

Chapter 11

A PROBLEM WITH BAYESIAN CONDITIONALIZATION

Richard Otte

University of California, Santa Cruz

1. INTRODUCTION

Bayesianism is often divided into a static and dynamic part. The static part is a theory of rational belief at an instant, and the dynamic part is a theory of rational belief change. In this paper I will begin by looking closely at the rule of conditionalization, which is central to the dynamic side of Bayesianism. I will argue that the dynamic side of Bayesianism needs to make distinctions central to traditional non-Bayesian epistemology; Bayesians may consider this unfortunate, since Bayesianism is supposed to be a way to avoid many of the problems that arise in classical epistemology. I will then give a counterexample that shows we can be rational even if we do not conditionalize in situations in which conditionalization is applicable. This counterexample also raises problems for the view that conditionalization is a way to manage our beliefs in pursuit of some ideal. Contrary to Bayesian doctrine, I will argue that we generally cannot use conditionalization to manage our beliefs. The reason for this is that conditionalization is an externalist requirement that aims at the goal of being in a good epistemic situation, but rationality contains only internal requirements; thus conditionalization is not required in order to be rational. I will then look at whether we can develop an internalist version of

conditionalization which does not have these problems. It turns out that although an internalist requirement of conditionalization can solve some problems, it faces other serious problems.

2. THE REQUIREMENT OF BAYESIAN CONDITIONALIZATION

If as a result of experience we come to believe E to degree 1, conditionalization tells us how to modify the rest of our beliefs to take account of learning E . The requirement of conditionalization is very simple: our new probability function, P_n , is obtained from our old probability function, P_o , by conditionalizing on our new evidence. If E represents our new evidence, we can write the requirement of conditionalization as: $P_n(*) = P_o(*|E)$. This requirement of conditionalization only applies when E is everything we have learned from experience and when our new probability of E is 1.

Although Bayesians often claim that all rational belief change is by conditionalization, upon closer inspection we find that things are not as clear as first appeared. Contrary to what is usually said, Bayesianism must admit at least two different ways in which belief may rationally be changed. This distinction is seldom stated; perhaps this is due to relying too much on generalizations such as 'all belief change is via conditionalization.' First, one may change belief in direct response to experience. Although this is rarely discussed in connection to Bayesianism, we must allow for the rationality of changing belief as a direct result of having some experience; otherwise there would be nothing to apply the rule of conditionalization to. Some beliefs may be completely based on experience, and other beliefs may be partly based on experience and partly based on other beliefs. What is important is that beliefs can be rationally changed as a direct result of experience, and not as a result of other changes in belief. There are no Bayesian requirements or rules that govern belief change directly based on experience; everything is permitted.

The second way one may rationally change belief is in response to some other change in belief. Bayesianism provides rules, either traditional conditionalization or Jeffrey's generalization of conditionalization, that govern this type of belief change. These rules presuppose a distinction between beliefs that are changed as a result of experience and beliefs that are changed only in response to some other belief change, and these rules only govern belief change that is not directly affected by experience. It may be that there is rational change of belief not based on either experience or other

belief change; memory, and a priori beliefs may be examples of these. If so, Bayesianism will also have to let them be the basis of rational belief change.

Since conditionalization only applies to beliefs that are changed in response to other beliefs being changed by experience, in order to state when conditionalization is applicable, we need to first distinguish beliefs that are changed in direct response to experience from beliefs that are changed in response to other beliefs. We can then give the requirement of conditionalization:

ROC: If E contains all and only those beliefs that have been directly affected by experience, and if $P_n(E) = 1$, then $P_n(*) = P_o(*|E)$.

Notice that conditionalization will not be applicable to situations in which some beliefs are at least partly based on experience and are not raised to degree 1. Richard Jeffrey developed a version of conditionalization that did not require what was learned on the basis of experience to be believed to degree 1. Jeffrey noted that learning experiences often directly result in partial beliefs. Jeffrey conditionalization is a generalization of ROC that tells us how to modify the rest of our beliefs in response to forming these partial beliefs directly on experience:

ROJC: If partition $E = \{E_1, E_2, \dots\}$ contains all of the beliefs that are changed in direct response to experience, then $P_n(A) = \sum_i P_o(A|E_i)P_n(E_i)$.

As with traditional conditionalization, Jeffrey's probability kinematics does not apply to beliefs that are directly based on experience; no restrictions are placed on belief change based on experience. According to ROC and ROJC, our beliefs are required to be in accord with conditionalization only in very specific circumstances.

3. A COUNTEREXAMPLE TO CONDITIONALIZATION

Conditionalization is normally presented as a requirement or necessary condition of rationality. A change of belief in a situation in which conditionalization is applicable cannot be rational if it violates conditionalization. In the following I will focus my discussion on ROC, but it should be clear that the problems I raise also apply to ROJC.

A problem arises because there are situations in which conditionalization is applicable according to ROC, but in which it is rational to not conditionalize. Consider examples with the following structure. Suppose an agent has a certain experience, and as a result comes to believe *E* to degree 1. The agent's old conditional probability of some proposition *A* given *E* is some number *r*. In this situation ROC is applicable, and would require that the agent's new degree of belief in *A* be *r*. But suppose the agent also believes (incorrectly) that belief *A* is directly affected by the experience. This false belief that *A* is based on experience may be a rational belief; it may be directly based on experience (and thus in *E*), or even in accord with conditionalization on previous belief changes. As a result, the agent feels free to rationally believe *A* to some degree other than *r*. In this situation the agent believes that the requirement of conditionalization does not place any restrictions on her belief in *A*. The agent is rational in these situations, even though she does not conditionalize and has violated ROC. She is not aware of having violated any rule of rationality and may be completely consistent; she has managed her beliefs well. Given other beliefs she may have, it may even be irrational for her to conditionalize. Examples of this sort are counterexamples to the requirement of Bayesian conditionalization, because the agents are rational in spite of not conditionalizing in a situation in which the rule of conditionalization is applicable.

Two examples with this structure involve Al the Bayesian rock climber. Suppose Al, a recent convert to Bayesianism, is rock climbing the Royal Arches route in Yosemite. Early in the climb Al looks at a piton in a crack and is thinking about whether it is trustworthy (would it hold him if he fell?). Al's subjective probability that the piton is trustworthy given that it is rusted is .7. Al looks at the piton, and as a result of experience believes to degree 1 that that it is rusted. Based on this belief change, Al modifies the rest of his beliefs and comes to believe to degree 1 that the piton is trustworthy; thus Al has violated ROC. However, based on experience Al mistakenly believes that his belief that the piton is trustworthy is directly based on experience. Wanting to be a good Bayesian, Al reflects on how to apply ROC to this situation. Since he mistakenly believes that his belief that the piton is trustworthy is directly based on experience, he mistakenly believes that his new beliefs are consistent with ROC. His mistaken belief about the source of his beliefs is rational, and thus it is rational for him to be certain the piton is trustworthy. But this rational change of belief violates ROC, and is thus considered irrational by traditional Bayesianism.

The above example does not depend on Al believing to degree 1 that the piton is trustworthy. He could have believed it to any lesser degree, as the following example illustrates. Later in the climb Al comes to a place where

he can not be certain where the route goes: he could climb a steep face with tiny holds to the left or a steep crack to the right. Al's subjective probability that the route goes left up the face given that he comes to such a fork on this route is .2 and he believes to .8 that the climb goes right up the crack given this fork. Since as a direct result of experience Al believes that he is at such a fork, in order to be consistent with ROC, Al would have to believe to .8 that the climb goes right and up the crack. However, after looking at the two options, Al finds himself believing to .9 that he should go left up the steep face instead of right up the horrendous looking crack. He is now convinced that the route cannot go up such a crack; it looks much too difficult for this route. As a direct result of experience, Al mistakenly believes that his belief to .9 that the climb goes left up the face is directly based on experience; in reality the belief is based only on his other beliefs and he is in a situation in which ROC is applicable. Since he is trying to manage his beliefs like a good Bayesian, he pauses to think about whether his belief that he should go left up the face is in accord with ROC. He immediately realizes that since that belief is directly based on experience, ROC does not place any restrictions upon it. Since Bayesianism permits any belief that is directly based on experience, he believes that his belief to .9 that the route goes left up the face is rational according to Bayesianism. Of course, since he believes this belief is directly based on experience, Al believes he will have to use ROJC to modify the rest of his beliefs to account for this, which he quickly does. Furthermore, Al believes that his belief to .9 that the climb goes left up the face is what any properly functioning mind would believe in these circumstances; the belief has warrant. As a result, it would be irrational for him to conditionalize and believe only to .2 that the climb goes left up the face.

These counterexamples are different from standard examples which show that Bayesian requirements are not sufficient for rationality. These examples show that the Bayesian requirement of conditionalization is not even necessary for rationality.¹ Furthermore, the last example shows that there are circumstances in which it may be irrational to form beliefs in accord with ROC.

4. MANAGING BELIEFS

The above counterexamples were based on the fact that we can be rational and yet mistaken about the basis of our beliefs. This possibility raises serious problems for the common view that we should manage our beliefs with Bayesian conditionalization. Bayesians often claim that conditionalization is

a requirement of ideal rationality; having our judgments in accord with ROC is an ideal that we should aim towards as rational beings.² Accordingly, we are supposed to use conditionalization to manage our beliefs in pursuit of this ideal.

However, there is a problem with this proposal. Even if conditionalization were a requirement of ideal rationality, we are generally unable to pursue it as a goal. In order to strive towards having beliefs in accord with ROC, we must be able to know in what situations we should apply it. We cannot use conditionalization to manage our beliefs and pursue the goal of ideal rationality without knowing which beliefs are at least partially directly based on experience (and are thus exempt from conditionalization) and which beliefs are based only on other beliefs (and are thus required to be in accord with conditionalization). But in general we are unable to reliably determine which beliefs are formed in direct response to experience and which ones are formed only on the basis of other beliefs. Through introspection it is very difficult to find beliefs that are based only on other beliefs. I have often attempted to determine which of my beliefs are directly based only on other beliefs and not on any experience, but I am unable to do so. For example, once a friend and I were lost on a dirt road in Joshua Tree while discussing Bayesianism. As we looked at maps and tried to figure out where we were, I attempted to determine if my beliefs satisfied ROC. To do so I had to determine which of the beliefs I was forming were based only on other beliefs and which beliefs were directly based on experience. Even though I was making a serious effort to determine what my beliefs were based on, I could not find a single belief that I was confident was directly based only on other beliefs and was not partly based on experience. Perhaps all of my beliefs in that situation were based on experience, but if so, it would be difficult to find many situations in which I have beliefs based only on other beliefs. As a result of my inability to determine which beliefs were based on other beliefs, I was unable to use conditionalization in that situation to manage my beliefs, even though I was consciously thinking about conditionalization. This inability seems to be the norm rather than the exception. For all we know, any of our beliefs could be directly influenced by experience and not based only on other beliefs. As a result, we cannot know if ROC applies to any of our beliefs, and thus we are unable to use ROC to manage our beliefs. If ROC is part of ideal rationality, then we have no idea how to begin pursuing this ideal.

This same problem arises if we attempt to manage our beliefs with Jeffrey's probability kinematics. As in the case with ROC, we generally are unable to know if the situation we are in is one in which ROJC applies. Like ROC, ROJC presupposes a distinction between beliefs that are directly based

on experience and beliefs that are based on other beliefs. For example, if a partial belief is directly based on experience, this degree of belief is rational and in accord with ROJC no matter what it is, even if it is different from what would be required if it were not directly based on experience. But if the belief is not based on experience, then it must be in accord with Jeffrey conditionalization. Since we cannot determine which of our partial beliefs are based on experience, we cannot use ROJC to manage our beliefs; for all we know almost all of our partial beliefs could be directly based on experience, and thus there would be no Bayesian requirements of belief change applicable to them.

These problems arise because ROC and ROJC are external requirements of rationality. In order to use ROC or ROJC to manage our beliefs we need access to the basis of our beliefs, which is not an internal matter. In general, we do not have access to whether our beliefs are based solely on other beliefs, or whether they are also partly based on experience.

Rationality is usually thought to involve only internal features of our cognitive life. We are not irrational for failing to account for things that we are unaware of. Even brains in vats and those whose brains are manipulated in various ways can be rational, if they properly deal with all that is accessible to them. For this reason some epistemologists classify Bayesianism as a version of coherentism, which is an internalist theory.³ Although they may be correct about the static part of Bayesianism, this is not a correct description of Bayesianism that includes ROC or ROJC. Since ROC makes reference to the basis of our beliefs, ROC is an external requirement, and any theory which requires agents to use conditionalization to manage their beliefs is an externalist theory. Because of this ROC is very different from other requirements of Bayesianism, such as coherence or even reflection. The static part of Bayesianism is an internalist theory; the agent is not required to have access to any non-internal factors in order to satisfy coherence.⁴ Similarly, an agent can satisfy reflection while having access only to internal factors. But when Bayesianism is expanded to a dynamic theory of belief change that includes ROC, then it is no longer an internalist theory. Since rationality requires only internal requirements, and since ROC is an external requirement, rationality does not require ROC. Our counterexample showed that one can satisfy all requirements of rationality and not satisfy ROC. Satisfying ROC may be necessary to be in a good epistemic situation, which externalists often cite as our epistemic goal, but it is not necessary for rationality.⁵ Since Bayesianism is usually construed as being a theory of rationality, this is in conflict with ROC. Of course, if ROC is required in order to be in a good epistemic situation, and if an agent can see how to fulfill ROC, then the agent should do what is required to get in that good epistemic situation. Being in a

good epistemic situation (which may include satisfying ROC) may be part of an ideal we should pursue if we are able. My claim is that since we generally are unable to know how to even begin working towards this good epistemic situation, we can be rational without pursuing it. Rationality may require us to aim at that ideal when we have some idea of how to get there, but rationality does not require us to aim at something that we have no idea how to pursue.

At this point one might object and argue that various Dutch book arguments demonstrate that rationality requires us to satisfy ROC and ROJC.⁶ But although these arguments may show that there is a problem in forming beliefs in accordance with some rule other than conditionalization, they do not reveal a problem for one whose violations of ROC do not follow any rule.⁷ In order to create a Dutch book for someone who violates ROC, one needs to know if her new belief will be greater or less than what conditionalization would require. But a Dutch book cannot be made against someone who mistakenly believes that a belief is based on experience and as a result violates ROC. The agent is not following any rule that will tell her when she will mistakenly believe some belief is based on experience and what degree of belief she will then form. As a result, we don't know if the new belief will be greater or less than what ROC requires. Thus the Dutch book arguments for conditionalization do not help with our problem.

The difficulty in distinguishing between beliefs based directly on other beliefs and beliefs based on experience may be partly responsible for the appeal of accounts that postulate an observation language that adequately describes our experience. On these proposals, the agent is supposed to have a belief in an observation language that captures the content of his or her experience.⁸ Other beliefs in a natural language are then arrived at by conditionalizing on beliefs in the observation language. This provides a natural way to classify the origin of a belief: beliefs in the observation language are directly based on experience, and beliefs not in the observation language are directly based on other beliefs (that are ultimately based on beliefs in the observation language). However, there is no reason to think humans must have such an observation language in order to be rational. And even if we did have this observation language, we would still need to require that only beliefs in the observation language be directly based on experience and not on other belief change, and that only beliefs in the natural language be directly based on other beliefs and not on experience; this is by no means implied by the existence of an observation language. Thus the same problem arises, since there is no convincing argument that experience only directly affects beliefs in the observation language.

5. INTERNAL RULES OF CONDITIONALIZATION

Although ROC is an external requirement, it may be possible to develop an internalist version of ROC that agents could use to manage their beliefs. ROC was an external theory because we are generally unaware of the basis of our beliefs. But we do have access to our beliefs about the basis of our beliefs, and thus beliefs about the basis of our beliefs are internal factors. One possible internalist requirement of conditionalization is the following:

IROC: If we are now certain that E contains all and only those beliefs that have been directly affected by experience, and if $P_n(E) = 1$, then $P_n(*) = P_o(*|E)$.

We can formulate an internalist version of ROJC as follows:

IROJC: If we are now certain that partition $E = \{E_1, E_2, \dots\}$ contains all and only those beliefs that have been changed in direct response to experience, then $P_n(*) = \sum_i P_o(*|E_i)P_n(E_i)$.

Principle IROC refers to the belief that E contains all and only those beliefs that have been directly affected by experience. Unlike ROC, Principle IROC requires we conditionalize only when we believe that E contains all and only those beliefs that have been directly affected by experience. As a result, requirement IROC does not face the above problem that ROC faced. It can be used by agents to manage beliefs, and agents are able to strive towards it as an ideal. A rational person could satisfy IROC without satisfying ROC. An example is our original counterexample of Al the Bayesian rock climber in which a person falsely believes that a belief is directly based on experience and the new degree of belief is not in accord with conditionalization. Although this belief change violates ROC, it is consistent with IROC and is rational. Furthermore, a person can satisfy ROC without satisfying IROC, although these are cases in which the person is irrational. If a person mistakenly believes ROC is applicable and fails to conditionalize, then that person is guilty of not managing his or her beliefs properly, even though ROC may be satisfied. In these situations ROC makes no requirements on rational belief, but IROC does make requirements. The agent is irrational in these situations, even though ROC is satisfied. These differences between ROC and IROC arise because ROC is an external requirement, and since rationality only has internal requirements, ROC is not a necessary condition of rationality. But since IROC is an internal requirement, it is not ruled out as a necessary condition of rationality.

A serious problem with IROC is that it is a very weak requirement. IROC places no requirements on an agent who is not certain of the basis of his or

her beliefs. Since it is common to not be certain of the basis of our beliefs, IROC is a rather empty requirement of rationality. But Bayesians certainly don't think of conditionalization as an empty vacuous requirement.

One might respond by claiming that I have misconstrued the requirement of conditionalization and what it means to learn something as a result of experience. Suppose we temporarily ignore whether a belief is directly based on experience, and simply look at situations in which an agent changes certain degrees of belief to degree 1. One way of looking at this proposal would be to view it as assuming that all of the beliefs that I learn as a result of experience change to degree 1. Conditionalization would then require that the agent modify her other beliefs to take account of changing these beliefs to degree 1. Consider the following:

IROC*: If E contains all of the beliefs that change to degree 1 ($P_n(E) = 1$), then $P_n(*) = P_o(*|E)$.

This proposal has the significant advantage that the agent does not have to know which of her beliefs are directly based on experience, or even have any beliefs about the basis of her beliefs; any belief that is changed to degree 1 is conditionalized on. Upon this interpretation of conditionalization we can manage our beliefs with conditionalization, because we can know which of our beliefs change to degree 1.

A problem with this proposal arises when we consider situations like the following. Suppose E is all of my beliefs that change to degree 1, and that $P_o(A|E) = .8$. But instead of believing A to .8 as required by IROC*, I believe A to .4. Am I irrational? If I think that my new belief in A is directly based on experience, then I have no reason to apply any rule of conditionalization to it. This is a counterexample to IROC*, because I am rational, but don't conditionalize according to IROC*. The problem with IROC* is that it excludes the possibility of beliefs based on experience that are not changed to degree 1, but rationality does not require we believe that all beliefs based on experience are raised to 1. As a result, IROC* is not successful.

Let us also consider a version of this proposal that deals with partial belief:

IROJC*: If partition $E = \{E_1, E_2, \dots\}$ contains all the beliefs that have been changed, then $P_n(*) = \sum_i P_o(*|E_i)P_n(E_i)$.

Principle IROJC* does not require that all learned beliefs or beliefs based on experience be raised to degree 1. However, this principle has a serious problem; it is trivially true. One always satisfies this rule, no matter what one's beliefs are. If I am not required to distinguish between what I learn by

experience and my other beliefs, in effect all of my new beliefs are based on experience. Thus IROJC* places no restrictions on rational belief change.

One might propose to avoid these problems by requiring that the agent have a rational belief about the source of her beliefs. We could thus modify IROC as follows:

IROC**: If E contains those and only those beliefs that S rationally believes are based directly on experience, and if $P_n(E) = 1$, then $P_n(*) = P_o(*|E)$.

Principle IROC** will not be vacuously true, because we may have rational beliefs about the basis of our beliefs.

However, a very serious problem now arises. Even though IROC** is an internal rule, an agent will not be able to use it to manage her beliefs. Suppose that for any belief, a belief about whether that belief is based only on other beliefs or is based on experience is a next higher order belief. If we let our ordinary beliefs be of level 1, then our beliefs about which of those beliefs are based on experience will be beliefs of level 2. Although this could be stated precisely, for our purposes we only need this rough intuitive notion. Of course, there will also be level 3 beliefs about which of our level 2 beliefs are directly based on experience, etc. Now suppose our agent decides to use IROC** to manage her beliefs. She must first form a rational second order belief about which of her first order beliefs are based on other beliefs and which are based on experience. But in order to rationally form that second order belief, she will have to form a rational third order belief about whether that second order belief is based on experience. One can easily see the regress our agent will fall into if she attempts to use IROC** to manage her beliefs. In order to use IROC** to rationally manage a belief of level n, the agent must first form a rational belief of level n + 1 about the source of her level n beliefs. Thus in order to manage her beliefs at level n she will need to first manage her beliefs at level n + 1. This is an instance of a general problem in applying any epistemic norm. It appears that IROC** is no better off than ROC; neither can be used by an agent to manage beliefs.

6. CONCLUSION

We began by clarifying the requirement of conditionalization and saw there are two basic ways that beliefs may be rationally changed: in response to experience and in response to other belief change.⁹ We presented a counterexample to the requirement of conditionalization and discussed the consequences of this counterexample for the view that conditionalization is a way to manage our beliefs in pursuit of some ideal. We then argued that since we are usually not able to determine which of our beliefs are directly

based on other beliefs, we are unable to apply rules of conditionalization to manage our beliefs. Bayesian requirements of belief change are external requirements and appeal to factors that are not generally accessible to epistemic agents. Conditionalization is based on an internally inaccessible distinction between beliefs that are directly based on experience and beliefs that are based on other beliefs, and Bayesians will have to adopt this traditional epistemic distinction if they wish to require a rule of conditionalization. This is unfortunate, because one of the appeals of Bayesianism is that it claims to not require most of the traditional epistemic distinctions.

Since rationality is an internal notion, it does not rely on external factors such as those required by conditionalization. Internalist versions of the requirement of conditionalization can be developed, and these avoid some of the problems that the externalist versions faced. But we saw that these rules are vacuous, have counterexamples, or face a regress problem. Thus neither external nor internal versions of conditionalization are helpful in managing beliefs.

ENDNOTES

¹ Once one sees the structure of these counterexamples one finds that it is very easy to construct other counterexamples. For example, suppose Susan's subjective probability that there is a rose in a garden given there is an assortment of flowers in the garden is .8. She sees the garden, and believes to degree 1 that there is an assortment of flowers in it. Based on this new belief, Susan modifies her other beliefs and contrary to ROC comes to believe to degree .9 that there is a rose in the garden. Susan mistakenly believes that her belief that there is a rose in the garden is directly based on experience, and that her new beliefs are consistent with ROC. Susan is rational in her mistaken belief, and thus it is rational for her to believe to .9 that there is a rose in the garden.

² It is not at all clear what ideal rationality is; for example, some might think one was ideally rational only if one believes all truths, or all logical truths, or....Furthermore, it is difficult to see how conditionalization can be a requirement of ideal rationality given that it is irrational to conditionalize in some situations (such as the previous counterexample). However, we will ignore these problems for now.

³ See Alvin Plantinga 1993.

⁴ Logical omniscience is required, but let us assume that this is not an external requirement.

⁵ See William Alston 1985 and 1988.

⁶ See van Fraassen 1989 and Teller 1973.

⁷ See Bas van Fraassen 1989.

⁸ See J.H. Sobel 1990.

⁹ Once this is seen, one realizes that some claims often made about Bayesianism are false. For example, the frequent claim that according to Bayesianism all rational belief change is by conditionalization is seen to be false; one can also change beliefs in response to experience. Another example is the claim that a rational agent should not assign a probability of 1 or 0 to any contingent proposition. The reason given is that if a proposition has a probability of 0 or 1, then it is impossible to change it by conditionalization; no matter what is conditionalized

on, the result will remain the same. In order to preserve the ability to change our minds about a contingent proposition it is claimed that we must not assign the proposition to a probability of 0 or 1. But from our description of Bayesianism it is easy to see the flaw in this argument. It is correct that by conditionalization an agent can never change a degree of belief in a proposition of degree 0 and 1. But it is also possible that as a result of experience the agent could change his belief from 0 or 1 to some other degree of belief. Nothing in Bayesianism restricts the beliefs we can form in direct response to experience. Thus the primary argument given to not assign degrees 0 or 1 to contingent propositions fails.

REFERENCES

- Alston, William. 1988. An internalist externalism. *Synthese* 74: 265-283.
Jeffrey, Richard. 1983. *The Logic of Decision*, 2nd ed. Chicago: University of Chicago Press.
Plantinga, Alvin. 1993. *Warrant: The Current Debate*. Oxford: Oxford University Press.
Sobel, J.H. 1990. Conditional probabilities, conditionalization, and dutch books. In *PSA 1990* Volume I, eds. Arthur Fine, Micky Forbes, and Linda Wessels, 503-515. East Lansing, MI: Philosophy of Science Association.
Teller, Paul. 1973. Conditionalization and observation. *Synthese* 26: 218-58.
Van Fraassen, Bas. 1989. *Laws and Symmetry*. Oxford: Oxford University Press.