of their components. In considering these alternatives, I try to rely on relatively uncontroversial ecology. I am interested in whether the individualist view of local communities—that they are merely phenomenological—can be rejected and replaced on the basis of common ground ecology. In Sections 4 and 5 I relate this discussion to the problem of boundaries. If local communities are causal systems, must they also be bounded? If so, what constitutes those boundaries? I then argue for the heuristic utility of a *dimensional perspective* on local communities. We can productively think of specific local communities as located in a three dimensional space, a space whose dimensions track ensemble effects, the extent of internal integration, and boundedness.

2. The Internal Organization of Communities. According to one line of thought, communities are organized if they are internally regulated. In turn, communities are internally regulated if their components interact in ways that constrain their own size, and which filter potential members of the community (in which case the community has internally determined “assembly rules”). It might seem that communities *must* be regulated in this sense. For it is possible to produce local field guides. Yet field guides are possible only if community membership is stable. But this stability cannot be explained by abiotic factors. Their impact is too variable. Rainfall, for example, varies dramatically on Black Mountain from season to season, and so too does the incidence of fire. This biota does not inhabit a physically invariant landscape. Yet it is roughly stable in both composition and abundance. As we shall see in the next section, there is a way of decoupling issues about the stability of communities from those about the stability of component populations. But one reasonable way of defining stability is in terms of membership. Stability in that sense, the idea runs, is best explained by the hypothesis that communities are regulated by “density dependent” interactions. The size of some given population—say, superb fairy wrens on Black Mountain—will fluctuate within bounds only if the factors that limit the fairy wren population become *more intense* as the population rises, and *less intense* as it falls. Yet abiotic factors are not intrinsically sensitive to the size of the populations they effect. Fire burns through fairy wren habitat without regard to the number of wrens present. In contrast, competition between the wrens for limited resources is density dependent. The more wrens, the harder such limits bite. Competition is bound to get more intense as population size increases, and less intense as it dips (see, for example, Cooper 1993, 2001,

2003; Pimm 1991). In brief, community stability is the result of a “balance of nature,” a balance deriving from the internal regulation of communities.

The persuasiveness of this argument depends on the idea that if the forces that affect a local population were to act independently of its size, it would be an amazing coincidence if abundance did not change over time. For very slight tendencies to increase or decrease result in crashes or plagues. An example of Haldane’s illustrates this point very vividly. For animals that breed once a year, a per generation increase of 1.01 would result in a population 21,000 times the starting population after a thousand years. Likewise a slightly negative rate of .99 would reduce a population to .000043 of its original number after the same period (Hixon et al. 2002, 1492). If populations persist, something internal to communities must damp down such fluctuations. There must be networks of biological interaction which filter membership and which constrain the demography of their members. This line of thought really represents a family of views that accord a “balance of nature” a central role in regulating the internal organization of communities. For as Stuart Pimm has shown, the notion of stability is itself ambiguous (Pimm 1991; Lehman and Tilman 2000). There is a family of notions, and hence a family of ideas about the extent to which communities are stable. But according to all the views in this family, communities have a genuine organization. The organisms already present and interacting make certain roles available and foreclose others.

Greg Cooper is rightly skeptical of this whole class of arguments: they depend on a crucial ambiguity (Cooper 2003). There is an undemanding sense of “stable,” where it means something like “the persistence of community membership.” On this reading, communities are indeed typically stable. The species composition on Black Mountain is similar year by year. But as we shall see, stability in this sense may not require community organization. There are two stronger notions of stability. One is the stability not just of community membership but population size. If communities are at true equilibrium, with population sizes varying only in minor ways around a mean to which they return, then they must indeed be regulated. But so understood, there is no reason to believe communities are stable. Another stronger notion is *long run* persistence of community membership: persistence over many generations. If community membership is stable over long periods, the inference to regulation looks plausible. But once more, in this sense communities may not be stable. In short, if there is evidence of limited variation around a mean in the population sizes of components of the community, or if there is evidence of long term stability in community membership (as some paleoecological studies claim: see Eldredge 2003), then we do indeed have evidence of equilibrating mechanisms. But neither limited movement around a mean nor

very long term persistence in specific local communities are obvious and established facts of ecological life.

There is an established fact: medium term persistence in community membership. But persistence on this scale may not depend on internal regulation (though the issues remain both technical and fraught: see Turchin 1999). The alternative to internal regulation is stabilization through migration: communities are typically demographically open, and thus a local population may persist by recruiting from neighboring communities. Open communities may be stabilized by “metapopulation dynamics.” If, for example, echidna populations vary independently of one another in a cluster of adjacent communities, a population fluctuating towards extinction can be rescued by migration from a neighboring community. Migration between communities can protect unregulated communities from random walks to extinction. The effects of density dependent internal regulation can be coarsely mimicked by a metapopulation of unregulated communities, provided that this metapopulation is spread over a heterogeneous landscape and provided that migration between populations is possible. So long as the trajectory of populations within a cluster are independent of one another, the stability of the metapopulation ensemble will be greater than the stability of a typical population within the ensemble. Populations without density dependence can persist for many generations (Murdoch 1994). We do not know the extent to which metapopulation dynamics explain the evident stability on which field guides depend. But the existence of this mechanism means that we cannot assume that persisting communities are internally regulated.

It is time to take stock. The biological characteristics of local habitats are fairly stable. Farming would be impossible if that were false. But while this fact is suggestive, in itself it is not sufficient to show that local assemblages are internally organized ensembles. For it may turn out that the stability that is observationally obvious is typically both coarse grained and short term. And if that were to be the characteristic extent of stability in local communities, metapopulation effects may explain it. In the next section, I focus on the idea that communities are causal systems because they have biologically important emergent properties.

3. The Emergent Property Hypothesis. A claim: communities are causal systems if those systems themselves have properties that are causally important in driving biological processes. I shall discuss this idea by exploring a family of famous hypotheses that link the diversity of a community to its stability. The crucial claim of the emergent property hypothesis is that these emergent properties are causally important: they drive ecological processes. The diversity-stability hypothesis is one attempt to show this. The idea that diversity adds stability to a community has

enormous intuitive plausibility. For diversity adds redundancy, and hence allows that community to survive fluctuations in the fortunes of its members. If only one population on Black Mountain pollinates the gumtree, *Eucalyptus rossii*, and if it were to suffer a serious decline, *rossii* would be unable to recruit new plants into its population. However, if there is a suite of eucalyptus pollinators, a fluctuation in one population would not ramify through the community as a whole. Redundancy buffers disturbance, and diversity adds redundancy.

This appealing picture seemed to be undermined by theoretical work of Robert May (1973). His models showed that more diverse communities were less stable, not more stable. However, the last decade has seen a revival of the diversity-stability hypothesis and its close relative, the idea that more diverse communities are more productive. The crucial move has been to sidestep May's result by a changed idea of the property that diversity stabilizes. May took the diversity-stability hypothesis to be a hypothesis about *population size*: in more diverse communities, the populations of the component species are more stable. However, David Tilman and others have argued that some *community level properties* are more stable in more diverse communities. In particular, Tilman argued that the biomass of more diverse communities is more stable than that of less diverse ones. For somewhat similar reasons, a diversity-productivity relationship looks plausible. For one thing, if diversity stabilizes productivity in response to disturbance, stability itself is positively correlated with productivity over time, for productivity will decline less in the face of drought, fire and similar disturbances. Moreover, habitats are heterogeneous in space as well as time. A habitat patch will exhibit small scale variation in its physical and biological characteristics, and so (the thought goes) in different micropatches different species will be more efficient. This helps explain how communities can retain diversity (as competitive superiority will not produce a monoculture) and explains why more diverse communities are more productive. For they are more likely to include the species that are best suited to the various local micropatches spread through the habitat.

Tilman's crucial theoretical idea is that of compensation. If one population declines in numbers, another population, using somewhat similar resources, expands, and hence stabilizes the overall productivity of the community. Importantly, the idea that populations compensate for one another's fluctuations does not depend on controversial ecological assumptions. In particular, it does not depend on the idea that decline in one population is caused by the competitive superiority of the expanding species. We get community level stability despite population level volatility because individual populations have somewhat overlapping resource requirements but quite different environmental tolerances. Tolerance dif-

ferences explain why populations fluctuate out of synchrony. Resource overlaps explain how one population can expand as a result of another contracting (thus leaving resources unused). The overall effect is to partially stabilize the overall productivity of the community. Furthermore, there is empirical evidence that supports this cluster of ideas. Tilman's own empirical work concentrates on Minnesota grassland plots, though he also reports African data which leads to similar conclusions. These long term studies support the key biodiversity-stability hypotheses. Species rich plots resisted drought better; overall biomass varied less in species rich plots; and species rich plots returned to the predrought biomass more rapidly than species poor plots (Tilman 1996, 358). In summary, Tilman argued that both theoretical models and experimental findings broadly support diversity-stability hypotheses when these are taken to be hypotheses about communities rather than populations (Tilman 1996, 1999; Lehman and Tilman 2000).

So far, then, the case for the emergent property hypothesis looks promising. The diversity-stability hypothesis has empirical support, and the crucial mechanism, compensation, does not rest on controversial assumptions about strong interactions between populations. However, there are problems. The experimental evidence in favor of the diversity-stability relationship depends on measuring plant biomass. There are serious doubts about whether these ensemble relationships hold when we consider the interactions between plants and animals, and between animals. When our attention shifts to herbivores and those that eat them, resource exploitation efficiency is by no means an obvious stabilizing mechanism. To the contrary, enhanced resource use can cause over exploitation and hence productivity collapses (Loreau et al. 2001, 807).

In short, though diversity-stability and diversity-productivity hypotheses are plausible, it is not a settled fact of descriptive ecology that more diverse communities are more stable (or productive). That is one reason to be wary of the conclusion that communities have causally important ensemble properties and must therefore be real causal systems. There is a second reason for caution about this inference. Even if more diverse communities are more stable, it is not clear that they are more stable *because* they are more diverse. Diversity may be a symptom of causally relevant properties of individual populations rather than a causally important property of ensembles.

On the one hand, suppose that stability does depend on compensation effects. The overall productivity of the community depends on a set of key processes. These include the acquisition of energy by primary producers; the flow of minerals to and from the abiotic substrate; decomposition by the detritivores; and the flow of organic material from organisms to organisms via predation, herbivory and similar activities. Ecosystem

function depends on these key processes. Communities which are more diverse, and hence have a variety of species with rather different tolerances that can compensate for one another while driving these processes, are thereby more stable. If this is the right story, diversity itself is genuinely causally important (Naeem 1998).

However, there is another possibility: the “sampling effect.” Species rich communities have more tickets in the relevant biological lotteries. Thus a diverse community is more likely to have species that resist disturbance. They are more likely to have drought resistant members. Likewise, species rich Serengeti plots are more likely to contain grazing resistant members. More generally: to the extent that resistance to disturbance, productivity, resistance to invasion, and other community properties depend on the presence of *specific taxa* from the regional species pool, richer communities are more likely to manifest these properties just because they are more likely to include the right taxa (Wardle 1999). Tilman notes, “Diversity itself is both a measure of the chance of having certain species present in a system and a measure of the variation in species traits in an ecosystem” (1999, 1470). The second of these properties—variation in traits in the community—is a genuine property of the community itself. Thus if a rich community has *a set of species* that can use the available resource envelope with maximal efficiency and is thus more able to resist invasion by lowering the level of resources below the point at which potential invaders can survive, then the diversity of the community is causally crucial to its resistance to invasion. Not so if diversity just increases the chance that an already dominant species is present already, and it is the presence of *that key species* which explains resistance to invasion.

This debate is still open, and difficult to settle. Chesson et al. point out that though the sampling effect and the complementarity effect are empirically distinguishable in homogenous environments, they seem to be predictively equivalent in heterogeneous environments. In homogenous environments, according to the sampling effect, no mixed community can be more productive than the best single species community. But in a heterogeneous environment both models predict that specific patches will be dominated by the best species for that patch, and both models predict “over yielding”; i.e., they both predict that mixed communities are more productive than single species communities (Chesson et al. 2001, 238–239). Chesson and his colleagues seem to take this as bad news for a sampling hypothesis explanation of the results of their models, models which support diversity-productivity relationships. But of course indistinguishability arguments cut both ways. So to establish an emergent property hypothesis, two quite onerous conditions must be satisfied. The covariation between the emergent property and its apparent effect must

be robust, not limited to a few kinds of systems. And the relationship must be genuinely causal. Tilman's hypothesis and its relatives remains plausible, but as with internal regulation, these hypotheses are not vindicated by common sense biology.

4. Boundaries. In the introduction, I noted that many communities do not have sharp, obvious boundaries. But if they have no boundaries at all, it is hard to see how phenomenological communities could be causal systems. Levins and Lewontin think that communities are both real and bounded. But those boundaries are not defined by sharp changes in physical conditions but by changes in interaction patterns. On the Levins-Lewontin conception, communities are systems of strongly interacting populations, where "strong" and "weak" are understood comparatively: the members of a community interact strongly with one another by comparison to influences on and from populations outside the community (Levins and Lewontin 1985, 138). Given this perspective of what a community is—a more or less closed network of interacting populations—boundaries are defined by interaction, not proximity. A boundary of a community is a zone in which interactions become fewer and weaker; it is a zone in which biological events—local increases and dips in population—have less impact on the populations of community members (see also Odenbaugh, forthcoming).

Even so, systems of strongly interacting populations will tend to occupy an identifiable physical space. Suppose A, B and C are strongly interacting populations, with the effects of C on A mediated by its effects on B. In most cases, the strong interaction condition will imply that the territories of the three populations largely coincide. For if they do not—if, say, the A and B only intersect moderately—many members of the A population will never meet a B, and vice versa. Since most biological interaction requires contact, populations that merely overlap moderately do not usually interact strongly. Not all biological interactions require intimate contact, and to the extent that this rough generalization fails, communities defined by strong interaction will not be spatially well defined. Even so, as a rule of thumb, interaction requires proximity. Think of such characteristic ecological interactions as predation, herbivory, mutual exchange of nutrients, pollination. None of these are interactions at a distance (though as Mark Colyvan reminds me, interaction need not be continuous: predators might interact with prey strongly only at a waterhole). Communities of strongly interacting populations will be roughly spatially identifiable.

Interaction strength is a matter of degree. Most plausible examples of communities come with a penumbra of intermediate cases: populations that by the interaction criterion are neither clearly members of a particular

community of interacting populations, nor clearly not part of that community. The interaction criterion will not sharply define most communities. But that is no problem for the view that communities are often causal systems. As Levins and Lewontin point out, evolutionary systems often do not have sharply defined boundaries either. Consider speciation. There is the full spread from unequivocally separate lineages to uncontroversial examples of regional variants of a single species. The existence of intermediate cases is no reason to be skeptical about the idea that species themselves are causally important systems.

Likewise the existence of intermediate cases is no challenge by itself to the idea that communities are networks of populations whose trajectories are under mutual influence (for a dissenting view, see Parker 2004; for further support of Levins and Lewontin's unperturbed response to intermediate cases, see Callicott 1996). But the Levins and Lewontin view does seem to presuppose that patterns of interaction are patchy: that most populations are parts of networks whose members interact with one another more strongly than they interact with populations outside the network. It is far from obvious that this condition is typically met. It is quite likely that ecological interactions are not clumped in ways that enable us to identify bounded communities, even taking into account the fact that community boundaries are vague.

One problem with the Levins-Lewontin criterion is that the interaction patterns of different components of putative communities may well not coincide. Even if communities are networks of interacting populations, they are typically demographically open. Migrants move in and out of most habitat patches. As we saw in Section 2, such movements are likely to have stabilizing effects. We have two populations rather than one if organisms of the same type are related by metapopulation dynamics rather than competition. The echidnas in Black Mountain are part of a different population, and hence a different community from the echidnas on the O'Connor Ridge (about a kilometer to the north) if they are a source population for the O'Connor Ridge echidnas. They then buffer that group against population collapses rather than competing with them for scarce resources. *Prima facie*, though, there is not much reason to expect the dynamics of echidna populations to match those of larger and more mobile organisms, or those of smaller and less mobile ones. Black Mountain kangaroos may well compete directly for resources with O'Connor Ridge kangaroos. It is hard to tell just how serious this problem is, for there are ecological processes that can generate coordinated patchiness across a habitat. Organisms do not just passively experience their environment, they actively change it. Organisms in part construct their own niches (Odling-Smee et al. 2003). As Paul Griffiths has pointed out to me, this is one mechanism through which an initially fairly homogenous territory

can turn into a mosaic of quite different patches. Niche construction can magnify an initial difference between patches, beginning a cascade that takes us from initially similar systems to a mosaic of quite different patches. Suppose, for example, that eucalypts rather than acacias happen to predominate in one zone (perhaps initially for quite accidental reasons). Eucalypts and acacias have different environmental effects (eucalypts grow more slowly; they support very different pollinators; they produce very different litter). Hence an initial difference can generate quite marked biological differences between adjoining patches, thus generating two somewhat closed networks of interacting populations. The extent to which niche construction increases landscape scale heterogeneity, creating mosaic effects, is not known. But it is one mechanism that might induce rough congruence between populations.

Another possibility is that ecological boundaries are defined by emergent properties. We may be able to roughly identify boundaries with zones in which the value of the emergent property changes. The diversity-stability literature suggests, for example, that there is an important difference between ensembles that are resource limited and those that are disturbance dominated (see, for example, Brown et al. 2001). If so, the zone in which resource limits give way to increased disturbance marks the border between communities. Likewise, if some form of the diversity-stability hypothesis is right, the zone in which a species rich assemblage gives way to a species poor assemblage marks the border between communities. For it will also mark the transition from a more to a less stable local system.

These ideas are not part of the common core of ecology. Moreover, even if they do identify important ecological processes generating real patchiness, it is not likely that every phenomenological community will have determinate boundaries. So what should we say about community ecology if communities often do not have objective boundaries? That is the topic of the next section.

5. Indexical Communities. Community ecology can often be carried out without commitment to the extent to which phenomenological communities are causal systems. In 1994, a relict population of the Wollemi pine—an *Araucariaceae* in the conifer family—was discovered in Wollemi National Park. There is an obvious conservation interest in the fate of this population, yet it may be part of no natural community. It may well be that these pines interact with a belt of associated plant, animal, fungus and microbial species. Likely enough, they in turn interact with others which in turn have links to still further populations. There is no special likelihood that the Wollemi pines are part of a bounded (even vaguely bounded) network of interacting populations. When we consider large

tracts of fairly similar habitat (such as the Wollemi National Park), there is no reason to suppose that ecological interaction is clumped into semi-independent nets of populations.

Even so, if these trees interact with a fairly stable set of organisms that are important to their demographic fate, the Wollemi pine population may well be part of an *indexical community*. The Wollemi pine indexical community consists in the set of populations with whom these pines interact directly, together with the other local populations which have important effects on those populations that effect the Wollemis. For example, the Wollemi indexical community will include those that attack its seeds. It may include the predators of those seed eaters, if those second hand interactions are strong enough to effect the flow of Wollemi seed. Given a choice of focal population and a decision about how fine grained our account of the fate of that population is to be, the extent of this indexical community is fixed. Identifying indexical communities enables us to pose some of the traditional problems of community ecology while remaining neutral on the extent to which the relevant phenomenological community is a causal system. But our interest in these communities is limited. Taking this stance is helpful only for predicting the fate of this focal population. We would have to redefine the community if our interest was in the predators of Wollemi seedeaters. Nor would we have good reason to accept (say) a stability-diversity hypothesis about such a indexical community. It makes no sense to apply to indexical communities our theory of the relationship between species richness and invasion resistance. For the Wollemi indexical community might be species poor, not because there were few species in the relevant habitat patch, but because only a few species have associations with Wollemi that are both stable enough and intense enough to affect its population dynamics. The criteria of inclusion in the Wollemi indexical community is: "Does this population significantly affect Wollemi population movements?" not "Does this population impact the local ecology as a whole?" There may be many species lurking around the Wollemi that make a big difference to the effects of fire, drought or invasion but which do not make much difference to the Wollemi. These pines live many hundreds of years and can afford to take a relaxed view of much seasonal variation. The Wollemi may live in spatial association with many species that are ecologically consequential but not part of their indexical community.

In short, we can ask questions about the distribution and abundance of particular populations in the context of their indexical community. But other elements of the research agenda of community ecology make no sense in this context.

6. The Space of Phenomenological Communities. There is no unequivocal

inference from the stability of communities to their having internal organization; nor is there unequivocal demonstration that they have ecologically salient emergent properties. Moreover, in many cases we cannot establish that they have objective bounds. Even so, recent theoretical developments have strengthened the case for thinking that communities are structured, and have ensemble properties, by recognizing the fact that organisms are profound agents of transformation both of their own and others' environments. Levins and Lewontin made this point forcefully, and their ideas have been developed by those working on niche construction (Levins and Lewontin 1985; Jones et al. 1997; Odling-Smee et al. 2003). Populations within a community can be linked via niche construction networks. One population can influence another by changing important features of the physical environment. Trees buffer the wind and modulate the impact of storms whilst providing shelter to many organisms. They stabilize the soil and influence water tables. These indirect ecological links expand the range of potential interactions in communities. Litter recycling is the cleanest example. Plants produce litter as a by-product of their life: fallen leaves, twigs, bark. A host of organisms live by consuming the litter, and as a consequence of these actions, they return crucial materials to the soil. These materials are absorbed by the vegetation, which in turn produces more litter. Thus "balance of nature" arguments understate the case for internal organization by focusing so exclusively on negative, direct biological interactions, of which competition and predation are the prime examples. Moreover, niche construction often involves ensemble effects. Soils are made organically, but not by any single population. They are made by a vast suite of very different animals, plants and fungi: ants and other burrowing animals turn over and redistribute soils; trees and other plants stabilize it; fungi, microbes and a vast army of small invertebrates make it by consuming litter. Soils are made, and not by any one population. The individualist view of communities looks plausible only when niche construction is neglected.

However, the crucial factors that distinguish a merely phenomenological community from a causally salient system come in degrees. Thus an assemblage may be internally regulated *to some extent*. The importance of internal regulation will depend on: (a) connectance, i.e., the proportion of the component populations that interact strongly enough to influence abundance in the community; (b) the mean strength of those interactions; (c) the extent to which regulating interactions are disrupted by outside disturbance. Likewise, an assemblage may have causally salient emergent properties to some extent. Suppose, for example, that there really is a diversity-stability relationship. If so, it will be a gradient phenomenon: the causal importance of the diversity-stability relationship will depend on the degree to which diversity buffers the community against distur-

bance, the range of properties which are buffered against disturbance, and the kind of disturbances whose effects are buffered. Boundedness, too, comes in degrees: a network of interacting populations can be more or less closed, more or less spatially coincident rather than merely intersecting.

Thus the factors that make phenomenologically identifiable communities explanatorily salient come in degrees. Moreover, these factors may well be at least partially independent of one another. If, for example, the stabilizing effects of diversity depends on compensation effects between populations, an assemblage can have causally important emergent properties without being internally regulated. On a true habitat island, local populations will occupy the same range, and hence be spatially coincident. Such an assemblage will be bounded, but the populations need not interact strongly. Nor need it have important ensemble effects. So a community can be determinately bounded but without much explanatory salience. In view of both the gradient and the (at least partial) independence of these factors, we should think of communities as occupying differing positions in a three-dimensional space. The three dimensions will be boundedness, internal regulation, and emergent property effects. An individualist assemblage in one corner of the space would consist of a set of populations that merely overlap, which do not significantly influence one another's demographic prospects, and which have no important collective impact on their environment. In the opposite corner of the space, we would find assemblages of spatially coincident populations strongly influencing one another's demographic fates, and with important ensemble effects.

This way of representing the nature of communities suggests an important research agenda. One important set of questions concerns the distribution of actual assemblages in Assemblage Space. There is no *a priori* reason to expect most actual assemblages to be in the same location in this space. But repeated patterns would be very interesting: are (say) all desert communities in roughly the same location? All open woodlands? Ecologists have already raised somewhat similar questions, in looking for correlations between physical characteristics of habitats and community level properties: for example, relations between resource availability and community richness. Thus they have asked such questions as "Are desert communities typically species poor?" (see Mikkelsen 2003 for a discussion of patterns like these). Assemblage space enables us to sharpen hypotheses about these correlations. Do desert communities tend to have similar causal profiles? Species richness would be just one symptom of such a shared profile. This way of representing communities also encourages us to ask important questions about the dynamics of actual assemblages. Some paleoecologists think that they have identified a "Pleistocene Paradox": ecosystems since the Pleistocene are less stable, less integrated than

earlier ecosystems (Boucot 1990). Similar issues arise at smaller temporal and spatial scales. For example, the reassembly of the Krakatau volcano communities obviously involves profound changes over time in the taxonomic composition of the reassembling communities (Thornton 1996). But does it involve changes in the extent to which they are causally salient systems as well? Hypotheses of this general form are very important: we want to know whether integration, ensemble effects, and boundedness change over time, and if so, what factors explain those changes.

There are also important open questions about the dimensions of Assemblage Space itself. It is possible that other dimensions are important. Ricklefs has recently argued that community ecology has been dominated by the assumption that local communities are under local control. He thinks this assumption of “local determinism” is and should be coming under increasing challenge (Ricklefs 2004). Island biogeography (and its more nuanced descendants) suggests that openness contributes very importantly both to richness and to stability (Ricklefs and Schuler 1993; Ricklefs 2004). So a candidate fourth dimension is openness to influence from the regional species pool.

In brief, then, this paper has both a substantive and a methodological conclusion. Substantively: it is possible but unlikely that most actual assemblages are clustered in the individualist corner of Assemblage Space. Methodologically: is likely to be productive to reconceptualize questions about communities as questions about location in Assemblage Space.

REFERENCES

- Boucot, Arthur J. (1990), *Evolutionary Paleobiology of Behavior and Coevolution*. Amsterdam: Elsevier.
- Brown, James H., Morgan Ernest, Jennifer Parody, and John Haskell (2001), “Regulation of Diversity: Maintenance of Species Richness in Changing Environments”, *Oecologia* 126: 321–332.
- Callicott, J. Baird (1996), “Do Deconstructive Ecology and Sociobiology Undermine the Leopold Land Ethic?”, *Environmental Ethics* 18: 353–372.
- Chesson, Peter, Stephen Pacala, and Claudia Neuhauser (2001), “Environmental Niches and Ecosystem Functioning”, in Andrew Kinzig, Stephen Pacala, and David Tilman (eds.), *The Functional Consequences of Biodiversity*. Princeton, NJ: Princeton University Press, 213–245.
- Cooper, Greg (1993), “The Competition Controversy in Community Ecology”, *Biology and Philosophy* 8: 359–384.
- (2001), “Must There Be a Balance of Nature?”, *Biology and Philosophy* 16: 481–506.
- (2003), *The Science of the Struggle for Existence*. Cambridge: Cambridge University Press.
- Cummins, Robert (1973), “Functional Analysis”, *Journal of Philosophy* 72: 741–764.
- Eldredge, Niles (2003), “The Sloshing Bucket: How the Physical Realm Controls Evolution”, in James P. Crutchfield and Peter Schuster (eds.), *Evolutionary Dynamics: Exploring the Interplay of Selection, Accident, Neutrality and Function*. Oxford: Oxford University Press.

- Godfrey-Smith, Peter (1993), "Functions: Consensus without Unity", *Pacific Philosophical Quarterly* 74: 196–208.
- (1994), "A Modern History Theory of Functions", *Nous* 28: 344–362.
- Hixon, Mark, Stephen Pacala, and Stuart Sandin (2002), "Population Regulation: Historical Context and Contemporary Challenges of Open vs. Closed Systems", *Ecology* 83: 1490–1508.
- Jones, Chris, John Lawton, and Moshe Shachak (1997), "Positive and Negative Effects of Organisms as Physical Ecosystems Engineers", *Ecology* 78: 1946–1957.
- Lehman, Clarence, and David Tilman (2000), "Biodiversity, Stability and Productivity in Competitive Communities", *American Naturalist* 156: 534–552.
- Levins, Richard, and Richard C. Lewontin (1985), "Dialectics and Reduction in Ecology", in Richard Levins and Richard C. Lewontin (eds.), *The Dialectical Biologist*. Cambridge, MA: Harvard University Press, 132–160.
- Loreau, L., et al. (2001), "Biodiversity and Ecosystem Functioning: Current Knowledge and Future Challenges", *Science* 294: 804–808.
- May, Robert M. (1973), *Stability and Complexity in Model Ecosystems*. Princeton, NJ: Princeton University Press.
- Mikkelsen, Greg M. (2003), "Ecological Kinds and Ecological Laws", *Philosophy of Science* 70: 1390–1400.
- Murdoch, William W. (1994), "Population Regulation in Theory and Practice", *Ecology* 75: 271–287.
- Naeem, Shahid (1998), "Species Redundancy and Ecosystem Reliability", *Conservation Biology* 12: 39–45.
- Odenbaugh, Jay (forthcoming), "Ecology", in Sahotra Sarkar (ed.), *Encyclopedia of the Philosophy of Science*. London: Routledge.
- Odling-Smee, John, Kevin Laland, and Marc Feldman (2003), *Niche Construction: The Neglected Process in Evolution*. Princeton, NJ: Princeton University Press.
- Paine, Robert T. (1966), "Food Web Complexity and Species Diversity", *American Naturalist* 100: 65–75.
- Parker, V. Thomas (2004), "The Community of an Individual: Implications for the Community Concept", *Oikos* 104: 27–34.
- Pimm, Stuart (1991), *The Balance of Nature*. Chicago: University of Chicago Press.
- Ricklefs, Robert E. (2004), "A Comprehensive Framework for Global Patterns in Biodiversity", *Ecology Letters* 7: 1–15.
- Ricklefs, Robert E., and Dolf Schuller (1993), "Species Diversity: Regional and Historical Influences", in Robert E. Ricklefs and Dolf Schuller (eds.), *Species Diversity in Ecological Communities*. Chicago: University of Chicago Press, 350–363.
- Thornton, Ian (1996), *Krakatau: The Destruction and Reassembly of an Island Ecosystem*. Cambridge, MA: Harvard University Press.
- Tilman, David (1996), "Biodiversity: Population versus Ecosystem Stability", *Ecology* 77: 350–363.
- (1999), "The Ecological Consequences of Changes in Biodiversity: A Search for General Principles", *Ecology* 80: 1455–1474.
- Turchin, Peter (1999), "Population Regulation: A Synthetic View", *Oikos* 84: 153–159.
- Wardle, David (1999), "Is 'Sampling Effect' a Problem for Experiments Investigating Biodiversity-Ecosystem Function Relationships?", *Oikos* 87: 403–407.

Copyright of Philosophy of Science is the property of Philosophy of Science Association and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.