

A More Fulfilling (and Frustrating) Take on Reflexive Predictions

Matthew Kopec^{†‡}

Even though social scientists continue to discuss the problems posed by self-fulfilling and self-frustrating predictions, philosophers of science have ignored the topic since the 1970s. Back then, the prevailing view was that the methodological problems posed by reflexive predictions are either minor or easily avoided. I believe that this consensus was premature, ultimately relying on an overly narrow understanding of the phenomenon. I present an improved way to understand reflexive predictions (framed in probabilistic terms) and show that, once such predictions are understood this way, the methodological problems they pose may turn out to be neither minor nor easily avoided.

1. Introduction. I will begin with a classic example. Imagine you are an economist studying Smith's bank, and you notice after crunching some numbers that your theories predict that Smith's bank will soon collapse. Unbeknownst to you, your theories are garbage, and the bank will be just fine if you keep your mouth shut. But you decide to publish your findings, and, upon hearing of your prediction, Smith's concerned customers quickly run to the bank in an attempt to withdraw their savings. Smith's bank collapses upon the loss of all of its liquid assets, and therefore your prediction came true.

If you believe that economists are sometimes concerned with testing theories, then you may ask: Is the collapse of Smith's bank evidence for

[†]To contact the author, please write to: Department of Philosophy, University of Wisconsin–Madison, 5185 Helen C. White Hall, 600 North Park Street, Madison, WI 53706; e-mail: kopec@wisc.edu.

[‡]I would like to thank John Basl, Malcolm Forster, Michael Goldsby, Casey Helgeson, and Joel Velasco for helpful comments and, especially, Dan Hausman, Elliott Sober, and Peter Vranas for discussions that helped me develop the project. I would also like to thank audiences at the University of Wisconsin–Madison and at the 2010 meeting of the PSA in Montreal for helpful questions and comments.

Philosophy of Science, 78 (December 2011) pp. 1249–1259. 0031-8248/2011/7805-0044\$10.00
Copyright 2011 by the Philosophy of Science Association. All rights reserved.

the truth of some or all of the theoretical machinery that predicted it?¹ The above example is one type of reflexive prediction, which is roughly a prediction that has some causal impact on the state of affairs predicted. I will argue that reflexive predictions like these can cause serious difficulties for those who claim that observations are evidence for or against the theories that make such predictions.

The plan of the essay is as follows. First, I will refine the definition of a reflexive prediction that was settled upon in the philosophical literature. Because I believe that an important feature of reflexive predictions has been largely ignored, I introduce the notion of a ‘weakly reflexive prediction’ framed in probabilistic terms that I take to be an important addition to the literature. Second, I show that some examples of such predictions cause evidential problems if confirmation relations are understood from a Bayesian or likelihoodist framework. In particular, if a reflexive prediction is disseminated, then observing that the prediction obtains may be (1) no evidence for or against, (2) evidence against, or (3) better evidence for the theory that led to the prediction, depending on the details of the case. I conclude by pointing out the work left to be done. Although philosophers have ignored reflexive predictions since the 1970s, I hope to show that they continue to threaten science with methodological problems, and they therefore deserve our attention.

2. A New Definition. While social scientists interested in reflexive predictions continue to use the definition of this phenomenon given by Merton (1948), the definition subsequently went through some important modifications in the philosophical literature (cf. Buck 1963; Grunbaum 1963; Romanos 1973; Vetterling 1976). The definition that was settled upon in the literature is due to Romanos (1973). It can be stated as follows:

Romanos’s Reflexive Predictions. A prediction is reflexive if and only if “the formulation/dissemination style of the prediction [is] a causal factor relative to the prediction’s coming out true or false” (Romanos 1973, 106).

By ‘formulation style’ Romanos means the way the prediction is made. For instance, it could be made in one’s mind, through a string of sounds that count as English words, or through ink marks that count as Japanese words, which all count as different formulation styles. The ‘dissemination style’, however, is the mode of reproduction and transmission of the prediction. For example, a prediction transmitted through network television

1. In what follows, I will use terms like ‘theory’, ‘theoretical machinery’, ‘hypothesis’, and ‘model’ loosely and interchangeably. Although there may be important differences, I take it that these are not crucial for my purposes here.

has a different dissemination style than one transmitted through word of mouth.

This definition requires preliminary modification. First of all, I believe it is better to think of predictions as abstract objects. A prediction on my account is the claim that a particular state of affairs will obtain. Although we commonly think of claims as speech acts, we cannot intend that meaning of ‘claim’ here. Specifying the particular mode of expression would be problematic because a claim can be expressed in any number of ways, from something as simple as saying it out loud to something as complicated as designing an institution around it. For instance, placing a child who arrives at a new school in a reading group for poor performers is one way of expressing the claim that the student will be a poor performer (to anyone who knows that the group is for poor performers). So it is best to think of a claim in the abstract, that is, as a proposition that holds that another proposition will obtain. Because of this, I do not see the need for a distinct concept of ‘formulation’. Instead I prefer to talk in terms of ‘modes of dissemination’, which I believe to cover the important aspects of both notions. Making this modification results in the following definition:

Revised Romanos’s Reflexive Predictions. A prediction is reflexive if and only if the mode of disseminating the prediction is a causal factor relative to the prediction’s coming out true or false.

While this definition is an improvement, I will show that it is overly restrictive and, therefore, ignores a whole class of phenomena that rightly deserve the label.

To determine what sorts of phenomena are isolated by this definition, we need to figure out exactly what Romanos means by a “causal factor relative to the prediction’s coming out true or false.” He gives some indication when he discusses the impact of self-fulfilling predictions on theory testing. A scientist has a theory T_1 , and this theory entails that the event e_1 will occur, and she thus forms the prediction P that holds that this event will occur. Now suppose that there is a particular mode of disseminating this prediction, and the event of disseminating the prediction in this way we will call e_2 . Romanos states: “Now I suggest that P will turn out to be a self-fulfilling (reflexive) prediction just in case there is some other, well accepted or likely theory, T_2 , such that according to T_2 , given certain conditions C (which obtain), e_2 is sufficient to bring about the occurrence of e_1 (i.e. the event originally predicted by P)” (1973, 107). I take it that focusing on there being a well-accepted or likely theory that suggests that e_2 is sufficient for e_1 is an error on his part. This seems to be a straightforward case of confusing what it is for P to be a reflexive prediction with what it is to know (or at least be justified in believing)

that P is a reflexive prediction. But putting that aside, Romanos clearly thinks that the relationship is one in which (if C obtains) e_2 is a sufficient condition for e_1 to obtain.

Romanos's sufficiency criterion causes a difficulty: e_2 's being sufficient for e_1 to occur does not necessarily make the prediction self-fulfilling because of the possibility of causal overdetermination. If the predicted event would have happened without the prediction being disseminated, then the prediction is not self-fulfilling in the traditional sense. The intuitive idea is that the prediction must play a crucial role in making it true, or else it is not "self"-fulfilling. So the relevant issue when dealing with reflexive predictions is whether the mode of dissemination is sufficient to make the truth-value of the prediction opposite of what it would be without it. Making these features explicit, we obtain the following definition, which marks off a class of predictions that I will call 'strongly reflexive predictions'.

Strongly Reflexive Prediction. A prediction is strongly reflexive if and only if the mode of dissemination is sufficient to switch the truth-value of the prediction from what it would be if not disseminated.

Although this is much improved, I will now show that this only isolates a subset of all reflexive predictions and is thus overly restrictive as a general definition.

To see how problematic the criterion of sufficiency is, consider the following scenario. Smith and Jones are running for president of the United States. After applying some theoretical machinery to an extensive amount of unpublished poll data, a well-respected analyst determines that Jones will win the election. Before our analyst decides to disseminate this prediction, the situation is as follows. The Americans who will vote in the election and who have decided for whom they will vote actually favor Smith by nine votes. There are only 10 undecided voters in the country, and each one plans to flip a coin: heads she will vote for Smith, tails she will vote for Jones. But if our analyst decides to disseminate the prediction over the television, a bandwagon effect (cf. Simon 1954) will cause 18 additional Americans to vote for Jones, even though they previously had no intention to vote at all. Neither the previously decided voters nor the coin flippers are affected by the news. Notice how disseminating this prediction in this way causes a difference. If the analyst says nothing, then the probability that Smith will win the election is greater than 99.9% since only one of the 10 coins would have to come up heads. But if she spreads her prediction that Jones will win the election over the airwaves, then the probability that Jones will win is greater than 99.9%. But notice that if sufficiency is what is required for a prediction to be reflexive, then the prediction in this case is not self-fulfilling. It is not the case that the

mode of dissemination is sufficient for the election results to flip since there is still some chance (no matter how small it is) that the truth-value stays the same, even after the analyst goes on television.

The previous example shows how restrictive the conditions stated in the definition of a strongly reflexive prediction are, since they rule out any prediction whose truth-value bears even slightly on chance. This definition is therefore unlikely to apply to any predictions made in the social sciences. It is no wonder that the excitement over such predictions ended back in the 1970s. Cases like the one above obviously have reflexive tendencies, and so we need a new concept that rightly includes such cases.

My proposal is to point out that the strongly reflexive predictions are a proper subset of a larger class that I will call 'weakly reflexive predictions'. The definition of this class must respect the truth-making or false-making tendency that predictions can have, even if that tendency is not sufficient to determine the difference. I propose the following definition:

Weakly Reflexive Prediction. A prediction is weakly reflexive if and only if the mode of dissemination is sufficient to change the probability of the predicted event occurring from what it would be if not disseminated.

If the mode of dissemination of a prediction raises the probability of the predicted event occurring, then the prediction is *weakly self-fulfilling*, and if it lowers the probability then it is *weakly self-frustrating*. With this modification, the new definition allows for cases like the election prediction above. It is important to note that unlike strongly reflexive predictions, weakly reflexive predictions vary in degrees. The election prediction has a very strong self-fulfilling tendency, even though it still merely counts as weakly reflexive.

I believe this concept of a weakly reflexive prediction is an important addition to the literature on reflexive predictions. As I will show in the next section, once we take weakly reflexive predictions seriously, we notice that the evidential issues surrounding reflexive predictions are much more complicated than previously thought.

3. Evidential Problems with Reflexive Predictions. In cases in which social scientists are interested in testing hypotheses using observations, the possibility that a hypothesis could be reflexive is problematic. In this section, I will examine the problems caused for someone who follows a Bayesian or likelihoodist confirmation-theoretic framework. Surprisingly, disseminating a reflexive prediction can change the evidential import of observing that the prediction obtains in any direction. I believe that similar problems will arise for testing under alternative frameworks, but I will not address these other frameworks here. For clarity's sake, I will present the frame-

work in a somewhat simplified manner. Also, for ease of presentation, I will focus predominately on self-fulfilling predictions since the problems are simply flipped in the case of self-frustrating predictions.

Bayesians and likelihoodists largely agree about how to consider the evidential import of a particular observation. For Bayesians, if a particular observation is more probable given a hypothesis than it is given the falsity of the hypothesis, then the observation is evidence for the theory. This allows us to determine whether an observation is evidence for a hypothesis by determining the value of the “likelihood ratio,” stated formally as

$$\frac{\Pr(O|H)}{\Pr(O|\sim H)},$$

where the numerator and denominator are conditional probabilities, O is the observation, and H and $\sim H$ represent the truth and falsity of the hypothesis, respectively.² For example, in the bank collapse example from the introduction, O would be the observation that the bank collapses, and H would be the macroeconomic theory that led the economist to predict that the bank would collapse. Likelihoodists prefer not to allow $\sim H$ into likelihood ratios since this “catchall” may cause technical difficulties. Instead, they move to testing specific (i.e., noncatchall) hypotheses against each other. So for them, $\sim H$ would have to be swapped with a particular hypothesis for comparison, say G. So, referring back to the bank example again, G might stand for a macroeconomic theory used by a rival school of economists. In both cases, if this ratio is greater than one, then the observation evidentially favors H, and the greater it is, the better the evidence O is in favor of H. If the ratio is less than one, the observation is evidence against H, and if it is equal to one the observation has no evidential import. When we are dealing with a token prediction that is reflexive, we have to consider another factor, which I will label D to stand for the particular mode of disseminating the prediction.³ For simplicity, I will let $\sim D$ represent that the prediction is not disseminated at all. Since D and $\sim D$ represent relevant information, both Bayesians and likelihoodists will maintain that these factors must be conditionalized on as

2. Throughout, we should rule out that O is identical with either H or $\sim H$ since that would cause technical problems because some of the relevant likelihood ratios will then be undefined.

3. Here I am using ‘token’ loosely since, on my view, predictions are propositions and propositions do not admit of the type/token distinction. By ‘token prediction’ I mean a token dissemination of a prediction. This avoids the confusion of speaking of ‘predictions’ and ‘token disseminations of a prediction’.

well, which means that the relevant quantities will be $\Pr(O|H\&D)$, $\Pr(O|\sim H\&D)$, $\Pr(O|H\&\sim D)$, and $\Pr(O|\sim H\&\sim D)$.⁴

First let's consider what this framework will tell us in a case in which a token prediction is strongly reflexive; let's say it is strongly self-fulfilling. Here is the scenario: our scientist is attempting to test her hypothesis H , and she disseminates the prediction that 'O will obtain' before making any observations. Since the prediction is strongly self-fulfilling, we know the predicted event O would not happen if the prediction were not disseminated but must happen if it is disseminated. So we know that $\Pr(O|D) = 1$, and it follows from this that if H and D are consistent (i.e., if it is possible for both the hypothesis to be true and the prediction to be disseminated), then $\Pr(O|H\&D) = 1$ as well. Similarly, if $\sim H$ and D are consistent, then $\Pr(O|\sim H\&D) = 1$. Since $\Pr(O|D) = \Pr(O|H\&D)$, then the truth-value of the hypothesis is irrelevant, or in other words it is "screened off" by the dissemination of the prediction. Therefore, the relevant likelihood ratio in this case becomes the "likelihood ratio: strongly self-fulfilling case":

$$\frac{\Pr(O|H\&D)}{\Pr(O|\sim H\&D)} = 1.$$

Since this ratio is equal to one, for both the Bayesian and the likelihoodist, the fact that O is observed bears no evidential value for testing the hypothesis. The case in which the prediction is strongly self-refuting is problematic in exactly the same way: just switch $\sim O$ for O in the above argument. Therefore, within a Bayesian or likelihoodist framework, a hypothesis that makes a strongly reflexive prediction cannot be tested by observing whether the predicted observation obtains.

Weakly reflexive predictions cause additional difficulties for testing hypotheses. If a token prediction is weakly self-fulfilling, then the probability of the predicted observation increases if the prediction is disseminated in that way; that is, $\Pr(O|D) > \Pr(O|\sim D)$. This allows a great deal of leeway in how the relevant likelihoods might change. To examine some of the possibilities, I will raise three toy examples, which, although highly unrealistic, exemplify the evidential worries. All three start out with the same scenario: there is an election in which three individuals, call them A , B , and C , are voting for one of two candidates, call them X and Y . You are an analyst working with the hypothesis that 'A will vote for X ', which I

4. The idea here is that since the information might be relevant to the evidential import of the observation, it would be illegitimate to conditionalize on H and $\sim H$ (or H and G) alone. A parallel explanation of why conditionalizing on certain additional facts is required, from the likelihoodist perspective, can be found in Sober (2009), where he discusses observation selection effects (a somewhat related phenomenon).

will call H, which leads you to make the prediction that 'X wins the election', and I will call the latter proposition O. So here the prediction is a proposition stating that O will obtain, that is, that X will win the election. If you do not tell the three voters about the prediction at all, then they will individually flip fair coins and vote for X if the coin lands on heads and vote for Y otherwise. We will call the event in the nondisseminated state $\sim D$.

We can figure out some of the probabilities involved in the scenario when the prediction is not disseminated. First, the probability that X will win the election is $1/2$ since that is determined by three coins being flipped. So $\Pr(O|\sim D) = 1/2$. The truth of the hypothesis in question makes a positive impact on the probability of X winning since if A votes for X then the only way X can lose the election is if both B's and C's coins land tails. So $\Pr(O|H\&\sim D) = 3/4$. However, if the hypothesis is false (i.e., A votes for Y), then that lowers the probability that O will obtain since the only way for X to win under such circumstances would be if both B's and C's coins land heads. So $\Pr(O|\sim H\&\sim D) = 1/4$. This state of affairs under nondissemination will stay the same for the following three scenarios.

In the first scenario, if the three voters are told about the prediction, let's say over a loudspeaker, the announcement will cause B to vote for X outright, but A's and C's behaviors are unaffected. Call the event of disseminating the prediction over the loudspeaker D. Now that B is voting for X, this raises the probability that X will win the election. There is only one way X does not win, and that is if both A's and C's coins land tails. So $\Pr(O|D)$ is $3/4$. Since the probability of the predicted observation occurring was caused by D being higher than under $\sim D$, this qualifies as a weakly self-fulfilling prediction. We also have enough information to figure out what kind of impact disseminating the prediction in that way will have on the evidential import of observing that X wins the election. Since B is voting for X, if A votes for X as well (i.e., if H is true), then X wins the election, and so O is certain to obtain. So $\Pr(O|H\&D) = 1$. If, however, A does not vote for X, the outcome of the election comes down to how C's coin lands. So $\Pr(O|\sim H\&D) = 1/2$. With this information, we can judge the way that dissemination has affected the evidential import of O. We have that $\Pr(O|H\&D)/\Pr(O|\sim H\&D) = 2$ and $\Pr(O|H\&\sim D)/\Pr(O|\sim H\&\sim D) = 3$. All these quantities are listed in table 1. Since the likelihood ratio under D is lower than the likelihood ratio under $\sim D$, telling our voters about the prediction over the loudspeaker decreases the evidential import of observing that X won the election.

If scenarios like this occur in the social sciences, this has the potential to raise difficulties for theory testing. The worry is that if a social scientist is engaged in testing a theory by making observations, she may believe

TABLE 1. IMPACT OF DISSEMINATION ON EVIDENTIAL IMPORT.

	Pr(O H)	Pr(O ~H)	Likelihood Ratio
Scenario 1:			
Nondissemination	3/4	1/4	3
Dissemination	1	1/2	2
Scenario 2:			
Nondissemination	3/4	1/4	3
Dissemination	3/4	1	3/4
Scenario 3:			
Nondissemination	3/4	1/4	3
Dissemination	1	1/4	4

the observation to be better evidence than it really is. In these probabilistic cases, the theory is not quite screened off by the self-fulfilling prediction, but a tendency in that direction exists. Now, one may think that the situation as it stands is not so bad since the observation is still some evidence in favor of the hypothesis. The next scenario shows that this need not be the case.

The second scenario is the same as the first if the prediction is not shared, but things are caused to change substantially if the prediction is spread over the loudspeaker. First of all, the announcement causes A to flip her coin first and tell the other two how she is voting. If she is voting for X, then the other two flip their coins and vote as they would have before the prediction was made. But if A is not voting for X, then this causes both B and C to vote for X. So $\Pr(O|H\&D) = 3/4$ and $\Pr(O|\sim H\&D) = 1$. So, we have the following ratios: $\Pr(O|H\&D)/\Pr(O|\sim H\&D) = 3/4$ and $\Pr(O|H\&\sim D)/\Pr(O|\sim H\&\sim D) = 3$ as before (see table 1). Notice from these likelihood ratios that while observing that X wins the election is evidence for H if the prediction is not disseminated, if the prediction is spread over the loudspeaker, observing that X wins the election is actually evidence *against* H. If cases like this exist in the real world, then this raises possibly serious methodological problems. If a scientist were in one of these cases but did not know it, an observation that she takes to be evidence for her theory could actually be evidence against it.

The final scenario shows, surprisingly, that weakly reflexive predictions need not have a negative impact on the evidential import of an observation. In scenario 3, if the prediction is not disseminated, then everything is as before. If it is spread by the loudspeaker, then, like in scenario 2, A is caused to flip her coin first and share the result. If A is going to vote for X, then B and C are caused to join her. But if A is going to vote for Y, then the other two flip their coins and vote as they would have before the prediction was made. So $\Pr(O|H\&D) = 1$, while $\Pr(O|\sim H\&D) = 1/4$, and so $\Pr(O|H\&D)/\Pr(O|\sim H\&D) = 4$. Since $\Pr(O|H\&\sim D)/\Pr(O|\sim H\&\sim D) = 3$, we have the following ratios: $\Pr(O|H\&D)/\Pr(O|\sim H\&D) = 4$ and $\Pr(O|H\&\sim D)/\Pr(O|\sim H\&\sim D) = 3$.

$H \& \sim D) = 3$ as in the two earlier scenarios, observing that X wins the election is actually better evidence for H if the prediction is spread in this way (see table 1). Since we can manufacture cases that move the evidential relation in each direction, it should be clear that it is also possible to dream up a scenario in which the evidential impact of the observation is exactly the same, even if the token prediction is weakly reflexive.

These scenarios are meant to expose the curious fact that a token prediction's being weakly reflexive does not dictate how we should modify our reaction to the observation after it has been disseminated. Even if a token reflexive prediction is disseminated, it could be that the observation that the prediction obtains is just slightly less or slightly more evidence for the theory used to make the prediction than if it was not disseminated. Maybe observing that the prediction obtains is not evidence at all, or maybe it actually disconfirms the theory being tested. The details of the case will determine this, which shows that much more work will need to be done if we are to figure out how to react to the knowledge that a prediction that has been disseminated is weakly reflexive. Once we have figured out how we should change our evaluation of the evidential import of the observation, we can make use of this knowledge in expanded models. Since the hypotheses, theories, or models at issue either ignore this relevant feature or take account of it incorrectly, they must be replaced with improved versions.

4. Conclusion. Philosophers have been strangely silent on the topic of reflexive predictions, evidenced by the fact that the topic disappeared from the philosophical literature after the 1970s. But the topic is still a lively area of research in the social sciences. Economists continue to discuss the self-fulfilling nature of currency crises and have been examining to what extent actors familiar with the assumptions of economic theory, like the self-interested agent, have been caused to act like these caricatures. Political scientists are still researching bandwagon and underdog effects. Educational psychologists continue to work on the Pygmalion effect and search for ways to minimize it or use it to society's advantage. All this interest is understandable from a scientist's perspective. If we want to make observations to test theories in economics, sociology, psychology, political science, and the like, then it is troublesome if the way we make predictions makes an impact on the observations themselves. I hope it is clear from this essay that there is an opportunity for philosophers to help in sorting out some of these methodological issues. First, we need to make sure that the problems are framed in a proper way—in a way that is useful given the types of problems faced by social scientists. Previous accounts clearly violate this, which might explain why social scientists have ignored the philosophical literature and why even philosophers have not taken

up the topic in some time. I believe the way I have framed the problem is much improved since it respects a large class of closely related phenomena originally left out of the discussion. The problem seems to be a genuinely methodological one, affecting the relationship between observations and hypotheses central in the sciences.

I will conclude by quickly mentioning two projects I believe philosophers and methodologically oriented social scientists ought to pursue. The first is to work out the difficulties posed by these predictions in frameworks other than the Bayesian or the likelihoodist framework. This is especially important since frequentist approaches continue to dominate the social sciences. My hunch is that the view from within these frameworks will be similarly disquieting, but it will take effort to work out the details. Second, we will need to extend the notion of a weakly reflexive prediction to another closely related class of phenomena, that of reflexive probabilistic forecasts. A statement like ‘there is a 90% chance X will win the election’ is not a prediction, per se, but rather a statement about current chances. But even though it is not a prediction, one can imagine cases in which such statements have reflexive aspects. For instance, the dissemination of the forecast might change the chance of the event either closer to or further away from 90%. The machinery I have used to expose the problems raised by weakly reflexive predictions cannot be generalized to cases like these. Instead, scoring rules for assessing the accuracy of probabilistic forecasts, like the Brier score, for example, would have to play a role. This is similarly important to work out since many of the “predictions” made in the social sciences are really probabilistic forecasts. I would wager that similar problems will arise once the notion of weak reflexivity is extended to this class of phenomena, although again, the details will need to be worked out.

REFERENCES

- Buck, Roger C. 1963. “Reflexive Predictions.” *Philosophy of Science* 30:359–69.
- Grunbaum, Adolf. 1963. “Comments on Professor Roger Buck’s Paper ‘Reflexive Predictions.’” *Philosophy of Science* 30:370.
- Merton, Robert K. 1948. “The Self-Fulfilling Prophecy.” *Antioch Review* 8:193–210.
- Romanos, George D. 1973. “Reflexive Predictions.” *Philosophy of Science* 40:97–109.
- Simon, Herbert A. 1954. “Bandwagon and Underdog Effects and the Possibility of Election Predictions.” *Public Opinion Quarterly* 18:245–53.
- Sober, Elliott. 2009. “Absence of Evidence and Evidence of Absence: Evidential Transitivity in Connection with Fossils, Fishing, Fine-Tuning, and Firing Squads.” *Philosophical Studies* 143:63–90.
- Vetterling, Mary K. 1976. “More on Reflexive Predictions.” *Philosophy of Science* 43:278–82.