

Draft copy. Please do not cite without permissions. Some rights reserved under a Creative Commons Attribution/Share-Alike license by Jon Lawhead, Columbia University Department of Philosophy.

0. Introduction

The pluralism of scientific models is both obvious and challenging. The proliferation of many different models of the same phenomenon in (e.g.) biology or climate science is so clear as to be impossible to miss, but it also presents interesting challenges for both working scientists and philosophers of science. As science has increasingly moved toward directly grappling with complex systems like the global climate, this working pluralism has become more and more difficult to ignore; sciences that deal with complex systems seem to be more rife with pluralism than ever. Given this evolving feature of scientific methodology, we might do well to reexamine some of the philosophical characterizations of science itself in light of contemporary advances.

In this paper, we shall examine some aspects of complexity theory's impact on the debate about the nature of scientific laws and science generally. I will argue that the notion of integrative pluralism-- introduced in Mitchell (2002), but most comprehensively articulated in Mitchell (2009)--can provide the skeleton of a robust account of science and lawhood that is better suited to deal with the kind of messy multi-level phenomena of interest to scientists studying complex systems. The ontology of "real patterns" introduced in Dennett (1991)--and leveraged to great effect by the defenders of structural realism community fits nicely with the sort multi-level account of scientific modeling and explanation that's necessary to generalize Mitchell's work into a fully-fledged theory of scientific work.

In **Section 1**, I shall sketch the background for IP's entrance into the debate on model

pluralism, with a particular emphasis on this debate's relevance to contemporary climate science. In **Section 2**, I'll draw on features of the "ontology of patterns" developed by structural realists to offer a concrete account of scientific laws which is distinctively suited to working with complex systems. While this section includes elements drawn from a metaphysical theory, it is primarily pragmatic in character: I have no real stake in the "real" element of the "real patterns" debate, except insofar as it is useful in thinking about scientific practice. In **Section 3**, I'll examine the relationship between this account of science and integrative pluralism as it has been articulated so far, and argue that it provides the basis for generalizing Mitchell's insights into a more comprehensive account of scientific reasoning in practice. Finally, we shall see how the foundational work done here helps us understand the contemporary pluralism in climate modeling.

1. Competitive and Compatible Pluralisms

Integrative pluralism (IP) constitutes fairly broad-reaching research program within the philosophy of science. Mitchell (2002) introduces IP as an attempt to bridge the gap between two positions that she has called "competitive pluralism" and "compatible pluralism" in Mitchell (1992). Both competitive and compatible pluralism are attempts to deal with the presence of a significant number of different theories that's often present within single scientific domains. For a striking contemporary example, consider the truly staggering number of different climate models in use today. Contemporary textbooks on climate modeling identify dozens of in-use models, even in the same *domain*. Houghton et. al. (2001) lists more than 30 widely used general circulation models (GCMs), each of which is supposed to model the global climate as a

single integrated system of oceanic and atmospheric circulation, but many of which operate with wildly different assumptions about the structure and dynamics of the system they are modeling--and that's just within GCMs, which are only single "class" of climate models. Earth Models of Intermediate Complexity (EMICs) tend to be more specialized than GCMs--they are designed, for instance, to specifically track the sea ice extent in the Antarctic, or to model changes in ocean salinity in response to increasing water temperature. EMICs are even more heterogenous than GCMs, and while there may be dozens of independent GCMs, there are literally *hundreds* of different EMICs in use by climate scientists around the world.

Given this proliferation of different models in sciences like climatology, how are we supposed to think about inevitable clashes of model assumptions? When one model represents the world as being *this* way while another represents the world as being *that* way (and we repeat that process a few dozen times), how are we to resolve the mounting tension? Concerns about whether we should *believe* the predictions generated by climate models (and if we should, how seriously we should take them) are partially underwritten by this pluralism (see Knutti (2008)); if the models don't agree on the nature of the system they're supposed to be representing, then why should we trust what they have to say about that system's behavior?

Discussions of how to understand this pluralistic approach to modeling the climate have begun to garner some attention (Parker (2006) for example), but the larger question of pluralism within science *in general* is worth confronting as well. Mitchell (1992) distinguishes between competitive and compatible pluralism as the two standard accounts of how we should think about scientific pluralism generally. Competitive pluralism, which she attributes to (among others) Kitcher (1991) and Beatty (1987) seems to be the mainstream view (insofar as there is one)

within philosophy of science. On this account, the multiple competing models (or theories, if you prefer) within particular domains are if not outright *pernicious*, something that we ought to strive to do away with in the limit. This is likely a consequence of the fact that philosophy of science has, as Mitchell notes, tended to focus on historical cases where such competitive elimination *was* in fact the outcome: Lamarckian evolution versus Darwinian evolution, aether theory of light propagation versus special relativity, the caloric theory of heat versus the kinetic theory, and the alchemical theory of composition versus atomic theory are all cases in which two or more competing models of one or another phenomenon arose and (after the dust had cleared) only one was left standing. We might think of this as the Thunderdome approach to model-building: two models enter, one model leaves.

In contrast, compatible pluralism has had relatively few active defenders. Mitchell argues that modern compatible pluralism has reached its apotheosis in the “levels of analysis” stance of Sherman (1988). Sherman argues that different models (or explanations) need not be thought of as competing unless they operate at the same “level of analysis,” where “level” seems to mean something like it did for Oppenheim and Putnam: a relatively distinct spatial scale of a system, defined in terms of its parts. On this view, we shouldn’t worry about the apparent contradiction between (say) brain-state level explanations for my behavior and chemical (or atomic) level explanations for that same behavior. There’s no contradiction, in other words, in giving an account of what’s going on in my head in terms of neurons firing and giving another account in terms of chemicals bonding and degrading. Chemistry and cognitive neuroscience can be thought of as operating on different *levels*, and thus aren’t direct competitors with one another despite dealing with what is really the same physical system in very different ways. They can

coexist.

Mitchell (2002) and later (2004, 2009) argues convincingly that this account of compatible pluralism based on levels of analysis is deficient. Her discussion focuses on explaining the behavior and structure of social insects like ants and bees, but it is worth noting that cases like climate science may provide an even more vivid reason to be concerned. As we noted, the plurality of explanations and models in climate science is not just a matter of what Sherman would recognize as inter-level pluralism, but rather also one of *intra*-level pluralism. We might be able to defuse the worry about potential competition between EMICs and GCMs (or perhaps even between different EMICs) through reference to Sherman-style compatible pluralism, but surely the same strategy is not available for smoothing over the apparent conflicts between different GCMs, all of which explicitly seek to operate on the same level of analysis.

Perhaps we should, then, endorse competitive pluralism (albeit maybe of the rather mild variety espoused by Kitcher (1991)), recognizing the climate model zoo as a necessary evil for the time being, but cast our eyes toward eventually eliminating it in favor of One Model to Rule Them All. This might be the case even if we want to allow for inter-level pluralism: we might want to shoot for one GCM, one sea ice extent model, one radiative balance model, and so on. But it seems to me that the fact that climate scientists *don't* appear to be thinking this way should give us pause: if climate scientists themselves are comfortable with a pluralism of models, why should philosophers be given the authority to countermand them? It is, at the very least, worth seeing if we can construct an account of scientific laws and modeling that respects working climatologists' comfort with persistent model pluralism. Let's give it a shot.

2. Toward a Pattern-Based Science

Let's start at the beginning with a rather banal observation: science is about the world.

Scientists are in the business of understanding the world around us—the *actual* world, not the set of all possible worlds, or Platonic heaven, or J.R.R Tolkien's Middle Earth. Of course, this isn't just limited to the *observable*, or *visible* world: science is interested in the nature of parts of the world that have never been directly observed and (in at least some cases) never will be.

Physicists, for instance, are equally concerned that their generalizations apply to the region of the world *inside the sun* as they are that those generalizations apply to their laboratory apparatuses.

There's a more important sense in which science is concerned with more than just the observed world, though: science is not just descriptive, but *predictive* too—good science ought to be able to make predictions, not just tell us the way the world is *right now* (or was in the past). A science that consisted of enumerating all the facts about the world *now*, as useful as it might be, wouldn't seem to count as a full-fledged science by today's standard, nor would it seem to follow the tradition of historical science; successful or not, scientists since Aristotle (at least!) have, it seems, *tried* to describe the world not just as it is, but *as it will be*.

This leads us to another (perhaps) banal observation: science is about predicting how the world changes over time. Indeed, a large part of how we *judge* the success (or failure) of scientific theories is through their predictive success; the stock example of Fresnel's success with the wave theory of light, as demonstrated by the prediction (and subsequent observation) of a bright spot at the center of the shadow cast by a round disk is a stock example for good reason—it was a triumph of novel predictive utility. General relativity's successful prediction of the actual orbit of the planet Mercury is another excellent paradigm case here; Mercury's erratic

orbit, which was anomalous in Newton's theory of gravity, is predicted by Einstein's geometric theory. This success, it is important to note, is not in any sense a result of "building the orbit in by hand;" as Ladyman and Ross (2007) observe, though Einstein did (in some sense) set out to *explain* Mercury's orbit through a general theory of gravitation, he did this entirely by reference to *general* facts about the world—the empirically accurate prediction of Mercury's orbit followed from his theory, but nothing in the theory itself was set with that *particular* goal in mind. The history of science is, if not exactly littered with, certainly not lacking in other examples of success like this; indeed, having surprising, novel, *accurate* predictions "pop out" of a particular theory is one of the best markers of that theory's success¹.

It is not enough, then, to say that science is about prediction of how the world will change over time. Science doesn't just seek to make *any* predictions, it seeks to make predictions of a particular sort—predictions with verifiable consequences—and it does this by attempting to pick out patterns that are in evidence in the world *now*, and projecting them toward the future. That is to say: *science is the business of identifying genuine patterns in how the world changes over time*. It is precisely this projectability that makes a putative pattern *genuine* rather than ersatz; this is why science is of necessity concerned with more than just enumerating the facts about the way the world is now—just given the current state of the world, we could hypothesize a virtually

¹ The Aharnov-Bohm effect, a surprising quantum mechanical phenomenon in which the trajectory of a charged particle is affected by a local magnetic field even when traversing a region of space where both the magnetic field and the electric fields' magnitudes are zero, is another excellent example here. This particular flavor of non-locality implies that the classical Maxwellian formulation of the electromagnetic force as a function of a purely local electrical field and a purely local magnetic field is incomplete. The effect was predicted by the Schrodinger equation years before it was observed, and led to the redefinition of electromagnetism as a gauge theory featuring electromagnetic *potentials*, in addition to fields. See Aharnov and Bohm (1959). Thanks to Porter Williams for suggesting this case.

infinite number of “patterns” in that state, but only *some* of those putative patterns will let us make accurate predictions about what the state of the world will be in (say) another hour.

2.1 Toy Science and Basic Patterns

Suppose we’re given a piece of a binary sequence, and asked to make predictions about what numbers might lie outside the scope of the piece we’ve been given:

S1: 110001010110001

Is there a genuine pattern in evidence here? Perhaps. We might reasonably suppose that the pattern is “two ‘ones,’ followed by three ‘zeros’ followed by ‘one, zero, one, zero,’ and then repeat from the beginning.” This putative pattern *R* is empirically adequate as a theory of how *this* sequence of numbers behaves; it fits all the data we have been given. How do we know if this is indeed a genuine pattern, though? Well, we can continue to watch how the sequence of numbers behaves, and see if our predictions bear out. If we’ve succeeded in identifying the pattern underlying the generation of these numbers, then we’ll be able to predict what we should see next: we should see a ‘zero’ followed by a ‘one,’ and then another ‘zero,’ and so on. Suppose the pattern continues:

S2: 0101100010101

Ah ha! In *S2*, the string of numbers continues to evolve in a way that is consistent with our hypothesis that the sequence at large is (1) not random and (2) is being generated by the pattern *R*. Of course, this is not enough for us to say *with certainty* that *R* (and only *R*) is the pattern behind the generation of our sequence; it is entirely possible that the next few bits of the string will be inconsistent with *R*; that is one way that we might come to think that our theory of how the string is being generated is in need of revision. Is this the only way, though? Certainly not:

we might also try to obtain information about what numbers came *before* our initial data-set and see if R holds there, too; if we really have identified the pattern underlying the generation of S , it seems reasonable to suppose that we ought to be able to “retrodict” the structure of subsets of S that come *before* our initial data-set just as well as we can predict the structure of subsets of S that come *after* our initial data-set. Suppose that we find that just before our initial set comes the string:

S_0 : 00001000011111

The numbers in *this* string are not consistent with our hypothesis that all the numbers in the sequence at large are generated by R . Does this mean that we’ve failed in our goal of identifying a pattern, though? Let’s think it through.

There’s an important question that we’ve been glossing over in our discussion here: for a pattern in some data to be *genuine* must it also be *global* (in the computer scientist’s sense of holding over the entirety of the data set in question)? That is, for us to say reasonably that R describes the sequence S , must R describe the sequence S *everywhere*? Here’s all the data we have now:

S_{0-2} : 000010000111111100010101100010101100010101

It is clear that we can no longer say that R (or indeed any single pattern at all) is the pattern generating all of S . This is not at all the same thing as saying that we have failed to identify a pattern in S *simpliciter*, though. Suppose that we have some reason to be particularly interested in what’s going on in a restricted *region* of S : the region S_{1-2} . If that’s the case, then the fact that R turns out not to hold for the totality of S might not trouble us at all; identifying a *universal* pattern would be *sufficient* for predicting what sequence of numbers will show up in S_{1-2} , but it is

by no means necessary. If all we're interested in is predicting the sequence in a particular region of S , identifying a pattern that holds *only* in that region is no failure at all, but rather precisely what we set out to do to begin with! It need not trouble us that the pattern we've identified doesn't hold *everywhere* in S —identifying that pattern (if indeed there is one to be identified) is another project entirely.

When we're investigating a sequence like S , then, our project is twofold: we first pick a *region* of S about which we want to make predictions, and then attempt to identify a pattern that will let us make those predictions. When we have a candidate pattern, we can apply it to heretofore unobserved segments of our target region and see if the predictions we've made by using the pattern are born out. That is: we first identify a particular way of *carving up* our target data-set and then (given that carving) see what patterns can be picked out. That any patterns identified by this method will hold (or, better, that we have *good reason* to think they'll hold) in a particular region *only* is (to borrow the language of computer programmers) a feature rather than a bug. It's no criticism to say that a putative pattern that we've identified relative to a particular carving of our subject-matter holds only for that carving; if our goal is just to make predictions about a restricted region of S , then identifying a pattern that holds only in that region might well make our jobs far easier, for it will give us license to (sensibly) ignore data from outside our restricted region, which might well make our task significantly easier[8].

Let's think about another potentially problematic case. Suppose now that we're given yet another piece of S :

S_3 : 001010**0**100010

S_3 is *almost* consistent with having been generated by R —only a single digit is off (the bolded

zero ought to be a one if R is to hold)—but still, it seems clear that it is not an instance of the pattern. Still, does this mean that we have failed to identify any useful regularities in S_3 ? I don't think it does. What's the difference between S_3 and S_0 such that we can say meaningfully that, in picking out R , we've identified something important about the former but not the latter?

Following Dennett (1991) and Ladyman and Ross (2007), we might begin by thinking of patterns as being (at the very least) the kinds of things that are "candidates for pattern recognition." But what does *that* mean? Surely we don't want to tie the notion of a pattern to *particular* observers—whether or not a pattern is in evidence in some dataset (say S_3) shouldn't depend on how dull or clever the person looking at the dataset is. As Dennett notes, though, there is a standard way of making these considerations more precise: we can appeal to information theoretic notions of compressibility. A pattern exists in some data if and only if there is some algorithm by which the data can be significantly compressed.

This is a bit better, but still somewhat imprecise. What counts as compression? More urgently, what counts as *significant* compression? Why should we tie our definition of a pattern to those notions? Let's think through these questions using the examples we've been looking at for the last few pages. Think, to begin with, of the sequence :

S_1 -2: 1100010101100010101100010101

This, recall, was our perfect case for R : the pattern we identified holds perfectly in this data-set. What does it mean to say that R holds perfectly in light of the Dennettian compressibility constraint introduced above, though? Suppose that we wanted to communicate this string of digits to someone else. One way—the easiest way, in a sense—to go about it would just be to transmit the string verbatim: to communicate a perfect *bit map* of the data. That is, for each digit

in the string, we can specify whether it is a 'one' or a 'zero,' and then transmit that information (since there are 28 digits in the dataset *SI-2*, the bit-map of *SI-2* is 28 bits long). If the string we're dealing with is truly random then this is (in fact) the *only* way to transmit its contents²: we have to record the state of each bit individually, because (if the string is random) there is no relationship at all between a given bit and the bits around it. Part of what it means to have identified a pattern in some data-set, then, is to have (correctly) noticed that there is a *relationship* between different parts of the data-set under consideration—a relationship that can be exploited to create a more efficient encoding than the simple verbatim bit-map.

The sense of 'efficiency' here is a rather intuitive one: an encoding is more efficient just in case it is *shorter* than the verbatim bit map—just in case it requires fewer bits to transmit the same information. In the case of *SI-2*, it's pretty easy to see what this sort of encoding would look like—we specify *R*, then specify that the string we're passing consists in two iterations of *R*. Given a suitable way of encoding things, this will be much shorter than the verbatim bit map.

This compressibility criterion is offered by Dennett as a necessary condition on patternhood: to be an instance of a (real) pattern, a data-set must admit of a more compact description than the bitmap. However, as Collier (1999) and Ladyman and Ross (2007) both point out, this cannot be the whole story; while compressibility is surely a necessary condition on patternhood, it cannot be both necessary *and* sufficient, at least not if it is to help us do useful work in talking about the world (recall that the ultimate point of this discussion is to articulate what exactly it is that science is doing so that we can see if philosophy has something useful to contribute to the

² Citing Chaitin (1975), Dennett (op. cit.) points out that we might actually take this to be the formal definition of a random sequence: there is no way to encode the information that results in a sequence that is shorter than the "verbatim" bit map.

project). Science cannot simply be in the business of finding ways to compress data sets; if that were so, then every new algorithm—every new way of describing something—would count as a new scientific discovery. This is manifestly not the case; whatever it is that scientists are doing, it is not *just* a matter of inventing algorithm after algorithm. There's something *distinctive* about the kinds of patterns that science is after, and about the algorithms that science comes up with. In fact, we've already identified what it is: we've just almost lost sight of it as we've descended into a more technical discussion—science tries to identify patterns that hold not just in existing data, but in unobserved cases (including future and past cases) as well. Science tries to identify patterns that are *projectable*.

Suppose we transmit the bitmap of *SI-2* and our recipient receives the following sequence:

SI-2: 1100010101100010101100??0?01

Some of the bits have been lost in transmission, and now appear as question marks—our interlocutor just isn't sure if he's received a one or a zero in those places. How can he correct for this? Well, suppose that he also knows that *SI-2* was generated by *R*. That is, suppose that we've *also* transmitted our compressed version of *SI-2*. If that's the case, then our interlocutor can, by following along with *R*, reconstruct the missing data and fill in the gaps in his signal. This, of course, requires more transmission overall—we have to transmit the bitmap *and* the pattern-encoding—but in some cases, this might well be worth the cost (for instance, in cases where there is a tremendous amount of latency between signal transmission and signal reception, so asking to have specific digits repeated is prohibitively difficult). This is in fact very close to how the Transmission-Control Protocol (TCP) works to ensure that the vast amount of data being

pushed from computer to computer over the Internet reaches its destination intact.

Ok, but how does this bear on our problem? Next, consider the blanks in the information our interlocutor receives not as *errors* or miscommunication, but simply as *unobserved cases*. What our interlocutor has, in this case, is a partial record of $S1-2$; just as before, he's missing some of the bits, but rather than resulting from an error in communication, this time we can attribute the information deficit to the fact that he simply hasn't yet *looked* at the missing cases. Again, we can construct a similar solution—if he knows R , then just by looking at the bits he *does* have, then our interlocutor can make a reasonable guess as to what the values of his unobserved bits might be. If an observer is clever, then, he can use a series of measurements on part of his data-set to ground a guess about a pattern that holds in that data set, and then use that pattern to ground a guess about the values of unmeasured parts of the data set.

At last, then, we're in a position to say what it is that separates $S3$ from $S0$ such that it is reasonable for us to say that R is informative in the former case but not in the latter, despite the fact that neither string is consistent with the hypothesis that R is the pattern underlying its generation. The intuitive way to put the point is to say that R holds *approximately* in the case of $S3$ but not in the case of $S0$, but we can do better than that now: given R , and a restricted set of $S3$, an observer who is asked to guess the value of some *other* part of the set will do far better than we'd expect him to if R was totally uninformative—that is, he will be able to make predictions about $S3$ which, more often than not, turn out to be good ones. In virtue of knowing R , and by measuring the values in one sub-set of $S3$, he can make highly successful predictions about how other value measurements in the set will turn out. The fact that he will also get things *wrong* occasionally should not be too troubling; while he'd certainly want to work to identify the

exceptions to R —the places in the sequence where R doesn't hold—just picking out R goes a very long way toward sustained predictive success. Contrast that case to the case in $S0$: here, knowledge of R won't help an observer make any deductions about values of unobserved bits. He can learn as much as he wants to about the values of bits before and after a missing bit and he won't be any closer at all to being able to make an educated guess about the missing data.

2.2 A New Unity of Science

Let's make this more concrete, and see how it might apply to scientific laws. Scientists are in the business of studying patterns in how the world changes over time. The method for identifying patterns varies from branch to branch of science; the special sciences differ in domain both from each other and from fundamental physics. In all cases, though, scientists proceed by making measurements of certain parts of the world, trying to identify patterns underlying those measurements, and then using those patterns to try to predict how unobserved cases—either future measurements or measurements in a novel spatial location—might turn out. Occasionally, they get a chance to compare those predictions to observed data directly. This is more common in some branches of science than in others: it is far more difficult to verify some of the predictions of evolutionary biology (say, speciation events) by observation than it is to verify some of the predictions of quantum mechanics (say, what state our measurement devices will end up in after a Stern-Gerlach experiment). More frequently, they are able to identify a number of different patterns whose predictions seem either agree or disagree with one another.

Just as in the case of our toy science in **Section 2.1**, it seems to me that science *generally* consists in two separate (but related) tasks: scientists identify a domain of inquiry by picking out

a way of carving up the world, and then identify the patterns that obtain given that way of carving things up. This is where the lengthy discussion from **Section 2.1** should be illuminating: not all scientists are interested in identifying patterns that obtain *everywhere* in the universe—that is, not all scientists are interested in identifying patterns that obtain for *all* of *S*. Indeed, this is precisely the sense in which fundamental physics is *fundamental*: it alone among the sciences is concerned with identifying the patterns that will obtain no matter where in the world we choose to take our measurements. The patterns that fundamental physics seeks to identify are patterns that will let us predict the behavior of absolutely any sub-set of the world—no matter how large, small, or oddly disjunctive—at which we choose to look; it strives to identify patterns that describe the behavior of tiny regions of space-time in distant galaxies, the behavior of the interior of the sun, and the behavior of the Queen of England's left foot. The special sciences are all, to one degree or another, concerned with identifying patterns that hold only in sub-sets of the domain studied by physics. A similar view of scientific laws is given in Maudlin (2007). Maudlin argues that scientific laws are best understood as what he calls LOTES—"laws of temporal evolution." This is largely consistent with the picture I have been arguing for here, and (not coincidentally) Maudlin agrees that an analysis of scientific laws should "take actual scientific practice as its starting point" (*ibid*, p. 10), rather than beginning with an *a priori* conception of the form that a law *must* take. Our point of departure from Maudlin's view, as we shall see, lies in our treatment of fundamental physics. While Maudlin wants to distinguish "FLOTES" (*fundamental* laws of temporal evolution) from normal LOTES on the basis of some claim of "ontological primacy" (*ibid*, p. 13) for fundamental physics, the view I am sketching here requires no such militantly reductionist metaphysics. The special

sciences are, on this view, emphatically *not* second-class citizens—they are just as legitimate as fields of inquiry as is fundamental physics. The sense of “fundamental” in “fundamental physics” should not be taken to connote anything like ontological primacy or a metaphysically privileged position (whatever that might mean) within the general scientific project. Rather (to reiterate) it is just an indicator of the fact that fundamental physics is the most *general* part of the scientific project; it is the branch of science that is concerned with patterns that show up everywhere in the world³. When we say that other sciences are concerned with restricted subsets of the physical world, we just mean that they’re concerned with picking out patterns in *some* of the systems to which the generalizations of fundamental physics apply. Ladyman and Ross (2007) put the point slightly differently, arguing that fundamental physics is fundamental in the sense that it stands in an asymmetric relationship to the rest of science. They claim that generalizations of the special sciences are not allowed to contradict the generalizations of fundamental physics, but the reverse is not true; if the fundamental physicists and the biologists disagree, it is the biologist who likely has done something wrong. They call this the “Primacy of Physics Constraint” (PPC). It seems to me that while this is certainly *true*—that is, that it’s certainly right that the PPC is a background assumption in the scientific project—the way I’ve

³ It is worth pointing out that it is indeed possible that there just are no such patterns in the world: it is possible that all laws are, to a greater or lesser extent, parochial. If that were true, then it would turn out that the goal underlying the practice of fundamental physics was a bad one—there just are no universal patterns to be had. Because of this possibility, the unity of science (in the sense we’ve been using it here) is an hypothesis to be empirically confirmed or disconfirmed. Still, even its disconfirmation might not be as much of a disaster as it seems: the patterns identified in the course of this search would remain legitimate patterns, and the discovery that all patterns are to some extent parochial would itself be incredibly informative. Many advances are made accidentally in the course of pursuing a goal that, in the end, turns out to not be achievable.

put the point here makes it clear *why* the PPC holds.

In contrast to fundamental physics, consider the project being pursued by one of the special sciences—say, molecular biology. Molecular biologists are certainly not interested in identifying patterns that hold everywhere in the universe; biologists have relatively little to say about what happens inside the sun (except perhaps to note that the conditions would make it difficult for life to prosper there). They are, instead, concerned with the behavior of a relatively small sub-set of regions of the universe. So far, the patterns they've identified have been observed to hold only on some parts of Earth, and that only in the last few billion years. However, it's worth noting that the search for habitable planets outside our own solar system is guided by the patterns identified by biologists studying certain systems here on Earth. This is an excellent case of an application of the kind of projectability we discussed above: biologists try to predict what planets are likely to support systems that are relevantly similar to the systems they study on Earth based on patterns they've identified in those terrestrial systems. It remains to be seen whether or not this project will prove fruitful. It's clearly no criticism of molecular biology to point out that it has nothing to say on the subject of what happens inside a black hole—that kind of system is (by design) outside molecular biology's domain of interest. Just as in the case of *SI-2* above, this restriction of domain lets molecular biologists focus their efforts on identifying patterns that, while they aren't universal, facilitate predictions about how a very large class of physical systems behave.

What exactly *is* the domain of inquiry with which molecular biology is concerned? That is, how do molecular biologists carve up the world so that the patterns they identify hold of systems included in that carving? It is rather unusual (to put it mildly) for the creation of a domain in this

sense to be a rapid, deliberate act on the part of working scientists. It is unusual, that is, for a group of people to sit down around a table (metaphorical or otherwise), pick out a heretofore unexplored part of the world for empirical inquiry, and baptize a new special science to undertake that inquiry. Rather, new sciences seem most often to grow out of gaps in the understanding of old sciences. Molecular biology is an excellent illustration here; the isolation of DNA in 1869—and the subsequent identification of it as the molecule responsible for the heritability of many phenotypic traits—led to an explosion of new scientific problems: what is the structure of this molecule? How does it replicate itself? How exactly does it facilitate protein synthesis? How can it be damaged? Can that damage be repaired? Molecular biology is, broadly speaking, the science that deals with these questions and the questions that grew out of them—the science that seeks to articulate the patterns in how the chemical bases for living systems behave. This might seem unsatisfactory, but it seems that it is the best answer we're likely to get: molecular biology, like the rest of science, is a work-in-progress, and is constantly refining its methodology and set of questions, both in light of its own successes (and failures) and in light of the progress in other branches of the scientific project.

This is an important point, and I think it is worth emphasizing. Science grows up organically as it attempts to solve certain *problems*—to fill in certain gaps in our knowledge about how the world changes with time—and is almost never centrally planned or directed. Scientists do the best they can with the tools they have, though they constantly seek to improve those tools. The fact that we cannot give a principled answer to the question "what parts of the world does molecular biology study?" should be no bar to our taking the patterns identified by molecular biology seriously. Just as we could not be sure that *R*, once identified, would hold in any

particular segment of *S* that we might examine, we cannot be sure of precisely what regions of the world will behave in ways that are consistent with the patterns identified by molecular biologists. This is not to say, though, that the molecular biologists have failed to give us any interesting information—as we saw, universality (or even a rigidly defined domain of applicability) is no condition on predictive utility. To put the point one more way: though the special sciences are differentiated from one another in part by their domains of inquiry, giving an exhaustive account of *exactly* what parts of the world do and don't fall into the domain of a particular science is likely an impossible task. Even if it were not, it isn't clear what it would add to our understand of either a particular science or of science as a whole: the patterns identified by molecular biology are no less important for our not knowing if they do or don't apply to things other than some of the systems on Earth in the last few billion years; if molecular biology is forced to confront the problem of how to characterize extraterrestrial living systems, it is certainly plausible to suppose that its list of patterns will be revised, or even that an entirely new science will emerge from the realization that molecular biology as thus far conceived is parochial in the extreme. Speculating about what those changes would look like—or what this new special science would take as its domain—though, is of little real importance (except insofar as such speculation illuminates the current state of molecular biology). Like the rest of the sciences, molecular biology takes its problems as they come, and does what it can with the resources it has.

If we can't say for any given special science what exactly its domain *is*, then, perhaps we can say a bit more about what the *choice* of a domain consists in—that is, what practical activities of working scientists constitute a choice of domain? How do we know when a formerly singular

science has diverged into two? Perhaps the most important choice characterizing a particular science's domain is the choice of what measurements to make, and on what parts of the world. That is: the choice of a domain is largely constituted by the choice to treat certain parts of the world as *individuals*, and the choice of what measurements to make on those individuals. Something that is treated as an individual by one special science might well be treated as a composite system by another[20]; the distinction between how human brains are treated by cognitive psychology (i.e. as the primary objects of prediction) and how they're treated by neurobiology (i.e. as aggregates of individual neural cells) provides an excellent illustration of this point. From the perspective of cognitive psychology, the brain is an unanalyzed individual object—cognitive psychologists are primarily concerned with making measurements that let them discern patterns that become salient when particular chunks of the physical world (that is: brain-containing chunks) are taken to be individual objects. From the perspective of neurobiology, on the other hand, brains are emphatically *not* unanalyzed objects, but are rather composites of neural cells—neurobiologists make measurements that are designed to discern patterns in how chunks of the physical world consisting of neural cells (or clusters of neural cells) evolve over time. From yet another perspective—that of, say, population genetics—neither of these systems might be taken to be an individual; while a population geneticist might well be interested in brain-containing systems, she will take something like *alleles* to be her primary objects, and will discern patterns in the evolution of systems from that perspective.

This echoes our discussion from **Section 1** about Sherman's levels-based approach to compatible pluralism. The problem is that, as we've seen, thinking of the divisions in terms of

“levels” isn’t appropriate, as it implies purpose-independent hard-wired divisions that the sciences latch on to. Discussions of whether or not two theories operate on the same level aren’t helpful, as the choice of how to carve up the world for a particular model is, like the world itself, messy and complicated. Rather than talking of levels, it’s better to talk of *spaces*, and rather than talking of *laws* it’s better to talk of *patterns*. Some patterns are easier to discern from one the perspective of one space, while others are easier to discern from another. This view is loosely akin to Don Ross' "rainforest realism:" a systematized version of Dennett's "stance" stance toward ontology (see Ross (2000), as well as Chapter 4 in Ladyman and Ross (2007), and Dennett (1991)). Ross' picture, like the one I have presented here, depicts a scientific project that is unified by goal and subject matter, though not necessarily by methodology or apparatus. It is one on which we are allowed to be frankly instrumentalist in our choice of objects—our choice of individuals—but still able to be thoroughly realists about the relations that hold between those objects—the patterns in how the objects change over time.

2.1 Pattern-Based Science: The Formal Proposal

The story of science is a story of progress through collaboration: progress toward a more complete account of the patterns in how the world evolves over time via collaboration between different branches of science, which consider different ways of carving up the same world. Individual sciences are concerned with identifying patterns that obtain in certain subsets of the world (though even within one science, a multitude of perspectives might be useful in discerning different patterns), while the scientific *project* in general is concerned with the overarching goal of pattern-based prediction of the world's behavior. Success or failure in this project is not absolute; rather, the identification of parochial or "weak" patterns can often be just as useful (if

not more useful) as the identification of universal patterns. Scientists identify patterns both by making novel measurements on accessible regions of the world and by creating models that attempt to accurately retrodict past measurements. The scientific project is unified in the sense that all branches of science are concerned with the goal of identifying patterns in how the physical world changes over time, and fundamental physics is fundamental in the sense that it is the most general of the sciences—it is the one concerned with identifying patterns that will generate accurate predictions for any and all regions of the world that we choose to consider. Patterns discovered in one branch of the scientific project might inform work in another branch, and (at least occasionally) entirely novel problems will precipitate a novel way of carving up the world, potentially facilitating the discovery of novel patterns; a new special science is born.

Consider the state space D of some system T —say, the phase space corresponding to the kitchen in my apartment. Suppose (counterfactually) that we take Newtonian dynamics to be the complete fundamental physics for systems like this one. If that is the case, then fundamental physics provides a set of *directions* for moving from any point in the phase space to any other point—it provides a *map* identifying where in the space a system whose state is represented by some point at t_0 will end up at a later time t_1 . This map is interesting largely in virtue of being valid for any point in the system: no matter where the system starts at t_0 , fundamental physics will describe the pattern in how it evolves. That is, given a list of points $[a_0, b_0, c_0, d_0 \dots z_0]$, the fundamental physics give us a corresponding list of points $[a_1, b_1, c_1, d_1 \dots z_1]$ that the system will occupy after a given time interval has passed. In the language of **Section 2**, we can say that fundamental physics provides a description of the patterns in the time-evolution of the room's *bitmap*: given a complete specification of the room's state (in terms of its precise location in

phase space) at one time, applying the algorithm of Newtonian mechanics will yield a complete specification of the room's state at a later time (in terms of another point in phase space).

This is surely a valuable tool, but it is equally surely not the *only* valuable tool. It might be (and, in fact, is) the case that there are also patterns to be discerned in how certain *regions* of the phase space evolve over time. That is, we might be able to describe patterns of the following sort: if the room starts off in any point in region P_0 , it will, after a given interval of time, end up in another region P_I . This is, in fact, the form of the statistical-mechanical explanation for the Second Law of Thermodynamics. This is clearly not a description of a pattern that applies to the “bitmap” in general: there might be a very large number (perhaps even a *continuous infinity*) of points that do not lie inside P_0 , and for which the pattern just described thus just has *nothing to say*. This is not necessarily to say that the project of identifying patterns like $P_0 \rightarrow P_I$ isn't one that should be pursued, though. Suppose the generalization identified looks like this: if the room is in a region corresponding to “the kitchen contains a pot of boiling water and a normal human being who sincerely intends to put his hand in the pot” at t_0 , then evolving the system (say) 10 seconds forward will result in the room's being in a region corresponding to “the kitchen contains a pot of boiling water and a human being in great pain and with blistering skin.” We can think of the “sincerely intends to put his hand in the pot” as being an assertion about location of the system when its state is *projected* onto a lower-dimensional subspace consisting of the configuration space of the person's brain. Again, this location will (obviously) be a *regional* rather than precise one: there are a large number of points in this lower-dimensional space corresponding to the kind of intention we have in mind here. Identifying these sorts of patterns is the business of the special sciences.

Not all regions will admit of interesting patterns in this way. This is the sense in which some ways of “carving up” a system’s space seem *arbitrary* in an important way. In many systems of interest to contemporary science--systems like the global climate--there will be a *very* large number of ways of specifying regions such that we won’t be able to identify any interesting patterns in how those *regions* behave over time. This is the sense in which some objects and properties seem *arbitrary* in problematic ways: carvings corresponding to (for example) grue-like properties (or bizarre compound objects like “the conjunction of the Queen of England’s left foot and all pennies minted after 1982”) just don’t support very many interesting patterns. Regions picked out by locutions like that don’t behave in ways that are regular enough to make them interesting targets of study. Even in cases like this, though, the patterns identified by fundamental physics will remain reliable: this (again) is the sense in which fundamental physics is *fundamental*. The behavior of even arbitrarily-specified regions—regions that don’t admit of any parochial patterns—will be predictable by an appeal to the bit-map level patterns of fundamental physics.

More precisely, then, the business of a particular special sciences consists in identifying certain *regions* of a system’s configuration space as instantiating enough interesting patterns to be worth considering, and then trying to enumerate those patterns as carefully as possible. A new special science emerges when someone notices that there exist patterns in the time-evolution of regions⁴ which have heretofore gone unnoticed. The borders of the regions picked out by the special sciences will be vaguely-defined; if the special scientists were required to give a complete enumeration of all the points contained in a particular region (say, all the possible

⁴ It might be appropriate to remind ourselves here that the *regions* under discussion here are regions of *configuration space*, not space-time.

configurations corresponding to “normal human observer with the intention to stick his hand in the pot of boiling water”), then the usefulness of picking out patterns of those regions would be greatly reduced. To put the point another way, there’s a very real sense in which the vagueness of the carvings used by particular sciences is (to borrow from computer science yet again) a feature rather than a bug: it lets us make reliable predictions about the time-evolution of a wide class of systems while also ignoring a lot of detail about the precise state of those systems. The vagueness might lead us to occasionally make erroneous predictions about the behavior of a system, but (as I argued in **Section 2**) this is not at all a fatal criticism of a putative pattern. The progress of a particular special science consists largely in attempts to make the boundaries of its class of carvings as precise as possible, but this notion of progress need not entail that the ultimate goal of any special science is a set of *perfectly* defined regions. To be a pattern is not necessarily to be a *perfect* pattern, and (just as with compression algorithms in information theory) we might be happy to trade a small amount of error for a large gain in utility. The scientific project consists in the identification of as many of these useful region/pattern pairings as possible, and individual *sciences* aim at careful identification of patterns in the evolution of particular regions.

There will often be overlap between the regions studied by one science and the regions studied by another. The “human with his hand in a pot of boiling water” sort of system will admit of patterns from (for example) the perspectives of biology, psychology, and chemistry. That is to say that this sort of system is one that is in a region whose behavior can be predicted by the regularities identified by all of these special sciences, despite the fact that the unique carvings of biology, psychology, and chemistry will be regions with very different shapes.

Systems like this one sit in regions whose time-evolution is particularly rich in interesting patterns, and scientists studying them will face very different model-building challenges than scientists studying less messy systems--a free electron in a vacuum, say.

3. Integrative Pluralism, Patterns, And Models

Let's return now to the problem we set out to consider in **Section 1**: the problem of how to understand the pluralism of models in fields like climatology. How can the account of science we've developed in the course of this paper help illuminate that puzzle? To answer that question, it's worth looking at how the view I've argued for here fits in with integrative pluralism (IP) as it's been articulated so far. Both views paint a picture of science that's flexible rather than rigid, messy rather than neat, and organic rather than mechanical; indeed, the story I've told here can (as I said) be taken as a contribution to the research program first defined in Mitchell (2002)--a contribution toward a novel way of seeing science in general that respects the kind of work being done in fields like climatology.

Mitchell writes "In complex natural phenomena, it is often the case that only parts of the array of contemporaneous scientific claims are in competition; other parts are compatible and, indeed, must be integrated into a multi-level explanation for explanation to succeed" (2009, p. 108). The heart of IP is the recognition that there are cases where competitive pluralism is correct, cases where the kind of guarded compatible pluralism she attributes to e.g. Kitcher (1991) is correct, and cases where *neither* approach is correct: "individual explanations making up the variety of explanations even for a single phenomenon...are not always competing; they are sometimes compatible and complementary" (Mitchell *op. cit.*, p. 109). Insofar as different perspectives on a single phenomenon--on the behavior of a single system--can contribute different information

about that phenomenon, they ought to be seen not as competing, not as coexisting in Sherman's non-interactive sense, where coexistence implies that they just have nothing to do with one another. Rather, different perspectives can *at least sometimes* be taken as not competing or just coexisting, but *complementing* one another. Integrating multiple complementary perspectives can help us understand the behavior of complex systems more fully, but only if we allow them to inform one another without demanding that they fight to the death in a winner-take-all competition.

The pattern-based account of science we've explored here provides a natural background against which to understand IP's emphasis on complementary integration. If, as I've argued, science consists in the related (but distinct) acts of carving the world into particular state spaces and then trying to discern the patterns in the time-evolution of paths through those state-spaces, it is only natural that robust understanding of systems that admit of many different interesting carvings--systems like the global climate--can come only through the integration of many different perspectives. An emphasis on patterns and carvings--rather than laws and levels--helps highlight the fact that even *very* similar-seeming models of the same system need not be thought of as being in competition with one another. Just as in the toy science we discussed in **Section 2.1**, a data-set collected by climatologists might admit of very many different approximate patterns, all of which can be useful in predicting the future behavior of the system represented by the data-set.

If we think of science's job as consisting in exploring and mapping these state-spaces, the idea that different models should be either competitive or compatible in an isolated way looks very puzzling: why shouldn't we allow for a variety of complementary perspectives, each of

which can be taken to illuminate the (incomplete) insights provided by all the others? The “mapping” language, aside from being a description of the project of working with abstract mathematical state-spaces, can be an instructive metaphor here. In just the same way that a number of different maps of the same region of a particular city might each give a very different picture of the same terrain (and provide the most accurate complete depiction when taken as an ensemble), a number of general circulation models of the global climate might give a different picture of the relevant dynamics, and provide a more accurate depiction when their predictions are treated as a complementary ensemble. We might demand a single comprehensive map of (say) New York City, but I’m not sure why anyone would want such a thing: it would be unusably cluttered, and thus no map at all. Questions like “is *that point* in space the intersection of 125th street and Saint Nicholas Avenue, or the 125th street D train subway stop, or the location of the lower Harlem McDonald’s restaurant?” are just ill-posed. Those possible answers aren’t in competition with one another, but neither are they totally divorced: the subway stop and fast-food restaurant are there partially *because* it’s an intersection of major streets. The different answers complement each other, and can paint a more complete (and useful) picture of the world when integrated than they ever could as parts of competing maps (or when all shoved into a single comprehensive map, for that matter).

The same seems true of scientific models. Different GCMs, for example, emphasize different aspects of the global climate’s dynamics in virtue of utilizing slightly different carvings *and/or* attending to different approximate patterns in the same carving. Trying to distill One Model to Rule Them All out of this complex, multi-level array of causal structures is a project that’s not just pointless, but potentially detrimental to the progress of the field--a waste of time and

resources that could be better spent exploring and integrating relationships between existing perspectives. Climate science is quite right to focus on developing a diverse family of models--a family that isn't competitive or coexistent, but *complementary*.

Of course, in some cases attempted integration will lead to synthesis--in the process of exploring the relationship between two different ways of looking at the same system, we might come to understand that the interesting aspects of both models can be better captured in a single model that integrates (and discards) some aspects of both. This is a strong virtue of the IP-approach to model-building: it is flexible enough to allow for instances of model elimination (and synthesis) when appropriate--it just doesn't demand that all instances of model pluralism be shoehorned into this competitive approach. In the end, this results in a more realistic, robust, and flexible picture of scientific methodology: one that's prepared to confront the challenges of the 21st century head-on.

3.1 Practical Lessons: Ensemble Modeling in Climate Science

Thus far, our discussion has been mostly abstract. I'd like to close with a few words about how the way of looking at science and model-building that I've been advocating here might dovetail with working scientific practice, with particular reference to climatology. Recall the puzzle: it doesn't seem like either competitive pluralism or compatible pluralism (in the Sherman sense) is sufficient to explain the practice of constructing climate models. They're simply too numerous and heterogeneous, and yet climate scientists seem to take this as a strength of the field rather than a defect to be remedied. How can we apply the foregoing discussion to help us understand what's going on here?

The key lies in looking at the way that concrete predictions are actually generated from

climate models. As (xxx) notes, it isn't quite right to think about individual climate models as being exactly fully-fledged models in their own right--or at least not in the same sense that we might think of (say) the Standard Model in physics as being a fully-fledged model in its own right. Climate models are a different sort of beast, crafted to thrive in a different sort of environment. Rather than delivering competing (or entirely independent) predictions, climate models are best thought of as delivering different aspects of one *ensemble* of predictions which can only be taken to reflect the way the world is when seen *as* an ensemble: climate models must be integrated with one another in order to deliver useful predictions.

The rationality of this approach looks obscure until we notice that the climate system itself is very unlike the systems being modeled in (say) physics, or even chemistry. It is a system that is rich in approximate patterns, and one that can benefit from examination from many different perspectives. It is, in other words, a paradigmatic complex system; it should come as no surprise that the practice of investigating it diverges from the practice of investigating more simple systems. Integrative pluralism as developed both here and by Mitchell can illuminate the practice of ensemble-based modeling by emphasizing that complex systems are best understood not as stacks of tier-like levels, but as fluid collections of mutually-constraining (and mutually-informative) *patterns*. The business of science is both the mapping of those patterns and the *integration* of those maps into ensembles in order to make the best predictions possible.

This suggests new challenges that climate scientists must face, and which have yet to be well-explored by philosophers of science. The challenge of how to understand this process of integration across many different complementary levels is perhaps the deepest, and it is one that is worthy of immediate attention. Fortunately, the formal tools for working with systems like

this are already under development in complexity theory: Bar-Yam (xxx) and (xxx) deal explicitly with the analysis of multi-level systems, and the integration of information about patterns in those systems drawn from different levels (where “level” has more in common with what I’ve been calling “perspectives” or “carvings” than it does with levels in the Sherman or Oppenheim & Putnam sense), and Wang et. al. (2012) demonstrates the superiority of ensemble-based approaches when dealing with models operating under rapidly fluctuating (or uncertain) boundary conditions--precisely the circumstance in which climate science finds itself. The future of the integrative pluralism research program lies in integrating insights like these from complexity theory proper into philosophy of science. There is a lot of groundwork to be done there:. A working definition of “complexity” that’s applicable to a wide-range of physical systems has yet to be established, and many of the existing literature and tools are highly abstract and mathematical in nature: it will be an achievement just to clearly articulate how to apply these tools and definitions to physical systems rather than strings of bits, or other mathematical objects. However, the perspective I’ve outlined here suggests a way forward, and lays the groundwork for what’s to come.

Works Cited

- Aharonov, Y., & Bohm, D. (1959). Significance of electromagnetic potentials in quantum theory. *Physical Review* , 485-491.
- Beatty, J. (1987). Natural Selection and the Null Hypothesis. In J. (. Dupre, *The Latest on the Best: Essays on Evolution and Optimality* (pp. 53-76). Cambridge, MA: The MIT Press.
- Blume-Kohout, R., & Zurek, W. (2006). Quantum Darwinism: Entanglement, Branches, and the Emergent Classicality of Redundantly Stored Quantum Information. *Physics Review A* (73).
- Chaitin, G. (1975). Randomness and Mathematical Proof. *Scientific American* , 47-52.
- Dennett, D. C. (1991). Real Patterns. *The Journal of Philosophy* , 27-51.
- Hooker, C. (. (2011). *Handbook of the Philosophy of Science, Volume 10: Philosophy of Complex Systems*. Oxford: Elsevier.
- Hooker, C. (2011). Conceptualizing Reduction, Emergence, and Self-Organization in Complex Dynamical Systems. In C. (. Hooker, *Handbook of the Philosophy of Science, Vol. 10: The Philosophy of Complex Systems* (pp. 195-222). Elsevier.
- Houghton et. al, (. (2001). *Climate Change 2001: The Scientific Basis*. Cambridge, UK: Cambridge University Press.
- Kitcher, P. (1991). The Division of Cognitive Labor. *Journal of Philosophy* , 5-22.
- Knutti, R. (2008). Why should we believe model predictions of future climate change? *Philosophical Transactions of the Royal Society A* , 4647-4664.
- Ladyman, J., & Ross, D. (2007). *Every Thing Must Go*. New York City: Oxford University Press.
- Maudlin, T. (2007). *The Metaphysics Within Physics*. New York City: Oxford University Press.
- Mitchell, S. (2002). Integrative Pluralism. *Biology and Philosophy* , 55-70.
- Mitchell, S. (1992). On Pluralism and Competition in Evolutionary Explanations. *American Zoologist* , 135-144.
- Mitchell, S. (2009). *Unsimple Truths: Science, Complexity, and Policy*. London: The University of Chicago Press.
- Mitchell, S. (2004). Why Integrative Pluralism? *Emergence: Complexity & Organization* , 81-91.
- Oppenheim, P., & Putnam, H. (1958). Unity of Science as a Working Hypothesis. In *Minnesota Studies in the Philosophy of Science* (Vol. 2). Minneapolis: University of Minnesota Press.
- Parker, W. (2006). Understanding Pluralism in Climate Modeling. *Foundations of Science* , 349-368.
- Ross, D. (2000). Rainforest Realism: A Dennettian Theory of Existence. In D. Ross, A. Brook, & D. (. Thompson, *Dennett's Philosophy: A Comprehensive Assessment* (pp. 147-168). The MIT Press.
- Sherman, P. (1988). The Levels of Analysis. *Animal Behavior* , 616-619.