

Reconceiving Eliminative Inference*

Patrick Forber^{†‡}

Eliminative reasoning seems to play an important role in the sciences, but should it be part of our best theory of science? Statistical evidence, prevalent across the sciences, causes problems for eliminative inference, supporting the view that probabilistic theories of confirmation provide a better framework for reasoning about evidence. Here I argue that deductive elimination has an important inferential role to play in science, one that is compatible with probabilistic approaches to evidence. Eliminative inferences help frame testing problems, an essential step that determines the context for evaluating statistical evidence. I illustrate this process with examples from molecular evolutionary biology.

1. Introduction. Scientists seldom consider all the possibilities. Successful theory testing requires restricting attention to a few promising alternatives, eliminating many possible hypotheses from consideration. Are there good rules of inference governing this process of elimination, or is this merely a case of opportunism in science?

I will argue that a process of elimination, call it *framing*, plays an important role in theory testing, a role neglected in current philosophical accounts. Reviewing recent accounts of eliminative inference shows that they all defend the standard picture, for they treat it as a distinctive pattern of reasoning that guides theory choice (sec. 2). These views have familiar problems that make eliminative inference controversial on epistemological grounds (sec. 3).

*Received November 2010; revised November 2010.

†To contact the author, please write to: Philosophy Department, Tufts University, Miner Hall, 14 Upper Campus Rd., Medford, MA 02155; e-mail: patrick.forber@tufts.edu.

‡Thanks to Peter Godfrey-Smith, Mark Richard, Elliott Sober, and Kyle Stanford for extensive comments on earlier drafts. Thanks also to audiences at the University of California, Irvine, Logic and Philosophy of Science Department and the Colorado State University Philosophy Department for useful feedback and to the anonymous referees for valuable criticism.

Philosophy of Science, 78 (April 2011) pp. 185–208. 0031-8248/2011/7802-0001\$10.00
Copyright 2011 by the Philosophy of Science Association. All rights reserved.

Framing relocates and repurposes eliminative inference (sec. 4). Instead of guiding theory choices, elimination helps frame testing problems, restricting focus to a few serious rivals. Framing thus aids the statistical evaluation of evidence and avoids some of the problems that face standard applications of eliminative inference. To support the philosophical account, I provide three examples from molecular evolution (sec. 5) and discuss eliminative inference from the Bayesian perspective (sec. 6). By relocating eliminative inference we preserve the view that elimination has a justifiable role to play in our best theory of science and sophisticate our account of how background theory guides testing procedures.

2. Chasing Holmes. The standard picture of eliminative inference, often called eliminative induction (EI), attempts to capture the pattern of reasoning made so perspicuous by Conan Doyle's (1890) fictional detective Sherlock Holmes, who counsels "that when you have eliminated the impossible, whatever remains, however improbable, must be the truth." Applied to science, EI thus guides *theory choice*. The idea that elimination works this way is commonplace, figuring prominently in the influential accounts of science offered by Mill and Popper. Dretske (1981) suggests that the process of eliminating relevant alternatives distinguishes knowledge from mere belief. It also has a strong appeal to scientists themselves. The elimination of alternative hypotheses through "crucial experiments" drives the process of *strong inference* (Platt 1964), an account of scientific method popular in ecology and evolutionary biology.

In the philosophical literature, Earman (1992), Kitcher (1993), and Norton (1995) defend the standard picture, identifying an important role for elimination in scientific reasoning. Yet there remains justifiable skepticism about the significance of EI. The contact between theory and data is, in practice, complex and rarely yields deductive falsification. Prevailing views take probabilistic theories of confirmation to provide better accounts of testing, and such accounts can accommodate the rare deductive inference. To motivate the need for a new account of elimination, let me begin by stating clearly the standard picture and discussing the problems that generate skepticism about the applicability of EI.

The standard picture treats EI as a two-step inferential pattern: (1) construct a space of possibilities and then (2) use observations to eliminate alternatives in that space. The goal of EI is theory choice: the process of elimination determines which hypothesis one should accept. Any account of EI must specify procedures for the construction of possibility space and for the elimination process.

All the accounts in the literature present a version of the standard picture, differing only in the explication of the two inferential steps. Earman and Norton marshal examples from the physical sciences to argue

for the importance of standard EI for guiding theory choices.¹ They explicate the construction step by appeal to local features of the community of practicing scientists, such as the nature of the explanatory domain and the sets of accepted auxiliary hypotheses (see, e.g., Earman 1992, 177). They explicate the elimination step in an inclusive way, allowing for both deductive and inductive ways of eliminating hypotheses, based on evidence or even theoretical principles.² Perhaps the only significant difference between their accounts concerns the epistemic goods EI can produce. Earman sees EI as augmenting epistemic progress within a Bayesian framework, whereas Norton argues that EI can also be a means of discovery. If all the hypotheses in the possibility space are eliminated, then this justifies formulating a novel alternative.

Kitcher turns instead to biology and provides a more complex explication of the two steps in standard EI.³ For the construction step, Kitcher identifies an algorithm for defining an exhaustive space of possibilities for explaining some regularity. This algorithm depends on a scientist's *individual practice*. An individual practice encompasses the complex set of commitments of a scientist, including a language, patterns of explanation, available techniques and instrumentation, and assignments of authority to colleagues, among other things (Kitcher 1993, 74–75, 221). Drawing on this set of resources, a scientist specifies all the possible explanatory dependencies. This constitutes the exhaustive possibility space. In contrast to Earman and Norton, Kitcher offers an individual-based rather than community-based explication of the construction step.

Kitcher also offers a different explication of the elimination step. For him, a set of representative observations drives purely deductive elimination of alternatives in the constructed space. Of course, Kitcher is aware of the complexities involved in bringing theories in contact with data. He appeals to the nuanced and reliable skills scientists have for making ob-

1. Their endorsement of the standard picture is explicit. Earman (1992, 167) designs his “sophisticated eliminativism” to fit within a Bayesian framework and requires (1) a finite partition of possibility space into observationally distinguishable members and (2) confrontations between theory and data. Norton (1995, 29) describes EI as a kind of argument with two kinds of premises: (1) premises that specify a possibility space and (2) premises that enable elimination.

2. Earman (1992, 178–81) claims that hypotheses can be “probabilistically eliminable,” a kind of Bayesian elimination that discards alternatives with very low credences. Norton (1995, 31) wants to allow elimination by “inductive inferences.” Both also want to countenance elimination by mathematical or theoretical principles. This weakened form of elimination is problematic, as I make clear below.

3. Kitcher's analysis of EI assumes a rich view about the nature of scientific practice, one I lack the space to treat in detail. I will focus on a few relevant points; for the complete treatment of his view of science, see Kitcher (1993).

servations to explain how they transform simple observations into the public data sets (1993, 222–23). These skills are demonstrably effective and help adjudicate disputes about the interpretation of observations (226–27). Provided that there are sufficiently many observations of the right kind, elimination occurs in a purely deductive Holmesian fashion (239).

Although the mechanics of Kitcher's account of EI are complicated, his view coheres with the standard picture.⁴ EI works via (1) the construction of a space of possibilities based on a scientist's individual practice and then (2) the use of skillfully obtained scientific observations to deductively eliminate those possibilities to inform a specific theory choice. In the ideal case the deductive inference leaves one hypothesis on the table and thus supports accepting that hypothesis as true.

On the standard picture, EI informs our theory choices; such inferences tell us what hypotheses we should accept based on what we take to be the space of alternative hypotheses and the conflicting evidence. These accounts are all chasing Holmes by giving plausible accounts of how such an intuitive process of reasoning functions in science. Yet there are problems with the standard picture that justify a reasonable skepticism about the viability of elimination. Reviewing these problems makes clear the need to sophisticate our account of how we use background knowledge.

3. Problems for Elimination. Both parts of EI, (1) construction of a possibility space and (2) elimination by evidence, face problems that underwrite skepticism toward EI. The first step can be confounded by uncooperative possibility spaces. Such spaces contain either too many hypotheses (an infinite number) or too few (they fail to include a relevant alternative). The second step can be confounded by testing holism or statistical evidence. Of these four problems, only two are problems specifically for elimination, and only one of those, the problem of statistical evidence, lacks a viable solution.

Consider the construction of a possibility space. Some inferential tasks, such as determining the exact value of a physical constant, have an infinite number of possible solutions. Eliminating one value, or even a large number of values, will not help. This problem has a clear solution. First, we should accept that there are some inferential problems in which standard EI may not work, such as the determination of physical constants. Otherwise, the infinite number problem can be overcome by imposing a finite

4. See Kitcher (1993, 237–42, 247–50); there he provides a kind of formalism for EI. He is also careful to include conditions that ensure a finite possibility space and defends the deductive approach against testing holism.

partition on possibility space. Indeed, this is the solution one finds in scientific practice.⁵

The more serious problem for the first step of EI concerns the possibility that theorists may fail to exhaust possibility space, what Stanford (2006) calls *the problem of unconceived alternatives*.⁶ If scientifically serious possibilities are left out of the constructed possibility space, then EI may produce misleading or mistaken theory choices. If Holmes leaves a suspect off his list, then his deduction should not be trusted and may very well lead him astray. While this is a serious problem, its scope is not restricted to elimination. Failure to exhaust possibility space presents a problem for any account of theory choice, for the accuracy of such choices depends on the available options.⁷ Thus, problems with the first part of EI do not underwrite targeted skepticism for eliminative inferences over and above general skepticism toward inductive methods.

The second part of EI, the elimination step, faces more severe problems. One concerns Duhemian *testing holism*. Testing holism undermines elimination because no hypothesis can make contact with empirical observations in isolation. A whole body of background theory is necessary to generate a confrontation between theory and data. Since such background theory encompasses a vast network, it seems that any conflict between theory and data can be reconciled by changing some element in the background. We have no justifiable reason for eliminating the hypothesis instead of the auxiliaries we must use to generate predictions. There are some sophisticated responses to testing holism.⁸ Yet, as with Stanford's problem, testing holism has implications for any strategy of assessing evidential relevance; it does not underwrite targeted skepticism about elimination.

There is a problem with the second inferential step that does support skepticism about elimination over and above other inductive methods,

5. Earman (1992, 175) argues that the EIs in gravitational research rely on such a finite partition to eliminate infinite sets of possibilities. Kitcher (1993, 248) claims that a finite partition can finesse this problem in biology.

6. Stanford's argument draws on the historical record of science to support an instrumentalist view of science. Yet he agrees with Earman, Norton, and Kitcher that EI plays a crucial role in science (Stanford 2006, 29–30).

7. Stanford argues that his problem undermines Bayesian inferences as well; see Forber (2008) for discussion.

8. Let me name two promising ones. Kitcher (1993, 249–50) argues that there are differential costs associated with different "escape routes" from a conflict between theory and data and that often the low-cost escape route involves eliminating the hypothesis under suspicion. Sober (1990, 1999, 2008) argues that by using contrastive testing, one can isolate parts of a theory responsible for specific predictions and thus provide directed tests of individual hypotheses.

and that is the challenge of using statistical evidence to drive EI. In scientific practice, evidence rarely provides deductive falsification but instead favors some hypothesis over others. Most, if not all, evidence in evolutionary biology favors one hypothesis over others, hence showing that the rival hypotheses are unlikely rather than deductively eliminating them. The evidence tends to be *statistical*. Many current theories of confirmation treat this favoring relation as the core of any theory of evidence and aim to explicate it using probability theory (Maher 2004; Joyce 2007). Can EI work with favoring evidence? If no, EI captures very little of scientific practice. If yes, EI requires rules that turn statistical evidence into eliminative evidence. Within probabilistic confirmation theories, such rules would be thresholds that identify how much evidence against must accumulate to eliminate a hypothesis.

The response to the problem of statistical evidence grants that the exact thresholds may be elusive, unavailable, or vague but insists that we can judge most cases unproblematically. Consider the following claim by Earman (1992, 164): “Nevertheless, such a hypothesis or theory may be probabilistically eliminable in that accumulating evidence can drive its probability so low that it is no longer worth considering.” Stanford (2006, 29) makes a similar claim: “Of course, eliminative inferences do not always exclude alternatives by rendering them impossible: they can instead simply show that some possibilities are much more likely than others.” These intuitive replies opt for a pragmatic threshold solution to the problem. However, such a solution obscures the difference between statistical and eliminative evidence. Evidential support is an epistemic matter, and so we need an epistemic justification of the rules or thresholds for elimination or at least some account of these notions of probabilistic eliminability and strong favoring.⁹

To further motivate this problem, consider the following argument from Sober (2005, 2008) against a rule of inference he calls *probabilistic modus tollens* (PMT). PMT is an extension of *modus tollens*, which applies when some hypothesis makes probabilistic rather than certain predictions. PMT inferences work like this: if some hypothesis H predicts O with low probability (the likelihood $P(O|H)$ is very low) and one observes O , then (probably) $\neg H$. PMT is not a viable rule of inference because many reasonable hypotheses make predictions that are very improbable. For instance, the likelihood of a mass extinction given a meteor impact may be

9. Notice that Kitcher ostensibly avoids this problem, but only by appealing to an account of scientific observation that can produce eliminative evidence. This process of observation must be inductive in character and so must explain how statistical evidence is converted into eliminative evidence. Appealing to the overall reliability of observation techniques is unconvincing.

exceedingly low, yet observing a mass extinction can be evidence for the impact hypothesis when alternative hypotheses—slow climate change, say—specify a lower likelihood for the observations. PMT fails because the absolute value of $P(O|H)$ is meaningless; it is the comparison of likelihoods that matters for assessing evidential support.¹⁰ Sober uses the failure of PMT to justify skepticism that there can be any general normative account of the rules for elimination with statistical evidence.

An old epistemological problem also supports this skepticism: Kyburg's (1961) lottery paradox (Sober 1993, 2008). The so-called lottery paradox reveals a tension between probabilistic credences and the acceptance (or rejection) of hypotheses. Suppose that we adopt an intuitive rule for elimination: reject H when $P(H) < y$, where y is the threshold where hypotheses are "probabilistically eliminable." For any elimination threshold y we can construct a fair lottery with one randomly chosen winning ticket out of x total tickets, where $x > 1/y$. For any ticket z , $P(z \text{ wins}) < y$; therefore, we reject the hypothesis ticket z wins. If we reject the hypothesis that ticket z wins for all x tickets, then we should accept the conjunction that no ticket will win. Yet this contradicts our belief that the lottery is fair and one ticket will win. There are certainly replies to the lottery paradox (Sorensen 2006). One reply appeals to utilities to explain away the paradox. Given the cost of lottery tickets, there is a point where investing in the lottery is not worth it. We make utility-based judgments when we claim to reject the hypothesis that ticket z will win, but that does not entail that we accept the conjunction that no ticket will win. The appeal to utilities works well for EI in science. Pursuing theories is not a cost-free enterprise, and thus taking low-probability alternatives off the table often yields the best epistemic return for our investment. Yet such a reply dodges rather than defeats the worry Sober raises. EI becomes a matter for decision theory rather than epistemology. If EI captures a general and justifiable epistemic pattern of inference in science, then we should strive to provide a general account of the rules for eliminating hypotheses based on statistical evidence. Without such an account the applicable scope of standard EI to scientific practice would be severely restricted.

To put the point another way, standard eliminative induction is not truly eliminative. In the cases in which we have statistical evidence, the

10. Sober also provides a simple counterexample based on tossing a fair coin. Suppose that one tosses a fair coin a large number of times, say a million. Given that the coin is fair— $P(\text{heads}) = P(\text{tails}) = 0.5$ —the probability of any exact sequence of heads and tails occurring is one-half to the millionth power; thus the likelihood of observing a specific sequence is exceedingly low. This likelihood obtains for each of all the possible sequences. One of those exact sequences must occur. According to PMT, from the observation of that exact sequence we should infer that (probably) the coin is not fair, which is absurd.

elimination step involves a pragmatic judgment that, all things considered, certain hypotheses should not be accepted or pursued. But EI provides no additional epistemic justification for theory choices beyond evidential favoring modeled by current confirmation theories. Instead, it is a concoction of pragmatic and epistemic criteria designed to provide a resolution to theory choices in a world of finite scientific resources. While I am sympathetic to this aim, there is a better way to divide the inferential burden between pragmatic and epistemic criteria.

4. Relocating Elimination. Let me summarize the state of play. Standard EI involves (1) constructing a possibility space and (2) using evidence to eliminate hypotheses in that space, ideally leaving a last one standing. Each inferential step faces problems. We may fail to articulate an exhaustive space of possibilities or to provide a finite partition of that space. Holism may undermine our ability to eliminate specific hypotheses, and statistical evidence may not permit elimination at all. The last problem presents a significant challenge because much of the evidence in science is statistical. Furthermore, the problem of statistical evidence justifies a targeted skepticism of EI beyond general inductive skepticism. I do not see a clear epistemic solution to the problem of statistical evidence, and the pragmatic threshold solution does not work well for theory choices. Yet there is a way to relocate eliminative inference that avoids the problem but still captures eliminative inferential patterns found in good scientific practice. Here I will propose and defend this relocation of eliminative inference. In the next section I will provide three examples to demonstrate that this new proposal describes successful science.

Perhaps eliminative inferences do not directly make theory choices but establish the boundaries for such choices. In Holmesian terms, instead of identifying whodunit, elimination determines who should (or should not) count as a suspect for investigation. On this approach, elimination constrains possibility space by ruling out those potential rival hypotheses that conflict with background knowledge. Evaluation of the focal evidence is done by using any of a variety of statistical tools to assess favoring relations among the remaining rivals. The statistical evaluation provides an assessment of epistemic support unencumbered by the stringent requirements for (pure or probabilistic) elimination. On the new proposal, eliminative inference plays a role in the first inferential step, whereas familiar confirmation-theoretic analyses suffice to represent the evaluation of the evidence. Call this application of eliminative inference *framing*.

In the context of evolutionary biology, framing involves the deductive elimination of possible evolutionary hypotheses that are incompatible with the background biology of a target system. Evolutionary theory provides the resources to construct a set of global alternatives for approaching the

evolutionary history of any system.¹¹ Yet not all of these possible hypotheses will work for a given target system. They may make substantive biological assumptions that do not cohere with what we know about that system. Incorporating the background biology of the target system, through the process of framing, eliminates some alternatives and yields a set of local rivals for the testing problem at hand. This is the set of viable evolutionary hypotheses given the background biological information. Once framing provides a set of rivals, statistical tools help determine which rival the data favor. Thus the overall chain of evidential reasoning has both an eliminative element and a statistical element.

Framing inferences require a distinction between *test focus* and *test background*. The focus of a particular test is the evolutionary outcome we seek to explain, say, the evolution of beak size in Galapagos finches. The focal data may include samples of beak sizes over many generations and the available food resources of specific finch populations. The background for this test would be information about finch beak development, phylogeny, demography, mating habits, and other aspects of their biology and ecology that biologists use to frame their evolutionary hypotheses about natural selection, drift, and constraint. For the purpose of one focal testing problem, biologists make a defeasible assumption to take the background biological information as certain. This information is not under direct investigation and simply is not in question given the test focus. This defeasible assumption permits the use of pure *modus tollens*, rather than the problematic probabilistic kind, for elimination.

Framing differs from the standard picture in three respects. First and foremost, it makes essential use of the distinction between focus and background. The decision to accept background biological claims as certain permits true deductive elimination. Second, it locates elimination in a different place in the inferential process: elimination is part of the construction of possibility space rather than the evaluation of evidence. Third, it has a different goal. Framing inferences make possible the effective statistical evaluation of data, whereas standard EIs make theory choices.

Of course, any biologist worth her salt will draw on background information about her study system to frame tests of her target hypotheses. Biologists can focus on only one part of the complex biological web at a time, and so they assume that background information about the rest of the web, usually provided by studies done by their colleagues and predecessors, is reliable. The status of background biological claims should, and usually does, have at least some sort of evidential support. Using eliminative inference to capture this commonplace feature of science

11. See Forber (2010) for further exploration of how-possibly explanations and their connection to confirmation.

thus faces a number of skeptical objections. Let me respond to one primary objection here and raise two others that I will discuss later.

The most pressing objection to this proposal concerns its normative force: what is the justification for taking the background information as certain? To put the objection more forcefully, simply deciding to take the background information as certain seems just as problematic as the pragmatic thresholds that are part of standard EI. Furthermore, without any criteria for determining what counts as background, this decision appears dangerously arbitrary: anything a biologist wants to count as certain for framing a test is relegated to the background. The justification for the certainty assumption is primarily pragmatic. Treating background information as certain permits the conditional assessment of the evidence without background uncertainties clouding the evaluation of the data. If some biological claims count as background for a specific testing problem and there is a good account of how to distinguish focus from background, then scientists are entitled to this pragmatic solution. In effect, we have a way to impose context-sensitive thresholds for accepting the background claims as certain. As long as there is some support for the background claims, we can take them as certain to frame a testing problem. If we hold science to a higher standard, claiming that auxiliary hypotheses and background claims must be true or strongly supported, we risk retarding the epistemic progress of the enterprise. We also close off any possibility of explaining the origin of such an enterprise, for the lack of support for background claims would undermine any epistemic support marshaled for any hypothesis. Allowing for liberal standards for framing arguments permits science to move forward on certain problems, however tentatively. This is not to say that anything goes. Indeed, eliminative framing arguments are closely scrutinized and are the source of many scientific controversies over what the evidence says.

This sort of pragmatic justification works for framing but not for standard eliminative induction. One reason is that the framing inferences isolate pragmatic considerations from the epistemic assessment of the focal evidence. This division of inferential labor respects the empiricist desideratum that the evaluation of evidence should be purely epistemic, or as epistemic as possible. Explicitly incorporating framing provides a picture of science in which evidential reasoning produces conclusions of the following form: if the framing argument holds, then the data support H_1 over the rivals H_2, \dots, H_n . The support relation is purely epistemic but circumscribed to a specific testing problem framed by a mix of pragmatic and epistemic concerns. Another reason is that the pragmatic solution for eliminative framing inferences avoids the need for general threshold rules that apply across all inferential contexts. The thresholds for specific eliminative inferences depend instead on the structure of the testing prob-

lem and the support for the background claims. These will vary considerably from case to case. Framing incorporates this context sensitivity.

In contrast, standard EI turns evidential reasoning itself into a decision problem. Epistemic support is combined with a general pragmatic solution to the problem of statistical evidence (probabilistic eliminability, strong favoring) to produce theory choices. But theory choices depend on more than just epistemic support. The cognitive and material resources at our disposal and the costs of accepting or pursuing particular theories should matter for these choices. Imposing a general credence threshold idealizes these factors away. Also, the plausibility of an eliminative inference depends on the details of the particular inferential context; a blanket threshold rule for elimination lacks sensitivity to these details. Hence, standard EI, by relying on a general pragmatic justification for elimination, obscures the complexity of theory choices. More important, when a threshold rule to make an eliminative theory choice is imposed (i.e., the evidence rules out those hypotheses so accept this one), information about the exact degree of support is lost. While this strategy can accelerate epistemic progress in science when applied to background claims, it should not interfere with the assessment of epistemic support for the focal hypotheses within a testing problem. Framing inferences thus relocate the pragmatic elements to permit a clear, yet conditional, epistemic evaluation of the data. In short, framing a test is pragmatic; assessing the evidential support is epistemic.

Giving an account of how scientists should determine what claims are background claims, and therefore what can fairly be used in framing inferences, is more difficult. In general, the structure of the testing problem will provide the salient criteria, but such criteria depend on the local features of inquiry into particular domains and will vary across the sciences. In the case of evolutionary biology, tests of adaptive hypotheses can treat phylogenetic relationships as background. One reason for this is that phylogenetic inquiry controls for selection in the evolutionary process to reconstruct ancestral relationships (Felsenstein 1985). Another is that many tests of adaptive hypotheses require background assumptions about common ancestry (Sober and Orzack 2003; Sober 2008). These are the sorts of concerns that provide the normative criteria for what counts as background information for a testing problem, and the examples below provide illustrations. Ultimately, the normative standards for assessing framing arguments depend on methodological details internal to particular sciences, hence the importance of discussing detailed examples.

Before I move to the examples, let me raise two more objections. The replies to these objections emerge from the discussion in the next two sections. There is a clear Bayesian objection: why invoke eliminative inference at all when a confirmation-theoretic model of how individuals

update using multiple sources of information should suffice? Briefly, using eliminative inference to capture framing provides a more accurate picture of science. I return to discuss the Bayesian perspective in more detail below (sec. 6).

There is also the objection that the current proposal is unnecessary: why invoke a focus/background distinction to augment eliminative inference when the standard picture should suffice? Distinguishing between test focus and test background and assuming that the background is certain have two virtues the standard picture lacks. The epistemic support for background information used for elimination in the framing process need not, and often does not, meet the stringent conditions for probabilistic eliminability or strong favoring. Alternative phylogenies for the finches (say) need not be eliminated or have extremely low credences for biologists to use one when framing the testing problem for the evolution of beak size. Also, framing inferences make clear how evidential assessment is context sensitive. The chain of eliminative reasoning applies only to the focal testing problem and can be (indeed, often is) reevaluated when the background information about a target system gets updated or revised. In the finch case, recent research into beak development may bear out basic assumptions or require refinement of the evolutionary hypotheses (see, e.g., Abzhanov et al. 2004, 2006). One of the cases below, the McDonald-Kreitman (MK) test, illustrates this dynamic in more detail.

5. Three Examples. Descriptive adequacy is an integral part of my argument for relocating eliminative inference. Also, analysis of real examples provides traction on the normative standards for evaluating framing arguments. Here, I support these claims. As illustrations of the eliminative process of framing, consider three examples from molecular evolutionary biology. The first comes from recent work by Fry and colleagues on the evolution of snake venom. The second looks at a textbook case, convergent evolution of adaptive lysozymes for foregut fermentation in langur monkeys and cows. The third example is the original MK test for positive selection on the *Adh* gene in *Drosophila*. Each study aims to give evidence for a specific evolutionary hypothesis. The chain of reasoning includes eliminative framing, for the biologists explicitly argue for certain biological assumptions about their target system, and these assumptions eliminate potential rival hypotheses. I will discuss each case in turn.

Recent work on the molecular evolution of snake venom provides evidence for an overall evolutionary picture in which some venoms evolved early as adaptations and more specialized venoms evolved later, after the diversification of the snake lineage (Fry, Lumsden, et al. 2003; Fry, Wüster, Kini, et al. 2003; Fry, Wüster, Ramjan, et al. 2003; Fry and Wüster 2004; Fry 2005; Fry et al. 2006, 2008). The evolution of snake venom is an




Figure 1. Phylogeny with evolutionary origins of several different toxin families marked on the tree. Dashed marks indicate equivocal evidence, and only a subset of known toxin families are included. (Figure adapted from Fry and Wüster [2004, 882] with permission.) Sequence comparisons provide strong evidence that several toxins in the venom arsenal evolved early in the snake clade (homologies), before the radiation of different lineages. The phylogenetic relationships are based on independent evidence from nuclear and mitochondrial DNA (Vidal and Hedges 2002).

interesting case because venom plays a specialized ecological role critical to survival. Some lineages (the vipers and elapids; see fig. 1) have evolved specialized front-fanged delivery systems and larger venom glands to increase the effectiveness of venom (see Vonk et al. 2008). These two lineages include the more fearsome venomous creatures, such as the American diamondback rattlesnake, the Australian taipan, and the African cobra. Given that venom currently plays such a central ecological role in these two snake lineages, when did the venom arsenal evolve?

To frame their testing problem, Fry and coworkers make three major assumptions based on the background biology of snakes. First, they assume that the venom arsenal evolves primarily by natural selection. Analysis reveals significant variation in venom composition. Snake venom is not a single chemical but a cocktail of proteins that shows high diversity across lineages (Fry, Wüster, Ramjan, et al. 2003). Certain cocktails affect prey differently. One set of toxins may be very effective against birds but not at all effective against amphibians or mammals. Furthermore, the existence of complex, delicate, and costly physiological mechanisms for the production and delivery of venom suggests that the traits are adaptations for prey capture or defense. This background information does not make adaptation a certainty, for a thorough functional analysis and evidence about the historical ecology are necessary to test the adaptive hypothesis (Fry and Wüster 2004, 880), but it does help justify using the claim as part of their framing argument.

This aspect of the framing argument for venom evolution raises a gen-

eral point about the justification of framing inferences. Assumptions about natural selection are associated with adaptationism and within an adaptationist research program require little, if any, independent justification. Relying solely on methodological principles to justify framing inferences is problematic. An assumption without any empirical support jeopardizes the entire case, especially when such assumptions eliminate rivals that we have some empirical reasons to consider. In the case of adaptation, we know that genetics, developmental constraints, and fluctuating environments can confound natural selection, and so we cannot rule out non-adaptive hypotheses without at least some evidence that these factors are not significant. In the venom case, Fry and coworkers have the sort of support they need to assume adaptation, and the biologists are also careful to point out that this is an assumption they are not testing. Not all methodological principles are suspect; some do have a role in science. Their viability will depend on the overall support we have for endorsing the background framework.¹²

Second, in part because of the redundancy of the cocktail and in part because of the overall physiology of snakes, the biologists take venom to be highly modular. That is, changes in venom composition can occur without major downstream effects on the development of the individual. Also because of the redundancy of the arsenal, venom toxins can evolve rapidly. Knocking out one component of the venom cocktail usually has relatively little adverse effect since other toxins still function normally.¹³ Third, they assume a resolved phylogeny for the snake clade based on different data from a separate study (Vidal and Hedges 2002). The phylogeny provides the necessary backdrop for testing hypotheses about ancestral character states, and this assumption is not tested by any of the venom sequence data. Instead, the focal testing problem involves using molecular data to place evolutionary events on the existing tree.

The framing inferences eliminate two possible evolutionary scenarios. The modularity of venom eliminates the strong selective constraint hypothesis. In molecular evolution, similarity across aligned protein or gene

12. Also, we should expect some diversity on this point since mathematical assumptions play an important role in physics (Norton 1995). Thanks to an anonymous referee for drawing my attention to the role of methodological principles.

13. More precisely, Fry, Wüster, Kini, et al. (2003) argue that the sequence data provide evidence for the “birth-and-death” model of gene family evolution, much like major histocompatibility complex evolution in vertebrates (Nei, Gu, and Sitnikova 1997). According to this model, the toxin family evolves rapidly by many gene duplication events. Most of these duplications become nonfunctional as a result of mutation and are evident in the genome as pseudogenes (gene death). Some evolve into a new functional version with different toxic effects on prey, thus increasing the family size of venoms (gene birth).

sequences can often be explained by selective (functional) constraint. For example, protein sequences of cytochrome C, an enzyme crucial to cell metabolism, show a high degree of similarity across the biological world, from yeast to humans (Margoliash and Smith 1965). This similarity is due to the strong functional constraints that make most potential amino acid substitutions deleterious. Also, many detected substitutions in cytochrome C are functionally similar and therefore effectively neutral (King and Jukes 1969). The assumed phylogeny rules out a late origin homology hypothesis. Another common way to explain sequence similarity posits a close phylogenetic relationship between species. On this hypothesis the gene is conjectured to belong to a recent common ancestor, and the short divergence time explains similarity. Many venom compounds under investigation belong to vipers and elapids, two lineages that diverged long ago (see fig. 1). The framing process leaves two rival hypotheses for the shared venom compounds in the arsenal: one early origin of the venom before the viper-elapid split (early homology) versus two independent late origins in each lineage (convergent evolution).

To test early homology versus convergent evolution for snake venom, Fry and coworkers compare protein sequences of various venom toxins across snake lineages and investigate whether venom can be found in other extant lineages. Evolutionary distances—a measure of how long two lineages have evolved independently—can be estimated by examining the differences across aligned protein or gene sequences. The evidence for their venom origin hypotheses rests primarily on sequence data. An example is the sequence comparison of one particular toxin family in the venom cocktail, the family of “three-finger toxins” (3FTxs). The 3FTx comparison provides evidence for an early origin and the rapid evolution of the toxin using two different methods of phylogenetic reconstruction (maximum parsimony and neighbor-joining methods; Fry, Wüster, Kini, et al. 2003, 117). They put the point of origin right after the divergence of vipers, the oldest snake lineage, and the rest of the snakes. As further support for the early origin of 3FTx, they isolate venom in the Colubrinae lineage (from a species presumed nonvenomous because of a lack of specialized delivery systems) that had the same conserved core 3FTx sequence (Fry, Lumsden, et al. 2003, 450–51). Isolating a conserved 3FTx sequence from this lineage shows that the 3FTx family originated early before much diversification. The more comprehensive studies compare sequences of multiple toxins (fig. 1). The sequence data provide evidence that many components of the venom arsenal evolved once early and the specialized front-fanged delivery system, found only in a few lineages, evolved multiple times later (Fry and Wüster 2004, 870).

The studies make comparisons across different snake lineages to test the competing hypotheses of homology versus convergent evolution for

several snake venoms. Also, they use several different statistical tools for reconstructing ancestral character states from extant sequence data. This is a common strategy in phylogenetic reconstruction and shows that there is more to statistical inference than just using the “right” statistical tool. The framing inferences constrain possibility space. Within the space of local rivals the different statistical tools converge, clearly favoring one rival over another.

A textbook study on molecular adaptation, the convergent evolution of lysozymes in cows and langur monkeys (Graur and Li 2000, 121–23), provides a second example of using eliminative inference to frame a testing problem. Among placental mammals these two lineages have evolved foregut fermentation that involves the use of a specialized form of lysozyme to aid digestion. Given that two separate lineages have evolved a molecular trait for the same function, there are three alternatives for explaining the evolution of these specialized lysozymes: both lineages share a common ancestor and the trait evolved once in the ancestral lineage (homology); the lysozymes and their corresponding use have evolved independently in each lineage because of positive selection for the same function (convergent evolution); and cows and langurs have coopted the use of general lysozymes, found in all placental mammals, without any intervening selection for a specialized protein (exaptation with neutral evolution). Notice that both of the first two hypotheses posit positive selection for the new specialized function, whereas the third posits no new positive selection and allows only for accumulation of neutral molecular changes. Again, the assumed phylogeny plays an important role in framing, for it eliminates the first alternative. Cows and langurs are distantly related. In addition, many intervening lineages more closely related to each species lack foregut fermentation and the specialized lysozyme, another point inconsistent with the expectations of the homology hypothesis.

In the original study (Stewart, Schilling, and Wilson 1987; Stewart and Wilson 1987), Stewart and coworkers test convergent evolution against exaptation plus neutral evolution by sequencing lysozyme proteins from several lineages. They determined that the specialized lysozymes in cows and langurs share four convergent (identical) and three parallel (functionally similar) amino acid substitutions arising independently in each lineage that affect the ability of lysozymes to conduct the foregut fermentation. The evidence thus favors convergent evolution by positive selection over exaptation with no new adaptive evolution. The lysozyme case combines phylogenetic and protein sequence data to make a convincing case for positive selection producing convergent evolution in cows

and langurs.¹⁴ Both the eliminative inference used to frame the test and the statistical evaluation of protein sequence data are necessary to identify the evidential support for positive selection.

A final example of framing is the classic test for positive selection in *Drosophila*. The original study by McDonald and Kreitman (1991) compares alcohol dehydrogenase (*Adh*) gene sequence data from three species of fruit flies to detect selection at the molecular level. They aim to contrast two rivals: the positive selection hypothesis versus the neutral evolution hypothesis. The selection hypothesis claims that positive directional selection for variants of *Adh* is responsible for certain amino acid differences between the species. The neutral evolution hypothesis predicts that mutations generally have no effect on evolutionary fitness, and thus variation accumulates by genetic drift, not selection. To effectively test selection against neutral evolution, McDonald and Kreitman use an eliminative inference that rules out one evolutionary hypothesis that would confound their statistical analysis. To isolate their framing inference I need to first present their test and their results.

McDonald and Kreitman construct the test by classifying possible changes in the gene sequence. Owing to the degeneracy of the genetic code, some DNA changes are *silent*; they do not change the amino acid sequence of the protein. Those nucleotide changes that do change the protein are *replacement* changes. By assumption, molecular evolution of silent changes fits the neutral model. The study aims to test whether the evolution of replacement changes occurs by selection or drift. Comparing the number of silent versus replacement differences across the sampled gene sequences provides an indication of evolutionary rates (Kreitman 2000, 546). On the neutral hypothesis we expect both silent and replacement sites to show the same evolutionary rate since both kinds of sites evolve by genetic drift, whereas we expect different rates of evolution on the selection hypothesis. Positive selection for adaptive changes in the protein should accelerate the rate of replacement evolution.¹⁵ The data indicate a faster rate for molecular evolution at replacement sites, and this clearly supports the positive selection hypothesis over the neutral rival (McDonald and Kreitman 1991, 654).

McDonald and Kreitman (1991, 653) note that an alternative hypothesis could explain their data and the faster rate of replacement evolution. If ancestral *Drosophila* populations went through a population bottleneck

14. Two studies provide further support for convergent adaptive evolution of lysozymes: Kornegay, Schilling, and Wilson (1994) isolated a similar specialized protein in a species of leaf-eating bird, and Messier and Stewart (1997) have resolved specific episodes of adaptive evolution in colobine monkeys.

15. See Eyre-Walker (2006) for a derivation of the MK test.

followed by rapid expansion, then these dynamics would produce an artifactual signature of selection in the extant gene sequence data. They eliminate this hypothesis because it requires sensitive demographic fluctuations to occur at precisely the right time in *Drosophila* history. They prefer the assumption that the effective population size has remained large throughout recent evolutionary history, and this assumption does have some support (Kreitman 1983). Thus, they explicitly make an eliminative inference to frame their test. Moreover, the elimination is crucial to the evidential support for positive selection since the data would fail to adjudicate between the bottleneck and selection hypotheses.

There has been some subsequent evaluation of McDonald and Kreitman's demographic assumption and the eliminative inference they make. In particular, Eyre-Walker (2002) explores the bottleneck alternative in greater detail. He provides a model that identifies the conditions under which fluctuations in effective population size can produce artifactual signatures for positive selection. As it happens, a fairly wide range of demographic conditions can produce artifactual evidence for selection as well as synonymous codon bias. This shows just how thin the empirical support for eliminative inferences can be. It also illustrates how the standards for framing arguments can change over the course of scientific inquiry. Now independent support for the demographic assumptions becomes more important for applying the MK test (Kreitman 2000). With respect to the original test on *Drosophila*, Eyre-Walker (2002, 2023–24) concludes, after considering several sources of information on demographic conditions, that the evidence in favor of positive selection still stands. Background biological information continues to support McDonald and Kreitman's eliminative inference.

6. Through the Bayesian Lens. Distinguishing test focus from test background, and the defeasible assumption to take the background biological information as certain, sets framing apart from the standard picture. On the new proposal, evidential reasoning has an eliminative element that helps constrain the local testing set and a statistical element that involves assessing how data favor one rival over other hypotheses in the testing set. In principle, we can represent the second element, the contrastive evaluation of data, with a variety of statistical tools, Bayesian or otherwise. Which statistical tool provides the best representation of confirmation is a thorny issue, and the arguments in this paper are designed to be neutral on this point. Yet given the prevalence and flexibility of Bayesianism, it is worth exploring whether the Bayesian epistemological framework can capture the eliminative framing inferences and what additional commitments this would require. This line of inquiry also provides

further support for differentiating between new (framing) and old (standard EI) applications of eliminative inference.

Earman (1992) argues that a Bayesian framework can accommodate standard EI.¹⁶ This is right. The framework can easily capture deductive elimination of hypotheses by evidence. If the likelihoods of specific observations given certain hypotheses take the value of zero and there is a finite set of such hypotheses in the possibility space, then Bayes's rule for updating credences handles elimination: data that have a zero likelihood ($P(O|H) = 0$) eliminate that rival if they are observed. However, this kind of deductive falsification is uncommon in scientific practice. To extend the scope of EI, Earman gives elimination a probabilistic gloss. Bayesian elimination occurs when evidence accumulates and drives the credence in some hypothesis extremely low. This kind of elimination requires a commitment to some threshold rule and so runs into Sober's worry that there can be no epistemic justification for such a threshold (sec. 3).

Hawthorne (1993) goes further than Earman and argues that Bayesian induction simply is eliminative induction. He offers a more detailed formal framework and uses various convergence theorems to explicate a Bayesian notion of "relative refutation" by probabilistic evidence. Disconfirming evidence pushes the credences of rivals very low and thus pushes the credence of (what we take to be) the true hypothesis very high. Hawthorne is right to identify a point of common ground between Bayesian and eliminative induction—evidence against rivals is evidence for my target hypothesis—but the notions of refutation and elimination get a similar probabilistic treatment in his framework and so are not the same as their deductive counterparts. Both Hawthorne and Earman must contend with the threshold problem to make EI work. This problem lacks a solution, and none is forthcoming, but their efforts aptly demonstrate the power and flexibility of Bayesian models of scientific inference.

Owing to this incredible flexibility, I do not doubt that such a model can capture framing as well. Yet fitting this application of eliminative inference into a Bayesian model requires conceptual resources beyond updating to adequately represent the inference process found in scientific practice. In particular, rather than threshold rules, a Bayesian model of eliminative framing requires a way of changing credence distributions for specific testing problems in response to the framing arguments. If such a transformation is driven purely by evidence, for background biological claims or hypotheses under test, then no resources external to the basic Bayesian model would be necessary. However, the transformation is not

16. Earman does regard certain elements of EI as non-Bayesian, specifically those that involve the exploration and partition of possibility space and certain eliminations of implicit alternatives.

driven purely by data but by distinguishing the test focus from the test background and assuming that background claims are certain. This assumption effectively raises the credences for background claims to one in order to enable their use in (deductive) eliminative inferences.

Evidence for the background biological claims raises their credences in the usual Bayesian way, but this evidence almost never suffices to raise the credence to the certainty required to deploy true elimination, as the examples illustrate. Eliminative framing inferences, in turn, reduce the credences of some possible hypotheses about the focal phenomenon to zero. Those hypotheses that are incompatible with the assumed certainty of background biological claims are thus eliminated from consideration. This permits scientists to evaluate statistical data with a constrained set of local rivals. By focusing on one element of the biological world and reducing much of the myriad complexities and empirical connections to background assumptions, biologists are able to make good epistemic progress. Such progress requires setting a context for investigation—distinguishing focus from background—and this dynamic is not readily captured in a Bayesian model based solely on evidentially driven belief change.

One could argue that the transformation of credences through framing arguments should not be modeled in such an extreme way. A more Bayesian-friendly way would avoid setting the credences for background biological claims to one and by deductive elimination setting the credences for incompatible focal hypotheses to zero. The credences for background claims should instead be set to $1 - \epsilon$ and, by deploying some Bayesian elimination rule, the credences in incompatible hypotheses set to ϵ , where ϵ is very small but greater than zero. This strategy allows one to represent the overall inference process in science with one master model. When hypotheses get rehabilitated by new evidence or when background biological information deployed in a framing argument changes, then the relevant credences can be updated accordingly. In fact, Earman proposes his probabilistic notion of elimination to handle the possibility of rehabilitation. Bayesian updating simply cannot change credences that have reached certainty (one or zero).

This strategy of modeling framing does not undermine the conclusion that framing arguments require extraevidential changes to credences, for eliminative framing inferences will still transform credences from their values given the evidence to $1 - \epsilon$ or ϵ . Again, the evidence for background biological claims does not suffice to explain this transformation since the available evidence almost always fails to drive the credence high enough. Instead, it is the defeasible assumption to take the background as certain, based on what evidence we have now, that changes the credence with respect to a particular testing problem. Moreover, the move to construct

a Bayesian master model in this way comes at a price: elimination must be weakened. We must again appeal to some threshold rule that can “eliminate” hypotheses by setting their credences to something greater than zero based on background claims that have credences less than one. This inherits the problem facing Earman’s notion of probabilistic eliminability. Identifying the threshold for elimination and certainty, the value for ϵ , requires confronting Sober’s worry that there is no good epistemic justification for any particular threshold; any specific value for ϵ appears to be arbitrary. True deductive eliminative inferences, in the context of one specific testing problem, require the defeasible assumption to take background information as certain.

On the basis of these considerations, the modeling strategy best suited to capturing framing is one that constructs *local* Bayesian-style models for specific scientific inferences. Such local models require both the framing dynamic—transforming credences based on the defeasible assumption to take background claims as certain—and standard Bayesian updating. This strategy raises deep questions about how these local inferential models might figure into the general process of scientific inquiry and how we might model this general process (see, e.g., Henderson et al. 2010). But to capture the style of scientific reasoning we see in framing, we must have inferential models that have a high degree of context sensitivity and can accommodate the true eliminative aspect.¹⁷

7. Conclusion. Framing, a new application of eliminative inference, plays an important but indirect role in evidential reasoning by constraining the space of rival hypotheses. Direct evaluation of evidence is done using the constrained space of rivals and standard statistical tools, Bayesian or otherwise. The examples from molecular evolution show how this inferential process is instantiated in scientific practice.

In comparison to standard eliminative induction, framing has a number of virtues. It avoids the problem of statistical evidence by locating elimination in a different place in the inferential process. Elimination helps frame testing problems rather than make theory choices. It also reveals how widespread eliminative inferences are in science (framing is ubiquitous in practice) and focuses attention on evaluating the normative standards for framing arguments. Finally, it sophisticates our account of how background theory and information are actively used in scientific reasoning.

17. Looking through the Bayesian lens also reveals another difference between framing and the standard picture. The crucial element for modeling framing is the transformation of credences based on the test focus, whereas for standard EI there are no extraevidential changes to credences and the crucial element is the threshold rules for elimination.

While framing may be a different application of eliminative inference and arguably a better account of how elimination matters in science, it is compatible with the standard picture. Suppose that the standard picture provides a viable solution to the threshold problem so that it can cope systematically with statistical evidence. Then it is possible that framing (using background information to eliminate potential rivals and frame testing problems) and standard eliminative induction (using evidence to eliminate competing theories and make theory choices) could work in conjunction. Yet if the threshold problem proves insoluble and eliminative induction has limited applicability, elimination still has an important role to play in the structure of scientific inference.

REFERENCES

- Abzhanov, Arhat, Winston P. Kuo, Christine Hartmann, B. Rosemary Grant, Peter R. Grant, and Clifford J. Tabin. 2006. "The Calmodulin Pathway and Evolution of Elongated Beak Morphology in Darwin's Finches." *Nature* 442:563–66.
- Abzhanov, Arhat, Meredith Protas, B. Rosemary Grant, Peter R. Grant, and Clifford J. Tabin. 2004. "Bmp4 and Morphological Variation of Beaks in Darwin's Finches." *Science* 305:1462–65.
- Doyle, Sir Arthur Conan. 1890. *The Sign of Four*. London: Blackett.
- Dretske, Fred. 1981. "The Pragmatic Dimension of Knowledge." *Philosophical Studies* 40: 363–78.
- Earman, John. 1992. *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. Cambridge, MA: MIT Press.
- Eyre-Walker, Adam. 2002. "Changing Effective Population Size and the MacDonald-Kreitman Test." *Genetics* 162:2017–24.
- . 2006. "The Genomic Rate of Adaptive Evolution." *Trends in Ecology and Evolution* 21 (10): 569–75.
- Felsenstein, Joseph. 1985. "Phylogenies and the Comparative Method." *American Naturalist* 125:1–15.
- Forber, Patrick. 2008. "Forever beyond Our Grasp?" *Biology and Philosophy* 23:135–41.
- . 2010. "Confirmation and Explaining How Possible." *Studies in History and Philosophy of Biological and Biomedical Sciences* 41:32–40.
- Fry, Brian G. 2005. "From Genome to 'Venome': Molecular Origin and Evolution of the Snake Venom Proteome Inferred from Phylogenetic Analysis of Toxin Sequences and Related Body Proteins." *Genome Research* 15:403–20.
- Fry, Brian G., Natalie G. Lumsden, Wolfgang Wüster, Janith C. Wickramaratna, Wayne C. Hodgson, and R. Manjunatha Kini. 2003. "Isolation of a Neurotoxin (α -Colubritoxin) from a Nonvenomous Colubrid: Evidence for Early Origin of Venom in Snakes." *Journal of Molecular Evolution* 57:446–52.
- Fry, Brian G., Holger Scheib, Louise van der Weerd, Bruce Young, Judith McNaughtan, S. F. Ryan Ramjan, Nicolas Vidal, Robert E. Poelmann, and Janette A. Norman. 2008. "Evolution of an Arsenal." *Molecular and Cellular Proteomics* 7 (2): 215–46.
- Fry, Brian G., and Wolfgang Wüster. 2004. "Assembling an Arsenal: Origin and Evolution of the Snake Venom Proteome Inferred from Phylogenetic Analysis of Toxin Sequences." *Molecular Biology and Evolution* 21 (5): 870–83.
- Fry, Brian G., Wolfgang Wüster, R. Manjunatha Kini, V. Brusica, A. Khan, D. Venkataraman, and A. P. Rooney. 2003. "Molecular Evolution and Phylogeny of Elapid Snake Venom Three-Finger Toxins." *Journal of Molecular Evolution* 57:110–29.
- Fry, Brian G., Wolfgang Wüster, S. F. Ryan Ramjan, Timothy Jackson, Paolo Martelli, and R. Manjunatha Kini. 2003. "Analysis of Colubroidea Snake Venoms by Liquid Chro-

- matography with Mass Spectrometry: Evolutionary and Toxinological Implications." *Rapid Communications in Mass Spectrometry* 17:2047–62.
- Fry, Brian G., et al. 2006. "Early Evolution of the Venom System in Lizards and Snakes." *Nature* 439:584–88.
- Graur, Dan, and Wen-Hsiung Li. 2000. *Fundamentals of Molecular Evolution*. 2nd ed. Sunderland, MA: Sinauer.
- Hawthorne, James. 1993. "Bayesian Induction *Is* Eliminative Induction." *Philosophical Topics* 21:99–138.
- Henderson, Leah, Noah D. Goodman, Joshua B. Tenenbaum, and James Woodward. 2010. "The Structure and Dynamics of Scientific Theories: A Hierarchical Bayesian Perspective." *Philosophy of Science* 77:172–200.
- Joyce, James. 2007. "Bayes' Theorem." In *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta. Stanford, CA: Stanford University, <http://plato.stanford.edu/archives/sum2007/entries/bayes-theorem/>.
- King, Jack L., and Thomas H. Jukes. 1969. "Non-Darwinian Evolution." *Science* 164:788–98.
- Kitcher, Philip. 1993. *The Advancement of Science*. Oxford: Oxford University Press.
- Kornegay, Janet R., James W. Schilling, and Allan C. Wilson. 1994. "Molecular Adaptation of a Leaf-Eating Bird: Stomach Lysozyme of the Hoatzin." *Molecular Biology and Evolution* 11 (6): 921–28.
- Kreitman, Martin. 1983. "Nucleotide Polymorphism at the Alcohol Dehydrogenase Locus of *Drosophila melanogaster*." *Nature* 304:412–17.
- . 2000. "Methods to Detect Selection in Populations with Applications to Humans." *Annual Review of Genomics and Human Genetics* 1:539–59.
- Kyburg, Henry. 1961. *Probability and the Logic of Rational Belief*. Middletown, CT: Wesleyan University Press.
- Maher, Patrick. 2004. "Probability Captures the Logic of Scientific Confirmation." In *Contemporary Debates in Philosophy of Science*, ed. Christopher Hitchcock, 69–93. Malden, MA: Blackwell.
- Margoliash, Emanuel, and Emil L. Smith. 1965. "Structural and Functional Aspects of Cytochrome in Relation to Evolution." In *Evolving Genes and Proteins*, ed. V. Bryson and H. Vogel. New York: Academic Press.
- McDonald, John H., and Martin Kreitman. 1991. "Adaptive Protein Evolution at the Adh Locus in *Drosophila*." *Nature* 351:652–54.
- Messier, Walter, and Caro-Beth Stewart. 1997. "Episodic Adaptive Evolution of Primate Lysozymes." *Nature* 385:151–54.
- Nei, Masatoshi, Xun Gu, and Tatyana Sitnikova. 1997. "Evolution by the Birth-and-Death Process in Multigene Families of the Vertebrate Immune System." *Proceedings of the National Academy of Sciences* 94 (15): 7799–7806.
- Norton, John D. 1995. "Eliminative Induction as a Method of Discovery: How Einstein Discovered General Relativity." In *The Creation of Ideas in Physics*, ed. Jarrett Leplin, 29–69. Dordrecht: Kluwer.
- Platt, John R. 1964. "Strong Inference." *Science* 146:347–53.
- Sober, Elliott. 1990. "Contrastive Empiricism." In *Scientific Theories*, vol. 14, ed. W. Savage, 392–412. Minneapolis: University of Minnesota Press.
- . 1993. "Epistemology for Empiricists." *Midwest Studies in Philosophy* 18:39–61.
- . 1999. "Testability." *Proceedings and Addresses of the American Philosophical Association* 73:47–76.
- . 2005. "Is Drift a Serious Alternative to Natural Selection as an Explanation of Complex Adaptive Traits?" In *Philosophy, Biology and Life*, ed. A. O'Hear. Cambridge: Cambridge University Press.
- . 2008. *Evidence and Evolution: The Logic behind the Science*. Cambridge: Cambridge University Press.
- Sober, Elliott, and Stephen Hecht Orzack. 2003. "Common Ancestry and Natural Selection." *British Journal for the Philosophy of Science* 54:423–37.
- Sorensen, Roy. 2006. "Epistemic Paradoxes." In *The Stanford Encyclopedia of Philosophy*,

- ed. Edward N. Zalta. Stanford, CA: Stanford University, <http://plato.stanford.edu/archives/fall2006/entries/epistemic-paradoxes/>.
- Stanford, P. Kyle. 2006. *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives*. Oxford: Oxford University Press.
- Stewart, Caro-Beth, James W. Schilling, and Allan C. Wilson. 1987. "Adaptive Evolution in the Stomach Lysozymes of Foregut Fermenters." *Nature* 330:401–4.
- Stewart, Caro-Beth, and Allan C. Wilson. 1987. "Sequence Convergence and Functional Adaptation of Stomach Lysozymes from Foregut Fermenters." *Cold Spring Harbor Symposium on Quantitative Biology* 52:891–99.
- Vidal, Nicolas, and S. Blair Hedges. 2002. "Higher-Level Relationships of Caenophidian Snakes Inferred from Four Nuclear and Mitochondrial Genes." *Comptes Rendus Biologies* 325:987–95.
- Vonk, Freek J., et al. 2008. "Evolutionary Origin and Development of Snake Fangs." *Nature* 454:630–33.