Law and Explanation in Biology: Invariance Is the Kind of Stability That Matters
Author(s): Jim Woodward
Source: *Philosophy of Science*, Vol. 68, No. 1 (Mar., 2001), pp. 1-20
Published by: The University of Chicago Press on behalf of the Philosophy of Science Association
Stable URL: http://www.jstor.org/stable/3081021
Accessed: 24/03/2009 14:30

Law and Explanation in Biology: Invariance is the Kind of Stability That Matters*

Jim Woodward†‡
Division of Humanities and Social Sciences
California Institute of Technology

This paper develops an account of explanation in biology which does not involve appeal to laws of nature, at least as traditionally conceived. Explanatory generalizations in biology must satisfy a requirement that I call *invariance*, but need not satisfy most of the other standard criteria for lawfulness. Once this point is recognized, there is little motivation for regarding such generalizations as laws of nature. Some of the differences between invariance and the related notions of stability and resiliency, due respectively to Sandra Mitchell and Brian Skyrms, are explored.

1. Introduction. Are there laws in biology? This is a question that has attracted a great deal of attention recently from philosophers of biology.[1] My aim in this paper is twofold. First, I want to illustrate how the account of explanation and invariance that I have applied elsewhere to examples from the physical and social sciences can also be applied to biological examples, and to explore the implications of that account for the role of laws in biology. Second, I want to compare my account with some related ideas, due to Sandra Mitchell and to Brian Skyrms, about the role of stability in assessments of lawfulness.

*Received August 2000; revised November 2000.

†Send reprint requests to the author, Division of Humanities and Social Sciences, 101–40, California Institute of Technology, Pasadena, CA 91125.

‡Earlier versions of this paper were read at MIT and at the University of Pittsburgh. I thank members of the audience at both institutions and especially Evelyn Fox Keller, Peter Machamer, and Sandy Mitchell for helpful comments.

1. In addition to Mitchell 1997 and 2000, see, for example, Beatty 1995 and 1997, Brandon 1997, Cooper 1998, Sober 1997, Waters 1998.

Philosophy of Science, 68 (March 2001) pp. 1–20. 0031-8248/2001/6801-0001$2.00

2. Mitchell on Lawfulness and Stability. In a recent series of papers (1997, 2000), Sandra Mitchell has argued that to answer the question of whether there are laws in biology we need to rethink the notion of law and its role in biological practice. She advocates replacing the standard law vs. accident dichotomy with a framework for the classification of explanatory generalizations that admits of degrees. One of the "dimensions" of scientific law which receives the most attention within this framework is what Mitchell calls stability, which is roughly the extent to which a generalization is contingent on conditions that are stable across space and time. According to Mitchell, biological generalizations like Mendel's laws are less stable (more contingent) than, say, paradigmatic fundamental physical laws like the conservation of mass-energy but more stable than generalizations like:

(2.1) All the coins in Goodman's pockets are copper.

Biological generalizations thus occupy an intermediate position on the "continuum of contingency". They are stable enough to play the same sort of role that laws play in other areas of science: they can function so as to represent "causal knowledge" and can be used to predict, explain and to guide interventions. (2000, 249) Because biological generalizations can function in these ways, we may legitimately describe them as laws, even if they lack features, such as exceptionlessness, traditionally ascribed to laws.

As Mitchell notes, these ideas are similar in a number of ways to ideas that I have defended elsewhere (see e.g., Woodward 1997, 2000). Like her, I have argued that the traditional way of thinking about generalizations in the special sciences (including biology) in terms of an all or nothing law/accident dichotomy is misguided. Like her, I too favor replacing this dichotomy with a framework for the classification of explanatory generalizations that admits of degrees. Moreover, as Mitchell notes, her notion of stability bears a family resemblance both to the notion of invariance (understood as the stability of a relationship under some set of interventions) which plays a central role in the approach that I advocate, and to Brian Skryms' notion of resiliency (Skyrms 1980). Despite this general similarity, I will suggest (Sections 8 and 9) that invariance, stability, and resiliency differ in important respects and that we need the notion of invariance rather than stability or resiliency if we are to distinguish adequately between explanatory and non-explanatory generalizations. More generally, while I agree with Mitchell that some notion of stability or robustness is the right notion to look at if we are to understand how biological generalizations function in explanation, I will try to show that the details of how exactly we understand this notion matter crucially.

3. Background. In the recent literature on whether there are laws in biology, genuinely substantive issues are often entangled with issues that seem largely verbal or terminological—in particular, issues about how strictly or permissively we should use the word "law." Like Mitchell, I think that the key to separating substantive from terminological issues is to ask what turns on whether we take generalizations like Mendel's to be laws of nature or to have some other status. What role or function do laws play that could not be played by non-lawful generalizations? My assumption in what follows is that much of the interest in this question derives from the idea that laws are required for successful explanation (or what I shall take to be the same thing) the representation of causal relationships; hence that the status of biology as an explanatory rather than a "merely descriptive" science turns on whether it contains laws. I thus begin with some motivating background and stage setting regarding the notions of "explanation" and "law."

Despite the vicissitudes of the DN model of explanation and its more recent descendants, the majority of philosophers probably continue to subscribe to the thesis that all explanation in some way requires or "involves" laws of nature. Much of the appeal of this nomothetic thesis (as I shall call it) derives from the fact that it seems undeniable that generalizations play an important role in many explanations and we lack a generally accepted positive theory of how generalizations function in explanation except as laws. In the absence of such an alternative account, even philosophers who are well aware of the great distance between the generalizations that figure in explanations in biology and the other so-called special sciences, and paradigmatic examples of physical laws feel forced to assimilate the former to the latter. If we want to vindicate the idea that generalizations like Mendel's laws can be used to explain and the only possibilities are such generalizations are either laws of nature or completely accidental and non-explanatory, we seem to have no alternative but to try to shoe-horn them into the category of laws of nature, however awkward the fit may be.

This dialectical situation is further complicated by the fact that there is no consensus on what features a generalization must possess if it is to count as a law.[2] While there are of course a number of familiar and frequently invoked criteria for lawfulness, some of these (such as the requirement that laws be exceptionless) are violated by paradigmatic laws while others, at least as standardly formulated, fail to distinguish between accidental and lawful generalizations.[3]

2. Both the absence of consensus and the limitations of some of the standard criteria for lawfulness are discussed in Salmon 1989.

3. This is true, for example, of the suggestion that laws "support" counterfactuals while accidental generalizations do not; see section 5.

Absent such a consensus, we lack a clear purchase on what we are arguing about when we ask whether there are laws in biology. In view of this, I propose to bypass the question of what sorts of features a generalization should possess in order for it to count as a "law" (and whether biological generalizations possess those features) and to focus instead on the following question: what are the characteristics that a generalization must possess if it is to figure in explanations? I take this to be the substantive issue that lurks behind the terminological disputes in the recent philosophical literature about how permissively we should use the word "law."

In order to answer this question I begin with a very brief sketch of some ideas about explanation that I have described in more detail elsewhere. My interest here is not in providing a defense of these ideas but rather in illustrating how they may be applied to the issue of how generalizations function in biological explanation.[4]

4. Explanation, Invariance, and Laws. On my view, the key feature that a generalization must possess if it is to figure in explanations is *invariance*. Invariance is a kind of robustness or stability property: a generalization is invariant if and only if it would continue to hold under some range of physical changes involving *interventions*. Heuristically, the notion of an intervention represents an attempt to capture, in non-anthropomorphic language that makes no reference to notions like human agency, the conditions that would need to be met in an ideal experimental manipulation of the value of some variable X performed for the purpose of determining whether X causes a second variable Y. Slightly more precisely, an intervention on X (with respect to Y) is a causal process that directly changes the value of X in such a way that, if a change in the value of Y should occur, it will occur only through the change in the value of X and not in some other way. This in turn requires, for example, that the intervention not be correlated with other causes of Y except for those causes of Y (if any) that are causally between X and Y and that the intervention not affect Y independently of X—that is, the intervention should not affect Y via a causal route that does not go through X. As an illustration suppose that (4.1) X is a common cause of Y and Z . Then a manipulation of X that changes Y will not count as an intervention on Y with respect to Z, since in this case the manipulation affects Z via a route (the route that directly connects X to Z) that does not go through Y.[5] On the other hand, if, in such a common cause structure, we were to change the value of Y by means

4. Readers interested in a fuller defense of these ideas are referred to Woodward 2000.

5. For a more precise characterization of the notion of an intervention and of the notions of a causal route and causal betweeness in interventionist terms, see Woodward 2000 and Woodward forthcoming.

of some randomizing device that is not correlated with and is causally independent of changes in the value of *X*, and also does not directly affect *Z*, this would count as an intervention on *Y*.

A generalization that relates changes in (or describes a correlation between) one set of variables and another is invariant if and only if it would continue to hold (or would be stable) under some intervention on variables figuring in that relationship. For example, if *X* is genuinely a cause of *Y* in the common cause structure (4.1), we would expect the generalization describing how changes in *X* are correlated with changes in *Y* to be invariant under at least some interventions that change *X*. By contrast, the correlation between the joint effects *Y* and *Z* of the common cause *X* is *not* invariant under any interventions on *Y* (or *Z*), since all such interventions will disrupt the relationship between *Y* and *Z*.

What is the connection between invariance and explanation? On my view, to explain an explanandum is to show how changes in it counterfactually depend on changes in the factors cited in the explanans, or to express the same idea in a slightly different way, to show how the explanandum would have been different if the factors cited in its explanans had been different in various ways. The relevant notion of counterfactual dependence (or of answering a what-if things-had-been- different question) is captured by counterfactuals the antecedents of which have to do with interventions: to explain why an explanandum *Y* takes some particular value we need to identify some variable *X* and a generalization *G* linking *X* to *Y* such that, according to *G*, some range of changes in the value of *X* that are due to interventions are associated with changes in the value of *Y*. This requires that the generalization *G* must be invariant under some interventions on *X* that change the value of *Y*.

When a generalization relating *X* and Y is invariant in this way, we may think of it as telling us how to manipulate *Y* if (perhaps contrary to actual fact) it *were* possible to intervene to change *X*. In this sense the theory just sketched embodies a "manipulationist" conception of explanation. Thus, in the common cause example above, one cannot appeal to the value of *Y* and the correlation between *Y* and *Z* to explain the value of *Z* because intervening on the value of *Y* is not a way of manipulating the value of *Z*.

It should be clear from these remarks why explanatory generalizations need to be invariant under interventions rather than meeting some other stability condition. The correlation between *Y* and *Z* in the common cause example (4.1) is stable or would continue to hold under many sorts of changes. For example, it is stable under changes in the value of *Y* produced by changes in *X* as well as under changes in such background conditions as the price of tea in China. Mere stability under some or even many changes is not sufficient for explanatoriness.

What implications does this account have for the role of laws in explanation? Paradigmatic laws, such as the field equations of General Relativity, are invariant generalizations and it is their invariance that endows them with explanatory import. Nomothetic models of explanation are thus correct in insisting that laws play a central role in some explanations. However, as we shall also see, a generalization can be invariant and hence figure in explanations even though it differs in important respects from paradigmatic laws and even though it fails to satisfy many of the traditional criteria for nomological status.[6] On my view, Mendel's laws, as well as many other explanatory generalizations drawn from the special sciences, have exactly this character. Among the traditional criteria for lawhood, only one—support for counterfactuals—is directly relevant to the question of whether a generalization is invariant and even this criterion requires considerable reinterpretation if it is to be acceptable (cf. Section 5). In particular, a generalization may be invariant and hence explanatory even if it has exceptions, even if it makes reference to particular objects and spatio-temporal locations, and even if it is not part of any systematic or unified theory. Conversely, a generalization can fail to be invariant even if it is exceptionless, contains purely qualitative predicates, is confirmable by a limited number of instances, and plays a central, unifying role in some theory.

Seen from the perspective of the invariance-based approach, the fuzziness and inadequacy of the traditional criteria for lawfulness are neither surprising nor alarming. If what we are interested in is which generalizations can figure in explanations, we don't have to address the question of whether such generalizations are "laws" and hence we don't need to settle the issue of which of the traditional criteria are defensible. Instead, we can simply bypass the traditional criteria and focus directly on the question of whether the generalizations of interest are invariant in the right way. In other words, it simply doesn't matter, independently of whether or not generalizations like Mendel's are invariant, whether we choose to regard them as genuine laws. We can, if we wish, stipulate that the word "law" must be used in such a way that all invariant generalizations are laws. If so, because they are invariant, Mendel's laws and many other biological generalizations will qualify as laws and the nomothetic thesis will be correct. Alternatively, we may choose to regard similarity to paradigmatic laws and satisfaction of most of the traditional criteria as necessary for lawhood. If so, generalizations like Mendel's will probably not count as laws. Nonetheless, as long as such generalizations are invariant in the right

6. It is also the case that an argument can meet the requirements for DN explanation and yet fail to provide information about counterfactual dependence and hence fail to be explanatory. Standard counterexamples to the DN model, involving explanatory irrelevancies or asymmetries, have this character (as argued in Woodward 2000).

way, they can figure in explanations. I take this to explain Ernst Mayr's observation that the question of whether there are laws in biology "is of little relevance for the working biologist" (1982, 32; quoted in Mitchell 2000, 249).

5. Explanation and Invariance in Biology. I turn now to some biological illustrations of these ideas involving Mendel's "law" of segregation:

(S) With respect to each pair of genes in a sexual organism, 50% of the organism's gametes will carry one representative of that pair and 50% will carry the other representative of that pair.

There are of course obvious dangers associated with generalizing from a single case. On the other hand, (i) space is limited, (ii) a very large portion of the recent discussion of whether there are laws in biology has focused on the status of this very generalization, (iii) the case for the account of explanation I favor and the role of invariance is, if anything, much easier to make in other areas of biology, such as molecular biology, where the association between successful explanation and the identification of invariant relationships that are potentially exploitable for purposes of manipulation and control is often explicitly made by biologists themselves.[7]

I begin by reminding the reader of some of the ways in which (S) figures in elementary explanatory evolutionary models. First, from (S) and the assumption of random mating, one can derive the so-called Hardy-Weinberg law, which tells us that in the absence of various evolutionary forces—mutation, migration, drift and selection—genotypic frequencies will reach equilibrium after one generation with the equilibrium frequencies depending on the allele frequencies with which we began. One can then use the Hardy-Weinberg law in conjunction with additional assumptions about differential fitness of various genotypes to explain how the frequencies of those genotypes will change in response to natural selection. Consider a single locus model in which the *A* allele is dominant and the *a* allele recessive with the *Aa* genotype identical in phenotype and fitness to the *AA* genotype and both superior in fitness to the *aa* genotype. By using (S) and the other assumptions of this model one can readily derive that after one generation the frequency of *a* will decrease by an amount that is a function of its initial frequency and it relative fitness. In particular, the change Δq in allele frequency *a* of will be:

(5.1) $\Delta q = q\,(1 - wq2)/\,1 - wq2 - (q - wq2)/1 - wq2 = -wpq2/1 - wq2$

where *p* is the initial frequency of the *A* allele, *q* the frequency of the *a*

7. As an illustration, see Weinberg (1985).

allele and w the relative fitness of the recessive homozygote. With no changes in relative frequencies this process will continue in subsequent generations until the a allele is eliminated. By contrast, in the frequently discussed case of heterozygote superiority, in which the heterozygote Aa is more fit then either of the homozygotes AA and aa, one can show that selection will lead to a polymorphic equilibrium in which both the A and a alleles are maintained in the population, the frequency of each being a function of the fitnesses of the various genotypes.

Models and derivations of this sort are commonly regarded as explaining why gene frequencies change or fail to change over time. Following the remarks in Section 4, I suggest that what makes such derivations explanatory is that they answer a set of counterfactual or what-if-things-had-been-different questions about their explananda. This gives us a sense for the factors or conditions on which these explananda depend. For example, given the assumptions that figure in the derivation of the Hardy-Weinberg law, one can see that changing the values of the initial allelic frequencies will not change whether an equilibrium is reached but will change the genotype equilibrium frequencies—we can see how these frequencies would have been different if the initial allelic frequencies had been different. We can also see from this derivation how matters would be different in populations that do not conform to Mendelian segregation. In such populations even in the absence of migration, selection, etc., the Hardy-Weinberg law will not hold. Instead, one form of the gene will completely replace the other form in the population.

Similarly, in the case in which there is selection operating at a single locus against a recessive homozygote one can see how the outcome would (or would not) have been different in various ways if the initial frequencies p and q of the alleles A and a and the relative fitness w of the recessive homozygote had been different. In particular, the derivation shows us that provided selection is allowed to run on long enough, and no counteracting forces are operative, the recessive homozygote will always be eliminated regardless of the particular values of p, q and w. By contrast, the rate at which the allelic frequencies change in response to selection depends on the exact value of the selection coefficient w and the initial allelic frequencies. In this case, as well as in the previous examples, the model shows us how the change Δq in allelic frequency per generation would change if the value of w or q were changed—that is, how Δq would have been different had such changes occurred.

Parallel remarks apply to case of heterozygote dominance. Here too, the sort of model described above allows us to see what the maintenance of polymorphic equilibrium depends on and how the equilibrium frequency would be different in various ways, depending on the relative fitness of the heterozygote in comparison with two homozygotes.

Suppose that we agree that above explanations involving (S) work by conveying counterfactual information in the manner just described. What features must (S) possess if it is to play a role in conveying such information? As suggested above, the key feature is that (S) must be stable or invariant in the right way in the population *P*. We can motivate this requirement in the following way. Since the general explanatory strategy employed in the models described above works by using (S) in conjunction with information about initial genotype frequencies and the fitness of various genotypes to explain patterns of changes in those frequencies, (S) must not only hold for the actual values of the genotype frequencies and fitnesses in *P* but must also be such that it continues to hold (at least approximately) under some range of changes in the values of those frequencies and fitnesses in *P*. That is, it should be true that some range of changes in, say, the frequency of the *Aa* genotype or in its relative fitness in *P* do not by themselves disrupt (at least over the time scale of changes in gene frequency we are trying to explain) the mechanism of normal Mendelian segregation. If (S) were to hold when, say, the initial frequency of the *a* allele in *P* is q^* but were to break down when the frequency of a differed significantly from q^* , (S) would not be invariant under changes in q in *P* and we could not appeal to it in support of counterfactual claims about the conditions under which the frequency of q would change. Similarly, if (S) were to hold only when the relative fitness of the *Aa* and *aa* genotype in *P* took certain very restricted values, but not more generally, the models described above would no longer show us how the rate of allelic frequency change depends on w. This condition—that (S) is such that it is invariant or would continue to hold under changes in q and w—is plausible for many populations conforming to (S). The mere fact that a genotype becomes more or less fit as a consequence of some environmental change or that it increases or decreases in frequency will not by itself cause a shift from Mendelian to non-Mendelian patterns of segregation.

I emphasize that what matters for the use of (S) to explain gene frequencies in some particular population *P* is whether (S) is invariant in the right way under interventions *in P* and not whether various *other* populations, different from *P*, conform to (S). In other words, the counterfactuals that matter for successful explanation, both in this case and more generally, have to do with what would happen under interventions that change the values of the explanans variables for the very system whose behavior we are trying to explain (what we might call "same object counterfactuals") , rather than with counterfactuals that describe the behavior of other systems ("other object counterfactuals").[8] If, as I claim, explan-

8. Many patently accidental generalizations "support" counterfactuals in some sense. For example, (5.4) "All of the balls in this urn are red" seems to support the counter-

atory generalizations are required only to support same object counterfactuals, they may correctly describe what would happen under interventions in some system of interest while failing to correctly describe (or while failing to apply to) the behavior of other systems. The idea that same object counterfactuals are the relevant ones from the point of view of explanation thus goes along with the intuition (which I mean to endorse) that the explanatory status of Mendel's laws with respect to those populations that do exhibit Mendelian segregation is not impugned by the existence of other populations for which those laws fail to hold.

6. Degrees of Invariance and Explanatory Depth. The conditions on explanation described above require that explanatory generalizations be invariant under *some* (but not necessarily *all*) interventions. Thus, in the examples above, it is not necessary, in order for (S) to be explanatory, that it be invariant under all possible interventions or changes in background conditions. Instead, it is necessary only that (S) tell us how gene frequencies would change under some range of interventions on initial frequencies and fitnesses. This use of the existential rather than the universal quantifier reflects an important feature of the notion of invariance that is crucial to its usefulness. This is that invariance is a relative notion or a matter of degree rather than an absolute or all-or-nothing notion: a generalization can be invariant under some interventions but not others. This feature of the notion of invariance allows us to avoid various problems that are due to the dichotomous character of the traditional law vs. accident framework.

Consider a non-biological illustration: a particular sort of spring S conforming to a version of Hooke's law (H) $F = -kX$ where X is the extension of the spring, F the restoring force it exerts, and k a constant characteristic of springs of sort S. Suppose that (H) is invariant under some range of interventions that change X for springs of sort S within the interval between x_1 and x_2. This means that (H) correctly describes what the restoring force of the spring would be under experimental manipulations of the extension of the spring within this interval. Nonetheless, it is clear that if we intervene to make the extension of the spring too large, the generalization (H) will break down—the spring will no longer exert a linear restoring force and in fact may break. While (H) is invariant under some interventions on the extension of springs of sort S, it is not invariant

factual (5.5) "If a ball were to be withdrawn from this urn it would be red." The account developed above associates the non-explanatory status of (5.4) with its failure to support "same object" counterfactuals involving interventions: it is not true that an intervention that changes whether a ball is in the urn will change whether it is red. Examples like this show that to merely ask whether a generalization has "counterfactual force" in evaluating its explanatory status is insufficiently discriminating—it matters crucially whether the right sort of counterfactual is supported.

under all such interventions. (H) will also break down under certain changes in background conditions—changes in variables that do not explicitly figure in (H)—such as extreme temperatures. In addition, (H) has narrow scope—there are many other sorts of springs whose behavior is not governed by (H). In all these respects, (H) is typical of most explanatory generalizations in the special sciences, including (S).

Nonetheless, despite the fact that (H) is not exceptionless and has narrow scope (and hence lacks several of the features traditionally required of laws of nature), we may, according to the account of explanation described above, appeal to (H) and to the extension $X = x$, of some particular spring s of sort S to explain why it exerts a restoring force of magnitude $F = f$ as long as the conditions described above are met—as long, that is as (H) holds for the actual extension and force exerted by the spring and as long as (H) is invariant under some range of interventions that changes the value of X from $X = x$ to some new value. Clearly, these conditions can be met even if (H) fails to be invariant under various other changes in its extension. A similar point holds for (S).

Although it seems plausible that (H) can be used to explain why some particular spring conforming to (H) exerts the restoring force that it does, such an explanation strikes us as rather shallow. I fully agree with this assessment and see it as closely connected to the fact that (H) is invariant under only a rather limited range of interventions. In my view, a deeper explanation of the behavior of the spring would be provided by a set of generalizations that are invariant under (what we might intuitively regard as) a wider or more important range of interventions and that might be used to answer a wider range of what-if-things-had-been-different questions about the spring.[9]

The picture that emerges from all of this thus involves both a continuum and a threshold. Some generalizations, such as (2.1) or (4.1) above are not invariant under any interventions at all and hence are not explanatory. In addition we can distinguish among generalizations that are invariant under at least some interventions with respect to the range or kind of interventions under which they are invariant. Fundamental laws which are invariant under a wide range of interventions are at one end of this continuum; generalizations like (H) are closer to the other end.

7. Invariance and Explanation in Biology Revisited. With the example of (H) and the idea that a generalization can be invariant under some but not all interventions and hence explanatory firmly in mind, let us return to the question of the status of Mendel's generalizations. In a very inter-

9. For an attempt to be more precise about the notion of invariance under a wide range of interventions, see Woodward 2000, and Hitchcock and Woodward forthcoming.

esting series of papers (e. g., 1995) John Beatty has advanced what he calls the Evolutionary Contemporary Thesis, according to which "all generalization about the living world are [either] just mathematically, physical or chemical generalizations or [if they are] distinctively biological describe contingent outcomes of evolution" (46–47). He takes this to mean that "there are no laws of biology. For whatever laws are, they are supposed to be more than just contingently true." (46)

Beatty illustrates this thesis by drawing attention to two related features of Mendel's law of segregation. First it has a number of "exceptions." One of best known involves meiotic drive, which occurs when an allele influences meiosis in such away that it has a greater than 50% chance of ending up in a gamete, rather than the 50% chance that Mendelian segregation would require. When Mendel's law is understood along the lines of (S), this phenomenon represents a genuine violation of the law (i.e., a case in which the antecedent but not the consequent of the law holds) and not a mere failure of the law to apply in the sense that its antecedent is not satisfied. Second, the widespread prevalence of Mendelian segregation is itself the result of the operation of natural selection. That is, if as appears to be the case, most gene pairs in sexual organisms conform to (S), this is because natural selection has operated in such a way as to produce this outcome—because segregation in accord with (S) conferred a selective advantage of some kind. If past histories of most organisms had been sufficiently different, and if they had been subject to sufficiently different selective forces, nature would have contained few if any genes which segregate according to Mendelian ratios. Thus whether violations of Mendel's laws occur at all and whether they are common or rare is contingent on the course of evolution. To express this idea in the language of this paper, (S) fails to be invariant under many possible changes in selection pressure. This does not mean just that there are at present organisms and populations in which meiosis fails to conform to (S), but rather that even for those types of organisms that presently conform to (S), there are possible changes, due to natural selection which would make it the case that those types of organisms or their descendants would violate (S).

Beatty infers from these features that Mendel's law is not "necessary" and hence is not a real law. Beatty does not explore the implications of this claim for explanatory practice in biology but if we accept both the nomothetic thesis and Beatty's arguments, it appears to follow that the law of segregation (and indeed all distinctively biological theory) is unexplanatory. The most common strategy for avoiding this conclusion has been to accept the nomothetic thesis and to search for a somewhat weakened or watered down notion of law, according to which generalizations like (S) may qualify as laws and hence as explanatory.

One of the central claims of this paper is that this response is unnecessary and that it simply distracts us from understanding how explanations that appeal to (S) work. In assessing the explanatory import of (S), what matters is not whether we decide to bestow on it the honorific "law," but rather the range of what-if-things-had-been-different questions it can be used to answer and the range of changes over which it is invariant. As long as (S) is invariant in the right way, it doesn't matter whether it has exceptions or is contingent on the course of evolution in the way that Beatty describes—it still can be used to explain. In particular, what is crucial to the explanatory status of (S) in the evolutionary models described above is that it be invariant under some range of interventions on the explanans variables q and w in the population P whose behavior we are trying to explain. Moreover, it should also be clear that this sort of invariance in (S) is perfectly compatible with (S) failing to be invariant under (other) sorts of changes, for example, under the sorts of changes in selective regime that Beatty describes that would disrupt normal Mendelian segregation. It is also compatible with (S) failing to hold in other populations. What matters for the explanatory status of (S) with respect to population P is not whether it has exceptions in other populations or whether its holding in P is contingent on conditions that could have been otherwise, but rather that the condition on which it is contingent (whether Mendelian segregation continues to provide a selective advantage in P) is itself changed or disturbed by some range of interventions that produce the kinds of changes in fitness or initial genotype frequency which we invoke when we appeal to (S) to explain.

The status of (S) is in this respect is thus parallel to the status ascribed to (H) in the previous section. We argued there that the legitimacy of appealing to (H) to explain the behavior of some particular spring s is *not* undercut by the fact that there are *other* springs that do not conform to (H) or by the fact that there are circumstances in which (H) would fail to correctly describe the behavior of s. Similarly, even though it is true that whether (H) holds for some particular spring is contingent on the internal structure of the spring, we still can use (H) to explain why the spring exerts the restoring force it does as long as we are concerned with a range of extensions that do not alter the internal structure of the spring. For the same reason, even if we agree with Beatty that (S) is not a real law because there are evolutionary changes over which it fails to be invariant, we can still legitimately appeal to it to explain in models and circumstances of the sort described above.

8. Invariance and Spatio-Temporal Stability. I turn now to a more detailed assessment of Mitchell's approach. As I remarked above, Mitchell advocates a framework in which generalizations differ along several different dimensions, one of which is "degree of stability". She writes

> there is a difference between Mendel's laws and Galileo's law . . . but it is not the difference between a claim that could not have been otherwise (a "law") and a contingent claim (a "non-law"). What is required to represent the difference between these two laws is a framework in which to locate different degree of stability of the conditions on which the relation described is contingent. The conditions upon which the different laws rest may vary with respect to stability in either time or space or both. (2000, 252)

What is the connection between these ideas about "stability" and invariance? It seems to me that the notions are quite different in motivation and that neither is necessary nor sufficient for the other. Consider structure (4.1) in which *Y* and *Z* are joint effects of the common cause *X*. In this case the relationship between *Y* and *Z* is not invariant under (any) interventions on *X*, hence not invariant at all. However, if the common cause structure itself is stable in the sense that occurs repeatedly in many or most regions of space and time or in the sense that a single instance of *X* has common effects in many spatio-temporal regions, then the relationship between *Y* and *Z* will be highly stable in Mitchell's sense.

In fact, it is easy to think of examples having this sort of structure; many of them involve cosmological regularities. Consider:

(8.1) All regions of space exhibit a cosmic background radiation of 2.7 degrees Kelvin.

This generalization is highly stable in Mitchell's sense. However, the explanation for why different regions of space conform to this regularity is that the radiation present in these different regions are effects of a single common cause: the conditions that prevailed in the very early universe at the time of the so-called big bang. If we could somehow intervene to alter the microwave background radiation in some particular region of space, this would not be a way of altering the background radiation in other regions. In this sense, although (8.1) is stable in Mitchell's sense, it not invariant under interventions. Because it fails to be invariant under (any) interventions, (8.1) is not a plausible candidate for a law of nature. In keeping with the account of explanation I defended above, it also seems clear that one could not legitimately appeal to (8.1) to explain why some particular region of space exhibits the background radiation it does. Instead, the explanation for this is to be found in the conditions prevailing in the early universe.

Does a similar point hold for biological generalizations? In an interesting recent paper (1998), Kenneth Waters distinguishes between two sorts of biological generalizations. Generalizations about *distributions* describe "historically based contingencies" concerning the distribution of biological entities or characteristics. Examples include generalizations about the prevalence of various kinds of circulatory systems across taxa or the generalization that major arteries in many organisms contain a larger amount of elastin than do other blood vessels. Waters contrasts such generalizations with generalizations describing *causal regularities* — an example is the generalization that blood vessels containing a large amount of elastin will expand when the amount of fluid in them is increased (1998, 19). Although Waters does not explicitly endorse the manipulationist account of the content of causal generalizations that I advocate, his examples fit very naturally into this framework. Changing the amount of fluid in a blood vessel containing elastin is a way of manipulating whether it expands or contracts and it is because this generalization furnishes information relevant to manipulation that it is appropriate to think of it as a causal or explanatory generalization. By contrast, the generalization that elastin is present in larger amounts in arterial blood vessels or that it is present in the blood vessels of all or most vertebrates does not describe a relationship that is even potentially relevant to manipulation and control and hence is not a causal or explanatory generalization.

The relevance of this to Mitchell's discussion is as follows: generalizations describing distributions that claim that some biological characteristic is very widely or universally shared by all organisms may exhibit a great deal of stability in Mitchell's sense but this fact will not by itself show them to be causal or explanatory generalizations. Conversely, a generalization can be causal or explanatory in the sense that it describes a relationship that is exploitable in principle for manipulation and control even though this relationship holds for, or applies to, only a biological structure that is not shared by many organisms. For example, on my view, a generalization describing the response of a particular sort of neural circuit to stimulation can qualify as causal or explanatory even if that circuit is not widely conserved in other organisms.

This difference is in turn connected to another difference between my view and Mitchell's. As we have seen, Mitchell thinks in terms of a single "continuum of contingency" with (8.2) the law of the conservation of mass-energy at one end and a generalization like (2.1) "All the coins in Goodman's pocket are made of copper" at the other. This framework seems inevitable if the degree of stability or contingency of a generalization just has to do with the stability across space and time of the conditions on which it is dependent. The difference in stability in this sense between (8.2) and (2.1) is clearly a matter of degree: assuming that there was a period of time during which (2.1) was true, (2.1) itself and the conditions on which it depends were stable over this spatio-temporal interval, however limited its duration. By contrast, as explained above, while generalizations can differ in degree of invariance, some generalizations including (2.1) as well as (8.1) are not invariant at all. These differ in kind rather than in degree from invariant generalizations—they fall below the threshold for explanatory status. If our interest is in capturing the features that a generalization must possess if it is to figure in explanations and tell us about the results of interventions, the threshold/continuum model seems more appropriate. It isn't the case that generalizations like (2.1) and (8.1) are somewhat explanatory but less so than (8.2) or that they provide some information about the results of interventions but less than (8.2) does. Rather (2.1) and (8.1) seem not to be explanatory and tell us nothing about the results of interventions. Unlike Mitchell's model, my account captures this.

Not only can a generalization be stable in Mitchell's sense without being invariant, it can also be invariant under some non-trivial range of interventions and changes without being particularly stable in Mitchell's sense. The generalization (H) is (or may be imagined to be) a case in point. (H) may hold only for a very specialized sort of spring; the particular conditions (the fact that the spring has been constructed in a very specific way out of a specific sort of material) on which its holding depends may occur only very rarely, or within a very small corner of the universe. Nonetheless, when (H) does hold for some spring, it may be relatively invariant in the sense that it will continue to hold under a fairly wide range of interventions on its extension and under many changes in background conditions. If so, (H) will be relatively invariant but not particularly stable.

As a second illustration, consider two biological mechanisms involved in gene expression and regulation. One of these is highly conserved—it is found in many different species of animals that are widely distributed in space and time. The other is highly specific to a particular kind of animal. A generalization describing the behavior of the highly conserved mechanism will be more stable in Mitchell's sense than a generalization that describes the less conserved mechanism—it will rank higher on the continuum of stability. Nonetheless, both generalizations may be relatively invariant with respect to the behavior of the mechanisms they describe. On my view, the mere fact that a biological mechanism is highly conserved does not mean that the generalization describing its behavior is more lawful or invariant or that it provides a deeper or better understanding of its behavior than does a generalization describing a less highly conserved mechanism. As argued above, a parallel claim holds for Mendel's laws.

These differences between stability in Mitchell's sense and invariance as I conceive it are a reflection of the fact that stability is a non-modal notion: whether a generalization is stable depends on whether the conditions on which it is contingent are in fact stable across space and time. By contrast, invariance is a modal or counterfactual notion. Whether a generalization is invariant depends not on whether conditions that would disrupt it in fact occur, but rather on whether the generalization *would be* disrupted *if* various conditions (involving interventions) *were* to occur. Thus (8.1) fails to be invariant not because conditions that disrupt it do in fact occur but rather because if certain changes were to occur, they would disrupt it. Mitchell's account threatens to treat any de facto regularity as a law, although perhaps one with a very modest degree of stability.

9. Invariance and Resiliency. I turn now to some brief remarks on the notion of resiliency in the sense of Skyrms 1980, and Skyrms and Lambert 1995, and its relationship to invariance. Both Mitchell and several other commentators[10] regard these notions as closely similar. My aim in this section is to defend the view that they are different in important respects. Abstracting away from formal details, the resiliency of a proposition has to do with extent to which its subjective probability remains stable or unchanged (or changes only by some small amount) as one conditionalizes on other truth functional propositions in some family, all of which are consistent with the original proposition and its denial. Resiliency is thus a measure of degree of epistemic entrenchment or "resistance to belief change" or centrality to our web of belief (Skyrms and Lambert 1995, 139–141), in the sense that it indicates the extent to which an agent's degree of belief in a proposition would change under changes in her other beliefs. Laws are just generalizations that are relatively highly resilient (under conditionalization on some appropriate set of other beliefs): "the necessity of laws, like the necessity of causes, is resiliency." (1995, 145)

While resiliency is thus an epistemic or doxastic notion which has to do with the relationships among an agent's (or perhaps a scientific community's) beliefs, invariance is a non-epistemic or "objective" notion, the characterization of which has to do with the way the world is, rather than with anyone's beliefs. In particular, invariance has to do with the extent to which a generalization would continue to truly describe the behavior of some system (or the relationship described by the generalization would continue to hold) under changes that are actual physical changes in the system, rather than under changes in agent's beliefs or evidence. It follows that a generalization can be invariant even though no one knows that it is and it can be widely believed that a generalization is highly invariant when in fact it is not. Assigning a generalization a central role in one's web of belief and treating it as resilient does not make it invariant. Conversely, a generalization can be invariant even though, given an agent's other beliefs and the available evidence, it is highly non-resilient. Consider some newly conjectured candidate for a fundamental law of nature for

10. See especially Cooper (1998), who explicitly associates resiliency in Skyrms' sense with my notion of invariance. Skyrms has argued for a similar view in correspondence and conversation.

which there is at present only very weak evidence and which contradicts some apparently well-established scientific claims. Belief in this generalization at least at present will be relatively non-resilient but of course this is compatible with the generalization being in fact highly invariant.

While resiliency is thus relativized to a set of beliefs in a way in which invariance is not, this is not the only difference between the two notions, and resiliency is not just a subjective or doxastically relativized version of invariance. Presumably, the doxastically relativized counterpart to invariance is the notion of *believing* a generalization to be invariant, where to believe that a generalization is invariant is just to believe that it would continue to hold under some class of interventions. Like the notion of resiliency, believing a generalization to be invariant must be defined by reference to an agent's belief state. However, having a highly resilient belief in a generalization is not the same thing as believing it to be invariant. The reason for this is that, unlike invariance, the notion of resiliency assigns no special significance to stability of belief under changes in other beliefs that have to do with the occurrence of interventions. It is perfectly possible for an agent's degree of belief in some generalization to remain stable under changes in many of her other beliefs *B* but not under changes in her beliefs that an intervention has occurred. In this case, the agent's belief in the generalization is resilient with respect to the other beliefs *B* but the agent does not believe the generalization to be invariant. Conversely, an agent may believe that a generalization would continue to hold under some class of interventions—and hence believe that it is invariant—but would readily give up belief in the generalization if she were to acquire various other beliefs or if certain kinds of evidence were to become available.

The significance of this point can be brought out by means of some examples. Consider the generalization:

(9.1) No human beings live on other planets of the solar system.

There is a wide variety of different kinds of evidence for (9.1) and in the case of most of us, it is relatively epistemically entrenched—we would remain committed to it under many possible changes in other beliefs. Nonetheless, despite its resiliency, (9.1) is no law of nature. On my view, (9.1) is not a plausible candidate for a law of nature because there are many possible interventions (e.g., establishment of a colony on Mars) which would render (9.1) false. Moreover, it also seems clear that to have a resilient belief in (9.1) is not at all tantamount to *believing* that (9.1) is invariant or is a law of nature. While most people's degree of belief in (9.1) is relatively resilient, few regard it as a law of nature. The question of whether there are interventions that would render (9.1) false seems largely independent of the extent to which (9.1) is resilient. A similar remark holds for cosmological generalizations like (8.1). (8.1) is supported by a wide variety of evidence, is centrally located in our web of belief, and hence (according to Skyrms and Lambert) relatively resilient, but it does not follow that it is invariant or even that it is believed to be invariant.

Consider another example. My present degree of belief in General Relativity would change considerably if I became convinced that (9.2) most experts in the relevant scientific community in 2050 will regard this theory as false. My degree of belief in GR is not resilient under a change in my present beliefs to (9.2). Nonetheless this failure of resilience has nothing to do with the extent to which the equations of GR are invariant. Instead, the issue of invariance has to do with whether there are physical changes that might actually occur in nature which would disrupt those equations. The change consisting in my learning (9.2) will change my belief in GR but this change does not represent a physical intervention or change which creates a system in which the equations of GR no longer hold. Similarly, it does not follow from the fact that I would be prepared to give up my belief in GR if I were to become convinced of (9.2) (or if various other sorts of evidence were to become available) that I presently believe that the equations of GR are non-invariant or not laws of nature.

Before leaving the topic of resiliency, there is an additional point that is worth making. This is that many of the examples that Skyrms uses to illustrate the connection between resiliency, on the one hand, and lawfulness and causal necessity, on the other, can be very plausibly interpreted as instead illustrating the connection between invariance, lawfulness, and causation. For example, Skyrms (1980, 18) notes (following Max Planck [1922] 1960) that "all attempts to affect" the probability of decay of a uranium atom are unsuccessful. He connects this observation with the claim that generalization (P) describing this probability of decay is resilient and takes this in turn to capture the sense in which this generalization is (or is regarded by us as) lawful or necessary. However, it is at least equally natural to take Planck to be making a claim about the invariance of the generalization P: no matter what we or nature do, there are no physically possible changes that will affect the probability of decay. Planck's language and the particular illustration he offers (that changing the temperature of the atom will not affect the probability of decay) make it clear that he is talking about physical changes in the world, and not (or not just) about the stability of an observer's belief about the probability of decay under changes in the information or evidence available to him. My suggestion is that it is the invariance of the probability of decay under physical changes that leads us to regard P as lawful or necessary.

In my view, it was an extremely important insight on Skryms' part to recognize and articulate the connection between the lawfulness of a generalization and whether it remains stable as other conditions are changed.

However the intuitive force of this idea is better captured by thinking of stability in terms of invariance rather than resiliency. Resiliency is an important and valuable concept in its own right but its proper role is in capturing ideas having to do with doxastic entrenchment and stability of belief under additional information rather than in capturing notions like "law" and "physical necessity."

10. Conclusion. My aim in this essay has been to argue that invariance rather than lawfulness is the key feature that a biological generalization must possess if it is to figure in explanations. Invariance differs from mere de facto stability and also from resiliency. A generalization can be invariant even if it lacks many of the features standardly assigned to laws. Once the role of invariance in explanation is recognized, much of the motivation for shoe-horning explanatory biological generalizations into the category of laws of nature disappears.

REFERENCES

Beatty, J. (1995), "The Evolutionary Contingency Thesis", in J. Lennox and G. Wolters (eds.), *Concepts, Theories and Rationality in the Biological Sciences.* Pittsburgh: University of Pittsburgh Press, 45–81.

——— (1997), "Why Do Biologists Argue Like They Do?", *Philosophy of Science* 64 (Proceedings): S432-S443.

Brandon, R. (1997), "Does Biology Have Laws? The Experimental Evidence", *Philosophy of Science* 64 (Proceedings): S444-S457.

Cooper, G. (1998), "Generalizations in Ecology: A Philosophical Taxonomy", *Biology and Philosophy* 13: 555–586.

Hitchcock, C. and Woodward, J. (forthcoming), "Explanatory Generalizations, Deep and Shallow".

Mayr, E. (1982), *The Growth of Biological Thought: Diversity, Evolution and Inheritance.* Cambridge: Belknap Press.

Mitchell, S. (1997), "Pragmatic Laws", *Philosophy of Science* 64 (Proceedings): S468-S479.

——— (2000), "Dimensions of Scientific Law", *Philosophy of Science* 67: 242–265.

Planck, Max ([1922]1960), *A Survey of Physical Theory.* Reprint. Translated by R. Jones and D. H. Williams. Originally published as *Physikalische Rundblicke* (Leipzig: S. Hirzel). New York: Dover.

Salmon, W. (1989), *Four Decades of Scientific Explanation.* Minneapolis:University of Minnesota Press.

Skyrms, B. (1980), *Causal Necessity.* New Haven: Yale University Press.

——— and Lambert, K. (1995), "The Middle Ground: Resiliency and Laws in the Web of Belief", in F. Weinert (ed.), *Laws of Nature: Essays on the Philosophical, Scientific, and Historical Dimensions.* Berlin: Walter de Gruyter, 139–156.

Sober, E. (1997), "Two Outbreaks of Lawlessness in Recent Philosophy of Biology", *Philosophy of Science* 64 (Proceedings): S458-S467.

Waters, C. K. (1998), "Causal Regularities in the Biological World of Contingent Distributions", *Biology and Philosophy* 13: 5–36.

Weinberg, R. (1985), "The Molecules of Life", *Scientific American* 253(4): 48–57.

Woodward, J. (1997), "Explanation, Invariance and Intervention", *Philosophy of Science* 64 (Proceedings): S26-S41.

——— (2000), "Explanation and Invariance in the Special Sciences", *British Journal for the Philosophy of Science* 51: 197–254.

——— (forthcoming) "Causality and Manipulation".