The Evidence-Observation Distinction in Observation Selection Effects

Word Count: 4955

Abstract

Previous discussions of observation selection effects (OSEs) have ignored the distinction between observation and evidence. Evidence for a hypothesis, I argue, is distinct from the observation of that evidence. This shows that the fact that evidence is unobservable does not entail that the evidence does not obtain. What is required for an OSE is that evidence is guaranteed, not that counter-evidence is unobservable. With the evidence-observation distinction in hand, apparent counterexamples fail. I then show that observer perspective can change whether or not an agent is subject to an OSE, even when knowledge is shared between perspectives.

1 Introduction

In this paper I defend the entailment model of observation selection effects (OSEs). This simple model states that when background conditions, in conjunction with the hypotheses under consideration, entail evidence E, then E does not favor either of the

hypotheses.[1] However, this model for understanding how OSEs work has been subject to several apparent counterexamples, including FIRING SQUAD and a modified version of Eddington's Fish example.

What these, and other discussants, have presupposed, however, is that evidence and observation of evidence are equivalent. I argue that evidence and observation of evidence are importantly distinct. Evidence is a fact or state of affairs, about the world, that obtains. But the fact that a state of affairs obtains does not guarantee that it is observed. In many cases evidence obtains, but we do not, or cannot, observe it. With this distinction clarified, the entailment model is shown to handle the apparent counterexamples deployed against it.

I then present an example from Francis Bacon which shows the surprising result that observer perspective can change whether or not an agent is subject to an OSE, even when both perspectives share all the same knowledge.

2 Eddington's Fish and the Entailment Model

Discussions of observation selection effects (OSEs) rightly begin with an example adapted from Eddington's 1939 *Philosophy of Physical Science*. (Eddington 1939)[2] In this example, FISHING, we imagine a biologist attempting to distinguish between two

[1]It is not necessary that selection biases, more generally, require that the background and hypotheses *entail* the evidence. For a more general discussion of the phenomena of selection bias see Berkson's Paradox.

[2]This version due to Sober 2003, 41f.

hypotheses:

> L: All fish in this pond are longer than 10 inches.
>
> S: Half of the fish in this pond are longer than 10 inches, the other half are shorter than 10 inches.

The biologist then observes this evidence:

> E: All fish caught in this pond were longer than 10 inches.

How should we evaluate the hypotheses in light of E? One way would be to express the relationship between the evidence and hypotheses would be through the Law of Likelihood:

> Law of Likelihood: E is evidence for H_i over H_j iff $pr(E|H_i) > pr(E|H_j)$

In this case, since all the fish caught were larger than 10", and this is more likely if L is true rather than S, we get the following inequality:

$$pr(E|L) > pr(E|S)$$

Thus, by the law of likelihood, E is evidence for L over S.

But now a further fact about how the observations were made is revealed:

> N: The net used to catch fish in this pond always catches fish, if it can, but it can't catch fish smaller than 10 inches because of the size of the holes in the net.[3]

[3]Most authors have given N as, "The net used can't catch fish smaller than 10 inches

Now E no longer appears to be evidence for L over S because:

$$pr(E|L \wedge N) = pr(E|S \wedge N) = 1$$

The simplest explanation for what has happened is that,

$$N \wedge L \vDash E \text{ and } N \wedge S \vDash E$$

Since N, in conjunction with either L or S, ensures that any fish caught will be larger than 10", it entails E. The alternative, $\neg E$, that some fish caught in this pond were shorter than 10", is ruled out.

In general:

> If, for background conditions B and hypotheses H_1 and H_2, $B \wedge H_1 \vDash E$ and $B \wedge H_2 \vDash E$, then E is not evidence for H_1 over H_2.

Since it is necessary to consider background conditions, we must supplement the Law of Likelihood with a total evidence requirement which makes it explicit that we must take these conditions into account:

> Law of Likelihood*: E is evidence for H_1 over H_2 with respect to background conditions B iff $pr(E|H_1 \wedge B) > pr(E|H_2 \wedge B)$

This model of observation selection effects is the one I clarify and defend in the rest of the paper. In the next section, I explain the FIRING SQUAD counterexample to the

---

because of the size of the holes in the net," but then need to build in a "catch" condition so that $N \vDash E$.

entailment model. I then show that, once we distinguish evidence from observation, the entailment model gives the right result for FIRING SQUAD after all.

3 Evidence, Observation, and Firing Squads

3.1 Firing Squad

The first objection to the simple model above is FIRING SQUAD (Sober 2003, 44f.). In this example, imagine a prisoner faced with a firing squad. Two hypotheses are being considered:

Aim: The firing squad is aiming at the prisoner.

Avoid: The firing squad is aiming to avoid hitting the prisoner.

After the shots are fired and the smoke has cleared, the prisoner makes her observation:

Alive: The prisoner is alive.

Since it is unlikely that the prisoner would still be alive, were the executioners aiming, but quite likely that she would be alive, if they were avoiding her, it seems that:

$$pr(\text{Alive}|\text{Avoid}) > pr(\text{Alive}|\text{Aim})$$

Thus her survival seems to be evidence that the executioners were aiming to miss.

But, if the entailment model formulated above is correct, there appears to be an OSE at work in FIRING SQUAD. The way the evidence was gathered guarantees that, if the prisoner observes anything, she is guaranteed to observe that she is alive. The background necessary for the prisoner to observe that she is alive is:

Survive: The prisoner survives.

Survive clearly entails Alive, so:

$$pr(\text{Alive}|\text{Aim} \wedge \text{Survive}) = pr(\text{Alive}|\text{Avoid} \wedge \text{Survive}) = 1$$

But this will not do. It seems clear that the prisoner would be correct in taking her survival to be evidence that the guards aimed to miss. What has gone wrong?

3.2 Evidence and Observation

In order to understand what has gone wrong, we must give an account of the difference between observation and evidence. These two concepts have frequently been conflated. From an internalist perspective it is often assumed that observations are the only things which could possibly be *used as* evidence. Evidence, after all, must be accessible to the agent and therefore must be a mental state (or similar). The only candidate for evidence about the world, then, is our phenomenal experience of it—namely observations.

But there is more to the story. If we take observations to be evidence then we can, classically, only have two evidential states with respect to E: either we have the evidence E or we do not. But once we move outside of the agent and into an external world, rich in evidence, that is not how observations *of* evidence work. No doubt we *do* need to observe in order to incorporate evidence, but that is not what evidence is.

Evidence is, on my account, a state-of affairs or fact about the world. It can be, and often is, independent of our knowledge, awareness, or observation of it. As a state-of-affairs it either is, or is not, the case. Thus for any state-of-affairs E, either E or $\neg E$.

Observation of that E, on the other hand, is distinct from E. While E or $\neg E$ is necessary, our observation of E or $\neg E$ is not. We might observe E, observe $\neg E$, or we might fail to observe anything, with respect to E.

The test of whether some claim is an observation claim or an evidence claim will, then, as a first pass, be whether it obeys excluded middle. If the claim about φ does obey excluded middle, then it is an evidence claim. If it does not and there is a third option—failure to observe φ—then it cannot be a claim about evidence and must be a claim about observation.

Here's an example to help understand this distinction: An astrobiologist seeks evidence of life on other planets, such as O_2 concentrations in the atmosphere. There is a fact-of-the-matter about the O_2 concentration. It is there whether anyone ever knows about it. This evidence, however is difficult to observe. It may never be observed. But this does not mean that a high O_2 concentration is not evidence for life. It merely means that our observation of this evidence is contingent.

As a first pass at giving a logic of observation, let me propose the observation operator '$\bigcirc$'.[4] $\bigcirc_\alpha\varphi$ should be read as the tenseless claim that α observes φ. One might plausibly know $\bigcirc_\alpha\varphi$ without being the agent α and without observing φ oneself. Crucially, $\bigcirc$ will have the following properties:

1. $\bigcirc_\alpha\varphi \vDash \varphi$

2. $\neg \bigcirc_\alpha \varphi \nvDash \neg\varphi$

[4] $\bigcirc$ may be understood as similar to the epistemic modal operator K and will have similar properties, though developing the connections will require further work.

Thus observation is factive; observation of some evidence E, by an agent, α, entails E (or alternatively, entails that E is the case). But failure to observe E does not entail that there is no evidence E (or alternatively, that E fails to obtain). This distinction is, I think, easily understood, but easily overlooked.

The same distinction between observation and evidence may be elucidated in terms of conditionalization upon learning new evidence:

> Conditionalization: for any time t_i and later time t_j, if proposition E represents everything the agent learns between t_i and t_j and $pr_i(E) > 0$, then for any H, $pr_j(H) = pr_i(H|E)$ (Titelbaum 2015, 92).[5]

Note that conditionalization is put in terms of *learning* E, not in terms of whether E or $\neg E$. Just as with observation, failing to learn E between t_i and t_j does not entail $\neg E$. One does not update upon failing to learn anything.[6] Thus failing to observe E is equivalent to learning nothing between t_i and t_j.

Conditionalization tells us that Bayes' rule is about what one would do *if* one had the evidence, not simply what one believes already. It sets a norm that, if one had some evidence, one ought do such-and-such with that evidence. It does not guarantee that one has that evidence. This diachronic norm forces us to consider the act of learning—observing—new evidence. Rather than just considering evidence as a static phenomenon, the observation of evidence causes a change in our epistemic state. Thus

---

[5]Other updating norms are, of course, possible, such as Jeffrey conditionalization.

[6]Note, however, that one might learn, between t_i and t_j, that one has (or has not) made an observation. φ may be a complex statement which contains $\bigcirc$.

we should be considering two states: the state prior to receiving the evidence, which is when the law of likelihood tells us how we *should* react to the evidence, and the state after receiving the evidence, which is when we *do* incorporate that evidence.

Since Bayesians posit this tight link between synchronic conditional credences and posterior diachronic credences, how one ought to update on evidence is built in from the beginning. One can, therefore, evaluate how one ought to react upon learning E without having to observe E. Similarly, one retains this judgment about how one should evaluate E whether or not one does, or even can, observe E. Thus the evidence and observation distinction is built into the Bayesian approach.

This parallel between evidence-observation and conditionalization will be a useful heuristic going forward.

3.3 Resolving Firing Squad

Let's take a closer look at FIRING SQUAD to see how the observation-evidence distinction is relevant.

First, let us use conditionalization to review the situation. Recall that each of Aim and Avoid, in conjunction with Survive entail Alive, thus Alive seemed to provide no evidence for Avoid over Aim. But as we are considering how evidence is used to update beliefs, we should consider the prisoner's epistemic situation before and after the shots are fired.

It may be easiest to consider the perspective of a bystander to FIRING SQUAD. This bystander will not be subject to an OSE. Her life and future observations are not threatened by the executioners' guns. Thus he can legitimately have this likelihood

argument in mind at t_0:

$$pr(\text{Alive}|\text{Avoid}) > pr(\text{Alive}|\text{Aim})$$

Thus, upon learning, between t_0 and t_1, whether or not the prisoner is alive, he can update on that evidence. Both Alive and ¬Alive are possible states of affairs that come into being between t_0 and t_1. He learns something new when the smoke clears and the prisoner has survived.

The prisoner is in exactly the same evidential position at t_0, prior to the volley.[7] Just as the bystander is in suspense as to whether Alive or ¬Alive, the prisoner too is in suspense. She does not know what will happen. She knows that she will not observe that she is not alive, but this does not rule out the possibility that she does not survive. Thus, when the smoke clears, she learns something new and surprising—she is alive! She then updates on the information learned between t_0 and t_1, increases her credence in Avoid, and lowers her credence in Aim.

Let us put this in terms of observation and evidence. What will the prisoner's background conditions include? Survive (the prisoner survives) will not be in the prisoner's background conditions. Survive would presuppose that the only evidence possible is Alive. But this is not the case. Alive and ¬Alive are both possible states of affairs. What will be among her background conditions is Survive*:

---

[7]Contra Sober who argues that, because the prisoner cannot make the observation that she is not alive, the prisoner and bystander are in different evidential positions (Sober 2003, p. 46, 50n20f.).

Survive*: I will not observe that I am not alive.[8]

In symbols, for prisoner p:

$$\neg \bigcirc_p \neg \text{Alive}$$

This states that the prisoner, p, will not observe that she is not alive. Thus the correct likelihood argument will be:

$$pr(\text{Alive}|\text{Avoid} \wedge \neg \bigcirc_p \neg \text{Alive}) > pr(\text{Alive}|\text{Aim} \wedge \neg \bigcirc_p \neg \text{Alive})$$

Recall that $\neg \bigcirc \text{Alive} \not\models \neg \text{Alive}$. Survive* does not entail Alive. Survive* thus makes no difference to the likelihood argument. The fact that the prisoner will not *observe* that she is alive does not entail that she will not *be* alive. Thus, since it is not guaranteed that she survives, her survival can count as evidence that the soldiers were not aiming to hit her.

Once we recognize the role that the evidence-observation distinction is making, the entailment model gives the correct result for Firing Squad after all.

3.4 The No-Observation Objection

But this response leads to another objection to the entailment model. If we are correct about how the entailment model ought to respond to Firing Squad, then it seems as though our intrepid ichthyologist is also not guaranteed to observe that there are large fish in the net.

[8]Survive* is equivalent to Weisberg's S': If I observe whether I survive, I will observe that I survive (Weisberg 2005, 816).

The objection is this: it is possible that the biologist dies, or simply fails to return to the pond to look in her net, thus the evidence is no more guaranteed for her than for the hapless prisoner above. The biologist knows that she will not observe anything other than large fish, but this does not entail that she will observe large fish. In the same way, the prisoner knows she will not observe that she is not alive, but this does not entail that she will be alive. If this is right then the biologist's background conditions do not seem to entail E (Weisberg 2005, 817). Since the evidence is not entailed, there will be no OSE and the fish in the net *are* evidence for L over S.

In order to answer this we must be careful to specify what the evidence is in FISHING and what is in the biologist's background conditions. Recall what N says:

> N: The net used to catch fish in this pond always catches fish, if it can, but it can't catch fish smaller than 10 inches because of the size of the holes in the net.[9]

And this, in conjunction with either L or S, does entail E, which states:

> E: All fish caught in this pond were longer than 10 inches.

Thus the biologist knows, given N, without needing to observe the evidence, what the evidence is: there are large fish in the net. N therefore entails, not that she observes large fish ($\bigcirc E$), but the evidence (E) itself: All fish caught in the pond are larger than 10". The biologist will, of course, also know that she will not observe anything other than large fish ($\neg \bigcirc \neg E$), but this is irrelevant as she knows, without observing E, that E obtains.

[9]See footnote 3.

Conditionalization can similarly explain the problem. If the biologist dies, or otherwise fails to return to the net, she learns neither E nor $\neg E$. Thus she cannot change her credences $pr(L)$ or $pr(S)$ on the basis of E or $\neg E$. One cannot update without observing evidence!

The No-Observation objection seems to depend on one of two confusions. First, we might think that, in order to generate an OSE, $\bigcirc E$ must have been in the biologist's background conditions. Second, we might have thought that $\neg \bigcirc \neg E \vDash E$. But both of these assumptions are incorrect and both depend on a conflation of evidence and observation.

Sober seems to make this first mistake, stating: "If you fish with Eddington's net, you are guaranteed to observe that the net contains fish that all are over 10 inches long" (Sober 2009, 77). If we mistake observation for evidence then it is easy to assume that N is equivalent to $\bigcirc E$. That is, the fact that the net will contain large fish is the same as the claim that we will observe that the net contains large fish. But of course, as the no-observation objection shows, the biologist is not guaranteed to observe anything! The net, however, is guaranteed to contain fish over 10".

Second, we might have thought that $\neg \bigcirc \neg E \vDash E$, and since $\neg \bigcirc \neg E$ is in the biologist's background, that is the reason E is entailed. But this is, once again, to think that failure to observe small fish is an observation of large fish.

The lesson here is that, if we were already inclined to think that evidence must be observed in order to make any difference to our arguments, we will think that there is no harm in using 'evidence' and 'observation' interchangeably. Not observing E will sound a lot like observing $\neg E$. Similarly, we might argue that observation entails evidence, thus there is no harm in using them interchangeably. But observation statements are not

equivalent to evidence statements and thus there is harm in conflating evidence and observation. We will think that necessitated evidence necessitates observation or that failure to observe is itself an observation. This is a mistake. What is the case in FISHING is that the *observation* that there are large fish in the net is not guaranteed, but the *evidence*—the fact that there will be fish in the net—is guaranteed. The fact that evidence can be guaranteed—and we can know that it is guaranteed—makes it such that evidence can make a difference to OSEs, even without being observed.

To sum up, in FISHING the evidence is guaranteed, but the observation of it is not. In FIRING SQUAD neither the evidence nor the observation is guaranteed. The objection was that observation is not guaranteed, and this is true. But I've argued that it is the *evidence* that must be guaranteed by background conditions, not the observation of it. Thus the biologist in FISHING is subject to an OSE, while the prisoner in FIRING SQUAD is not.

4 Survivor Bias and the Power of Prayer

The evidence-observation distinction allows us to see a further consequence of observation selection-effects: parties can differ in what their evidence supports, even if both parties know all the same facts. To illustrate this, let's start with a case of shipwreck survivors and the power of prayer given by Sir Francis Bacon in *Novum Organum*:

> It was well answered by him who was shown in a temple the votive tablets suspended by such as had escaped the peril of shipwreck, and was pressed as to whether he would then recognize the power of the gods, by an inquiry, But where are the portraits of those who have perished in spite of their vows?

(Bacon 1620/2000, XLVI)

There is an OSE at work here. Clearly, only survivors give testimony (in the form of votive tablets). Those whose prayers were not answered did not survive to give testimony. Thus it is guaranteed that the evidence of the votives will always be from sailors who prayed and survived. Bacon's story is well explained by the entailment model of OSEs. A background condition for the testimony is the survival of the testifier. Thus the background condition entails that a sailor who leaves a votive must have survived, no matter the efficacy of prayer.

Now, to see the surprising result, let's consider a hypothetical conversation between a single sailor and the skeptical temple visitor, call it SHIPWRECK:

> Sailor: As the ship was sinking I wasn't sure I was going to make it. So I prayed and, lo and behold, I was saved! Surely, as I might not have been saved, my survival provides evidence that my prayer was effective.
>
> Visitor: I'm very happy you survived, but I'm sorry to say that your survival is no evidence that your prayer was effective. After all, if you hadn't survived, you wouldn't be standing here telling me your story. You couldn't have told me that you prayed and weren't saved, thus, because the fact that I met you entails that you survived, it does not tell me anything about whether prayer is effective.
>
> Sailor: But you must admit that *I* might not have been saved, and that's all that the argument requires. It's not the fact that we met that matters, but that I survived when I might not have.
>
> Visitor: True, the fact that we met doesn't matter to *your* argument. But

that fact matters very much to my argument. You must admit that I wouldn't know anything about your prayer or your shipwreck without your being able to tell me. I am not surprised that, given I met someone who was in a shipwreck, I met one of the survivors. I could not possibly have met someone who did not survive.

Sailor: Then it seems we're at an impasse. I can tell you everything about my situation and it doesn't change the fact that you could only have learned these things by meeting a shipwreck survivor. You can tell me all about the fact that no other evidence was possible for you and it changes nothing for me. How strange that we agree on all the facts, but differ in what our evidence supports.

What has happened here? Unlike the prisoner and bystander in FIRING SQUAD, in which neither of the parties was subject to an OSE, in SHIPWRECK the sailor is not subject to an OSE, while the visitor is. Let's take a closer look at the sailor and visitor's likelihood arguments:

Let our two hypotheses be:

Effective: Prayer is efficacious.

¬Effective: Prayer is inefficacious.

And the sailor's evidence:

Alive: Sailor s is alive.

Formally, we may now give the sailor's argument as:

$$pr(\text{Alive}|\text{Effective} \wedge \neg \bigcirc_s \neg\text{Alive}) > pr(\text{Alive}|\neg\text{Effective} \wedge \neg \bigcirc_s \neg\text{Alive})$$

The sailor could not have observed that she did not survive, but as this does not entail that she does survive, it is no challenge to the argument, according to the entailment model. She is in an analogous position to the prisoner in FIRING SQUAD.

The visitor's situation is crucially different in two ways. First, there is an additional relevant fact in her background conditions:

> Meet: The visitor, v, meets a sailor, σ, who has been shipwrecked, but v cannot meet sailors who did not survive shipwreck.

This is parallel to the fact N in FISHING. Just as N guarantees that there will be fish in the net, the visitor does meet a sailor who has been shipwrecked. And just as N guarantees that the fish caught will be larger than 10", the visitor can only meet sailors who have survived shipwreck.

But note that, while Meet does guarantee that the visitor meets a sailor, it does not guarantee that he meets this particular sailor, s. The claims in N and Meet *describe* the evidence, but do not mention the particular fish that was caught or the particular sailor who was met. This leads us to the second difference between the sailor and the visitor's situations: the visitor does not care which particular sailor he meets. In the visitor's case, then, the evidence is:

> Alive*: The sailor, σ, that v meets, is alive.

The evidence, for the visitor, is not that some particular sailor, s, survived, but that a

sailor, σ, survived and is met. Similarly, in FISHING, no particular fish was guaranteed to be caught, but the fish that was caught was guaranteed to be large. We now have all the necessary ingredients to give the visitor v's argument:

$$pr(\text{Alive*}|\text{Effective} \wedge \text{Meet}) = pr(\text{Alive*}|\neg\text{Effective} \wedge \text{Meet}) = 1$$

The two sides of the argument are equal because Meet $\vDash$ Alive*. Thus there is an OSE for the visitor.

Note that this is different from the bystander's position in FIRING SQUAD. In that case the evidence was the same for both prisoner and bystander. Prisoner P's survival was what mattered to both the prisoner and bystander in that case. Thus the evidence that the prisoner survived was not guaranteed.

The fact that the visitor is subject to an OSE, while the sailor is not, remains true when the two meet and share information. It is not relevant to the sailor that Meet be known, because Meet only entails that v meet some sailor σ, not that she meet sailor s. Similarly, the fact that sailor s survived is crucial to the sailor's argument—if she had not survived, Effective would not be supported. But the fact that sailor s survived is irrelevant to the visitor's argument—if sailor s had not survived, the visitor would have met some other sailor who had survived. Thus while the two share all the same facts, they do not take those facts to provide equal support for the efficacy of prayer.

This surprising result follows because the arguments that the sailor and visitor make are different and use different statements of the evidence to come to their respective conclusions. The sailor's argument uses the *de re* claim that she survives, while the vistor's argument uses the *de dicto* claim that someone survives. More needs to be said

about whether this can result from a difference in the facts known by the agents. It is at least plausible that the difference rests in a difference in known facts such as, "I am subject to an OSE," said by v, which differs in content from "v is subject to an OSE." Such facts will be unsharable by the agents.

5 Conclusions

In this paper I have distinguished observation from evidence and shown how this can help clarify when there is, or is not, an OSE. I have given three kinds of scenarios:

1. Firing Squad Bystander: Can observe E and can observe $\neg E$. No OSE.
2. Firing Squad Prisoner and Shipwreck Sailor: Cannot observe $\neg E$, but E is not guaranteed. No OSE.
3. Shipwreck Visitor and Fishing: E is guaranteed. OSE.

These three possibilities make it clear that the unobservability of evidence is necessary, but not sufficient for an OSE.

Although I suggest the development of a modal treatment of observation, there is a quick heuristic for distinguishing an evidence claim from observation claim. If the claim about φ does obey excluded middle, then it is an evidence claim. If it does not and there is a third option—failure to observe φ—then it cannot be a claim about evidence and must be a claim about observation.

Finally, observer perspective seems to matter in cases such as the two Shipwreck scenarios. Whether this results from an unshared—or unsharable—fact or results from the accessibility of arguments to agents is an avenue for further exploration.

References

Bacon, Francis. 1620/2000. *The New Organon.* Edited by Lisa Jardine and Michael Silverthorne.

Eddington, Arthur Stanley. 1939. *The Philosophy of Physical Science.* CUP Archive.

Sober, Elliott. 2003. God and Design: The teleological argument and modern science. In *God and design: the teleological argument and modern science,* edited by Neil A Manson. Routledge.

———. 2009. Absence of Evidence and Evidence of Absence: Evidential transitivity in connection with fossils, fishing, fine-tuning, and firing squads. *Philosophical Studies* 143 (1): 63–90.

Titelbaum, Michael. 2015. Fundamentals of Bayesian Epistemology. *Unpublished manuscript.*

Weisberg, Jonathan. 2005. Firing Squads and Fine-Tuning: Sober on the design argument. *The British Journal for the Philosophy of Science* 56 (4): 809–821.