

Tell me you love me: bootstrapping, externalism, and no-lose epistemology

Michael G. Titelbaum

Published online: 2 April 2010
© Springer Science+Business Media B.V. 2010

Abstract Recent discussion of Vogel-style “bootstrapping” scenarios suggests that they provide counterexamples to a wide variety of epistemological theories. Yet it remains unclear why it’s bad for a theory to permit bootstrapping, or even exactly what counts as a bootstrapping case. Going back to Vogel’s original bootstrapping example, I note that an agent who could gain justification through the method Vogel describes would have available a “no-lose investigation”: an investigation that can justify a proposition but has no possibility of undermining it. The main suggestion of this article is that an epistemological theory should not permit no-lose investigations. I identify necessary and sufficient conditions for such investigations, then explore epistemological theories that rule them out. If we want to avoid both skepticism and no-lose investigations, we must eschew either Closure or epistemic externalism.

Keywords Bootstrapping · Closure · Epistemic externalism · Skepticism

1

Jonathan Vogel (2000) describes a “bootstrapping” counterexample to reliabilism. For our purposes, we can take the target view to be:

Reliabilism: A belief is justified just in case it is formed by a reliable process.

Vogel’s counterexample involves Roxanne, a woman who forms beliefs about the contents of her car’s gas tank by reading its gas gauge. Roxanne has not bothered to ascertain whether the gauge is a reliable indicator of the tank’s contents. In point of

M. G. Titelbaum (✉)
Department of Philosophy, University of Wisconsin-Madison, 5185 Helen C. White Hall,
600 North Park Street, Madison, WI 53706, USA
e-mail: titelbaum@gmail.com

fact, the gauge is reliable, but Roxanne has neither justification to believe this is true nor justification to believe it is false.

Over the course of many days, Roxanne reads the gauge repeatedly. On each occasion, she reasons as follows:

- On this occasion, the gauge reads “F”.
- On this occasion, the gas tank is full.
- On this occasion, the reading on the gauge matches the contents of the tank.

The gauge does not read “F” on every occasion; the example shows a pattern of reasoning that will yield the same third proposition each time. The crucial point is that according to reliabilism, Roxanne is justified in believing each of the above propositions (or their analogues) on each occasion. Because Roxanne’s perception is reliable (we may suppose), the belief that the gauge reads “F” is justified once she looks at the gauge. Because the gauge is reliable, reading the gauge gives Roxanne justification for the belief that the tank is full. From these two justified beliefs, Roxanne may deduce the justified belief that the reading on the gauge matches the tank’s contents.¹

After Roxanne has engaged in this pattern of reasoning many times, she can deduce:

- I have read the gauge many times, and each time its reading has matched the contents of the tank.

By induction, Roxanne then forms the justified belief that:

- The reading on the gas gauge always matches the contents of the tank.

From which a quick deduction yields:

- The gauge is reliable.

Roxanne started out without any justification to believe that her gas gauge is reliable. Yet if reliabilism is correct, Roxanne can gain such justification simply by reading the gauge over and over and working through Vogel’s bootstrapping reasoning. This looks like a serious problem for reliabilism.

One might think that the problem comes from Roxanne’s inductive reasoning step—perhaps the reliabilist can avoid bootstrapping by denying that induction from justified premises always yields justified conclusions.² However, we can easily construct bootstrapping examples with no inductive step. Suppose that for some reason Roxanne knows before ever checking her gas gauge that

- The gauge is either reliable or anti-reliable.

¹ Vogel puts this conclusion in terms of the gauge’s “reading accurately.” To me, that suggests a causal connection that Roxanne is not in a position to deduce from the first two premises. My gas gauge is not accurately reading the current female composition of Harvard’s freshman class if my gauge reads 1/2 and one half the class is female. So I have rephrased Vogel’s characterization of Roxanne’s reasoning in terms of more neutral “matching” terminology.

² The thought process that led me to this article began with a conversation in which Kenny Easwaran made just such a suggestion.

where being “anti-reliable” means the reading on the gauge *never* matches the contents of the tank. As soon as Roxanne observes the gauge *once* and deduces that

- On this occasion, the reading on the gauge matches the contents of the tank.

she can be justified in believing the gauge is reliable. No inductive step is required. All we need to get the bootstrapping going are the ability of the gauge to produce justified beliefs, the ability of perception to produce justified beliefs, and

Closure: If each premise in a set is justified, any proposition jointly entailed by the set is justified as well.

(In symbols: $[Jp_1 \& Jp_2 \& \dots \& Jp_n \& (\{p_1, p_2, \dots, p_n\} \models q)] \rightarrow Jq$)

The importance of bootstrapping may seem mitigated by the fact that no one in the literature endorses reliabilism as I have described it.³ Goldman (1976) argued briefly for the view, only to reject it for lack of a defeaters condition. But adding a reliabilist-friendly defeaters condition to the reliabilism I’ve described won’t save the view from bootstrapping; it would be contrary to the spirit of reliabilism to claim that Roxanne’s lack of initial evidence that the gauge is reliable counts as a *defeater* for the proposition that it is. Moreover, (Cohen 2002) shows that a variety of epistemological theories besides reliabilism give rise to bootstrapping examples. Cohen concludes that a theory will permit bootstrapping just in case it allows a source to give an agent justification without that agent’s being antecedently justified in believing the source is reliable.⁴ (Reliabilism, for example, allows the gas gauge to justify beliefs about the tank’s contents without Roxanne’s having antecedent justification to believe the gauge is reliable.)

Cohen defends his conclusion by working through a number of epistemological theories with the relevant feature and generating bootstrapping examples for each. Yet this falls short of a general argument that he has identified the correct class of views; for instance, it leaves open the possibility that an epistemology might have bootstrapping problems even if it requires reliability information before an agent can gain justification from a source.⁵ As Cohen admits (p. 321), his efforts are hindered by our lack of a precise characterization of the problem cases—almost no attention has been paid to defining necessary and sufficient conditions for “bootstrapping.”

Insufficient attention has also been paid to the question of *why* it’s a bad thing for a theory to permit bootstrapping. This question is particularly important in light of Van Cleve’s argument (2003) that any theory immune to bootstrapping will be susceptible to skepticism. If we are forced to choose between bootstrapping and

³ I am grateful to an anonymous *Philosophical Studies* referee for suggesting that I explicitly address this concern.

⁴ Cohen’s discussion, like Vogel’s, actually concerns knowledge rather than justification. I have adapted his conclusions to the justificatory case.

⁵ Cohen is also rather vague on whether Closure endorsement is required for an epistemological theory to generate bootstrapping. At Cohen (2002, p. 320) he writes, “We have seen two related ways in which the problem of easy knowledge arises for theories that allow for basic knowledge—by the closure principle, and by bootstrapping.” The implicit contrast between *two* ways suggests that bootstrapping susceptibility does not depend on Closure. Yet Cohen’s bootstrapping examples all involve justification-transferring deductive inferences, and the basic knowledge views he surveys embrace Closure.

skepticism, we need to know what costs a theory incurs by permitting bootstrapping. The bootstrapping literature largely trusts our intuitive rejection of bootstrapping processes. It is sometimes pointed out⁶ that bootstrapping allows an agent to use a process in establishing its own reliability. But it's unclear that such circularity is always vicious—doesn't some of our justification for believing perceptions involve past deliverances of perception? It is also sometimes suggested that bootstrapping allows Roxanne to conclude that her gauge is reliable without performing any independent checks, such as correlating the gauge's readings with dipstick observations of the tank's contents. But very little is said about what "independent" comes to here—what conditions must the dipstick observation meet to clear the bar?

Perhaps there are many bad things about bootstrapping; perhaps there is no *the* thing that is wrong with an epistemological theory that permits it. But let me suggest *one* problem a theory has if it makes bootstrapping possible. There's an old idea in epistemology that some risk must attach to any reward: If an investigation can't undermine a conclusion, it can't support it either. Now consider Roxanne's situation. Suppose that the true epistemology allows Vogel's bootstrapping procedure to yield justification if the gauge is reliable, and suppose that Roxanne (a seasoned epistemologist) knows this. Suppose further that before she makes any observations of the gauge, Roxanne can predict what kinds of observations she'll make and what kinds of reasoning she'll use. She doesn't know precisely what reading the gauge will give on each occasion, but she knows that whatever the reading, she will conclude that it matches the contents of the tank. Finally, suppose Roxanne knows in advance that *which* precise gas levels are indicated at which times won't yield any clues as to the gauge's reliability; perhaps her car will be driven or partially filled up by her brother each day before she checks its gauge so that there will be no pattern in the readings.

Roxanne knows that after she has made many observations, she will believe that the gauge is reliable. Her epistemological knowledge tells her that if the gauge is indeed reliable, her belief that it is will be justified. On the other hand, Roxanne knows that nothing in the course of her observations will give her any justification for believing the gauge is not reliable. Considering the proposition p that the gauge is reliable, Roxanne knows that if p is true, her investigation will justify it for her, but if p is false, the investigation will provide no evidence against it. And Roxanne knows all this before the investigation of p begins, when by stipulation she lacks justification for believing either p or $\sim p$.

In short, an epistemological theory that allows Roxanne to bootstrap permits a *no-lose investigation*. But the true theory of justification shouldn't permit no-lose investigations.

2

This article's main suggestion is that true epistemological theories do not permit no-lose investigations. Later our goal will be to identify classes of theories that meet this *desideratum*. First, however, we need to clarify the claim in a number of ways.

⁶ By Fumerton (1995), for example.

Up to this point I haven't been precise about the notion of justification in play. This is because I think the points in this article work for a number of notions of justification, as well as justification-like notions such as having warrant for, having support for, having evidence for, and even being in a position to know. Anywhere the word "justification" appears, or its abbreviation "*J*", you should feel free to substitute any of those notions—subject to a few conditions.

First, I will assume that knowledge entails justification. I will not assume that knowledge is justified true belief. I will also not assume that a proposition's being justified for a particular agent entails that the agent believes the proposition. That is, I am working with a notion of *propositional* rather than *doxastic* justification. We attribute doxastic justification with locutions like "the agent's believing *p* is justified;" an agent has doxastic justification for *p* only if the agent believes *p* and does so for good reason. Propositional justification is evoked by "*p* is justified for the agent;" *p* is propositionally justified for an agent *whenever* that agent has adequate justificatory resources for *p*, whether the agent avails herself of those resources or not.⁷

Our notion of justification may be global or local; it may require surpassing some evidential threshold or it may not. Any of the following could be our notion of justification: has *some* evidence for *p*, has *prima facie* justification for *p*, has *pro tanto* reason to believe *p*, has all-things-considered justification for *p*, is in a position to know *p*, etc.

Finally, I am going to overlook the distinction between an investigation that removes justification to believe *p*, an investigation that provides justification to *disbelieve* *p*, and an investigation that provides justification to believe $\sim p$. I will describe an investigation that does any of these as "undermining" *p*. In general, I will coarse-grain descriptions of an agent's justificatory situation with respect to a proposition so as to work with just four categories: *p* is justified for the agent, $\sim p$ is justified for the agent, both, or neither. (The "both" option will be available on only some notions of justification.)

So much for clarifying our notion of justification; what about no-lose investigations? Notice that our main suggestion is a conditional with an existential antecedent: Given a particular epistemological theory, *if* there exists a situation that meets the conditions for a no-lose investigation, *then* that theory is false. This suggestion will be plausible only if every case meeting those conditions is epistemologically repugnant.

To that end, we should start by noting that a better name for the class of cases might be "some-win-no-lose investigations" (though we'll stick with the catchier moniker). There's nothing troubling about the existence of an "investigation" of *p* that has no hope of providing justification for *p* or for $\sim p$. We should restrict our attention to cases in which the undermining of *p* is not possible but the justification of *p* is.

Even then, we might think there are investigations that are "no-lose" in some sense but are perfectly permissible by the true theory of justification. Suppose I'm

⁷ Compare the discussion in Kvanvig and Menzel (1990), as well as Goldman's version of reliabilism for "ex ante justification" at the end of his (1976).

about to drill for oil below the spot on which I'm standing, and there is indeed oil just beneath the surface. There's a sense in which this investigation is guaranteed to produce justification for the claim that there's oil beneath me—it will in fact produce such justification. Yet this is a perfectly good investigative process on any plausible theory of justification.

But this is not a no-lose investigation in our sense. When we say that an investigation of p shouldn't be *guaranteed* not to undermine p , the fact that an investigation will justify p in the actual world isn't strong enough to produce the kind of guarantee in which we're interested. This suggests that our conditions for a no-lose investigation might be spelled out using counterfactuals—perhaps a no-lose investigation is one that fails to undermine p not only in the actual world but also in close possible worlds. Yet that move would force us to choose among various theories of counterfactuals, and would also bring in questions about how to identify *this very investigation* across possible worlds. We would quickly find ourselves dealing with the generality problem (about the correct level of description for identifying epistemological processes), which is already a problem for a variety of epistemological views.⁸

A better alternative is to spell out the guarantee in a way that sets aside process descriptions and focuses exclusively on conditions in the actual world. When we require an investigation to have the possibility of undermining p , we should focus not on metaphysical possibility but instead on epistemic possibility for the agent. We should ask whether the agent entertains any epistemically possible worlds in which the investigation undermines p ; that is, we should ask whether the agent knows in advance that undermining is ruled out.

Pursuing this line yields a simple list of necessary and sufficient conditions for a no-lose investigation. Suppose an agent knows at t_1 that between that time and some specific future time t_2 she will investigate a particular proposition (which we'll call p). Her investigation counts as a no-lose investigation just in case the following three conditions are met:

- (1) p is not justified for the agent at t_1 . ($\sim J_1 p$)
- (2) At t_1 the agent knows that $\sim p$ will not be justified for her at t_2 . ($K_1[\sim J_2 \sim p]$)
- (3) At t_1 the agent knows that if p is true, p will be justified for her at t_2 . ($K_1[p \rightarrow J_2 p]$)

The second of these conditions captures the sense in which a no-lose investigation is guaranteed to have no justificatory downside. The first and third conditions provide the possible upside: if p is true, the agent will go from p 's not being justified for her at t_1 to p 's being justified for her at t_2 . Our main suggestion is that, setting aside a few small exceptions to be discussed later, a correct epistemological theory will not allow investigations satisfying all three of these conditions.⁹

⁸ See Conee and Feldman (1998).

⁹ It's important that our definition of a no-lose investigation employs two times, at the first of which the agent lacks justification for the proposition in question. This addresses a discussion in which Vogel takes up something like the suggestion that what's wrong with Roxanne's procedure is that it permits no-lose investigations:

To illustrate this definition of no-lose investigations, consider the following story:

The Court Jester: Noblemen from Italy have arrived in the King of England's court, bringing with them the jester Giacomo. The King has heard a rumor that Giacomo is quite the ladies' man,¹⁰ but the King knows the rumor's source is unreliable and so lacks justification to believe it.

To settle the matter, the King orders the jester to regale the court with tales of his amorous conquests. The King's instructions are very precise: If Giacomo is indeed a ladies' man, the tales are to be true; if not, the jester is to make up false tales that sound convincingly real. The King knows the jester will obey these orders—to disobey is punishable by death, and nearby Italian nobles who know the truth about Giacomo will be happy to expose any disobedience. So the King knows that whether the jester is a ladies' man or not, His Highness will hear nothing this evening that convinces him otherwise.

As the King expects, the jester spends a long evening describing broken hearts left littering the landscape. In fact, Giacomo is a ladies' man and all his tales are true. At the end of the evening, is the King justified in believing this?

An epistemological theory's answer to this question will depend on whether it holds that the King is justified in believing what the jester says.¹¹ If so, then once the King is justified in believing that Giacomo wooed the Lady Gwendolyn, that Giacomo

Footnote 9 continued

"Point. The problem with Roxanne's procedure is that it could not have possibly yielded any other result than the one it did, namely, that the gauge is reliable.

"Counterpoint. It is not clear that the putative defect really is one. The process by which I know I am conscious when I am is surely a reliable one, yet that process could not return a verdict other than that I am conscious." (Vogel 2000, p. 615)

Whatever the "process by which I know I am conscious" is, surely it is available to me whenever I am conscious. So there is no time at which I satisfy the conditions for a no-lose investigation with respect to the proposition p that I am conscious. If I am capable of having any knowledge at a time (and thereby satisfying the second and third conditions), I have justification to believe p at that time and so fail to satisfy the first condition.

The point of our no-lose investigation definition is to characterize a class of objectionable investigations. Vogel's consciousness case doesn't satisfy that definition, while his gas gauge case does. So I think Vogel has sold short the suggestion that Roxanne's procedure is objectionable because it creates a no-lose investigation.

¹⁰ After all, it is the *Italian* court—what better place to court Italians?

¹¹ Depending on the epistemological theory in play and the conditions under which it counts testimony as reliable, we might tweak the story's details so as to make Giacomo's reports clear the bar. For instance, we might imagine that it is a deep-seated fact about Giacomo that he is a ladies' man (perhaps this has to do with his relationship with his mother), so that the only close possible worlds in which his tale of wooing Lady Gwendolyn (say) is false are worlds in which he wooed another lady instead, and reports that exploit truthfully to the King. Or we can imagine that whenever *any* jester arrives at *any* court he is given orders like the King's, and that (unbeknownst to the King) all jesters are ladies' men. In that case, testimony in circumstances similar to Giacomo's will be broadly reliable, which might make a difference to the justification-conferring status of Giacomo's testimony. (Thanks to Ted Hinchman for discussion of these points.)

wooded the Maid Jean, etc., the King will (by Closure) have justification to believe that the jester is a ladies' man.

The point of the story is that this conclusion is absurd. One should not be able to gain justification for a proposition just by ordering up favorable evidence, even if that evidence happens to be true. Any epistemological theory that says one can is incorrect. This is brought out by the fact that the King's investigation, if capable of providing justification, could be worked up into a no-lose investigation. Suppose that according to the true epistemological theory the King's procedure provides him with justification if Giacomo is reliable—and suppose the King knows that. Now consider the King just after he decides what orders to give the jester. The King's source for the rumor is unreliable, so he lacks justification to believe the proposition p that Giacomo is a ladies' man. This satisfies the first condition for a no-lose investigation. The King also knows that the jester is a seasoned performer who will be under serious duress and so will not give any indication that he is not a ladies' man. This satisfies the second condition. Finally, the true epistemological theory tells the King that if the jester is indeed a ladies' man, his reports will provide the King with justification. So the King knows that if p is true, he will wind up with justification to believe p . This satisfies the third condition. The true epistemological theory should not allow the King to gain justification by listening to Giacomo, in part because if the King could do so he could engage in a no-lose investigation.

I will leave it to the reader to verify that if Roxanne is well-informed epistemologically and Vogel's bootstrapping process goes through, it creates a no-lose investigation for her as well.

3

One might think that the possibility of no-lose investigations could be ruled out immediately by something like Bas van Fraassen's Reflection principle. (van Fraassen 1995) The Reflection principle itself will probably not do the job, because it concerns an agent's current and future credences and our criteria for no-lose investigations are not obviously about credences.¹² But there is a principle in the vicinity that might apply to our case:

Epistemic Reflection: Given two times t_1 and t_2 and a proposition p , if the agent has justification at t_1 for the proposition that she will have justification for p at t_2 , then the agent has justification for p at t_1 . ($J_1[J_2p] \rightarrow J_1p$)

Epistemic Reflection is a highly plausible principle¹³ as long as we make allowance for some well-known kinds of exception to Reflection. Arntzenius (2003) argues that Reflection fails when an agent is subject to memory loss or the threat thereof, and such cases will also create exceptions to Epistemic Reflection. For example, suppose I have evidence at t_1 that favors p but also have a defeater for that

¹² Though "The agent would be rational in assigning a high credence to p " might be one of the notions that could play the role of Jp in our no-lose investigation definition.

¹³ Compare the "Evidence for evidence is evidence" principle in Feldman (2006).

evidence. Suppose also that I know I will forget the defeater (but not the evidence) between t_1 and t_2 . I have justification at t_1 to believe that p will be justified for me at t_2 , but this does not give me justification for p at t_1 . One may also argue (via the Sleeping Beauty Problem) that Reflection fails in cases in which p is context-sensitive, or in which some of the evidence relevant to p is context-sensitive even if p itself is not.¹⁴ Such cases will also create exceptions to Epistemic Reflection.

There are also exceptions to Reflection in which the agent suspects she may be irrational at future times. However, these cases do not create exceptions to Epistemic Reflection because Epistemic Reflection concerns the propositions that are *justified* for an agent at various times, not what the agent actually believes at those times.

These exceptions to Epistemic Reflection can also provide exceptions to our claim that the correct theory of justification will not allow no-lose investigations. For a silly example, take the proposition p “There are memory-erasers who want belief in their existence to be justified.” Suppose that at t_1 I have evidence for p but also have a defeater for that evidence (so that I meet the first condition for a no-lose investigation). Suppose further that I know of some specific future time t_2 that I’m not going to get any evidence against p between now and then (so that the second condition is met).¹⁵ Finally, suppose that if p is true the memory-erasers will remove the defeater from my memory so that I have justification to believe in them at t_2 (thereby meeting the third condition). Under our definition, this example involves a no-lose investigation, yet such arrangements will be possible on any epistemological theory that allows for defeated justification.

To avoid memory-loss and context-sensitivity counterexamples, I hereby amend our main suggestion to apply only to cases in which: (1) p and all the agent’s relevant evidence concerning p are context-insensitive; and (2) the agent knows at t_1 that every proposition relevant to p that is justified for her at t_1 will also be justified

¹⁴ See Titelbaum (2008) for more detail, including a discussion of the notion of “context-sensitivity” needed here.

¹⁵ As a side point, it is extremely important that our no-lose investigation conditions concern the agent’s t_1 knowledge about *some specific* future t_2 . That is, the conditions say that the agent knows of some specific future t_2 that such-and-such, rather than that the agent knows there will exist a future t_2 such that such-and-such.

Without a *caveat* like that one can generate counterexamples to Reflection-like principles such as Avoid Certain Frustration, as Hájek shows in his (2005). One might also be able to generate unobjectionable no-lose investigation examples. For instance, take some numerical property and consider the proposition that it is possessed by at least one positive integer. I might investigate this proposition by taking the positive integers in order one at a time and determining whether the integer in question has the property. If some positive integer has the property, my investigation will eventually reveal that to me (supposing I have an indefinite amount of time available), but if no positive integer has the property I will never gain justification for that negative conclusion. Notice, however, that there is no specific future time such that I know of it that if the proposition is true I will have justification to believe it by that time; it is crucial to the example that I do not know of some specific n that if any integer has the property then some integer less than n will.

for her at t_2 .¹⁶ The Roxanne and Giacomo examples either meet these conditions or can be made to do so by slightly rewriting context-sensitive premises.

With the proper caveats in place, Epistemic Reflection should be adopted as part of any correct theory of justification. But that alone won't make a theory immune to no-lose investigations. Epistemic Reflection concerns cases in which the agent has justification at t_1 for the proposition that she will have justification for p at t_2 . But in a no-lose investigation the agent has justification at t_1 only for the proposition that *if p is true* she will have justification to believe it at t_2 . This is insufficient to justify p for her at t_1 by Epistemic Reflection; so Epistemic Reflection does not put our third no-lose investigation condition in tension with the first.

4

Having laid out precise necessary and sufficient conditions for no-lose investigations, we can now ask what kinds of epistemological theories make such investigations impossible.

First, any theory that says no agent is ever justified in believing anything will clearly avoid no-lose investigations.¹⁷ But I take this to be an unappealing option.

Second, one could avoid no-lose investigations by denying Closure. For example, one might have a theory on which an agent's justification for a proposition is always relative to a set of alternatives. If moving from premises to entailed conclusion changes the set of relevant alternatives, justification may not be preserved and Closure may be violated.¹⁸ For example, when Roxanne forms her initial belief that

¹⁶ The second clause of our amendment is stronger than it needs to be just to ban memory-loss cases, but that added strength helps rule out other putative counterexamples to our main suggestion. For example, on some understandings of propositional justification an agent can lack propositional justification for a proposition because she lacks the cognitive capacity to understand how her epistemic resources bear on that proposition. This will cause trouble both for Epistemic Reflection and for our main suggestion in cases in which an agent's cognitive capacities (her space of concepts, her ability to follow complex reasoning, etc.) may diminish between t_1 and t_2 .

To take an extreme example of diminished capacity, suppose Luke is considering the proposition that there is still good in his father and plans to investigate this proposition by confronting Vader. Luke knows that if the proposition is true, he will have justification for it by the end of their confrontation, but if the proposition is false Vader will kill him. Since the latter eventuality will leave Luke dead, he clearly won't have propositional justification for his proposition's negation. So this case satisfies the second condition for a no-lose investigation. Moreover, if we assume that Luke lacks justification for his hopeful proposition before the confrontation, this example meets the first and third no-lose conditions. However, it runs afoul of our stipulation that Luke must know at t_1 that he will retain at t_2 justification for all the propositions that are justified for him at t_1 —if he dies, Luke's propositional justifications will all disappear. (Thanks to Juan Comesaña and Alex Hyun for discussion.)

¹⁷ One might think that no-lose investigations could also be avoided by denying that agents ever have knowledge, or even (perhaps more plausibly) by denying that agents ever have knowledge of the future. After all, our second and third no-lose investigation conditions require the agent to have knowledge at t_1 about what occurs at t_2 . However, the "K"s in those conditions could be replaced by "J"s and all of our arguments would still go through—I have used "K"s rather than "J"s only because epistemic possibility is usually understood in terms of knowledge. It is far less plausible to hold that agents never have *justification* for beliefs about the future.

¹⁸ I am grateful to Jonathan Schaffer for discussions of this view.

the tank is full, the only alternatives under consideration are (1) the gauge reads “F” and the tank is full, (2) the gauge reads “1/2” and the tank is half-full, etc. When she later considers whether the gauge is reliable, the set of relevant alternatives includes cases in which the gauge reading mismatches the contents of the tank. While she had justification to believe that the tank was full relative to the initial set of alternatives, she lacks justification for the proposition that the gauge is reliable relative to this expanded set of alternatives. So Roxanne cannot justifiably infer the reliability of the gauge from her individual justified beliefs in the gauge’s reports.

Disavowing Closure can thwart Vogel’s Roxanne example. But can a theory escape all possible no-lose investigations by denying Closure? The answer appears to be “yes.” All of our no-lose investigation examples involve a proposition p that is equivalent (given the agent’s background knowledge) to the proposition that her epistemic process is reliable. (For example, relative to the King’s background knowledge the jester is a ladies’ man just in case his testimony is reliable.) Now suppose some reliabilist denies Closure, and we try to construct an investigation that is no-lose on his view.¹⁹ The epistemic process in question will have to report on some matter other than its own reliability, and the agent will then have to *infer* proposition p from that report. But if the reliabilist denies Closure, he can deny that p is justified for the agent after the inference, thereby blocking the no-lose example.²⁰

Crispin Wright’s theory (2004) accepts Closure—at least for the justificatory notion he calls “warrant”—but avoids no-lose investigations by another tack. Wright would not grant the reliabilist’s claim that Roxanne can gain warrant for beliefs about the contents of the gas tank just by reading the gauge; for Wright Roxanne would need antecedent evidence that the gauge is reliable. Yet Wright avoids the threat of skeptical regress here by holding that there are some fundamental epistemic processes (such as perception) which we are entitled to accept as reliable without any evidence to that effect. Then why can’t we generate a no-lose investigation whose p concerns the reliability of one of these special epistemic processes? Because Wright believes we are *always* entitled to accept the proposition that such a process is reliable; that entitlement is not earned by any process we go through. So for a p of this sort we will never be able to construct an example that satisfies the first condition for a no-lose investigation; the agent in question will always have warrant for p at t_1 .

¹⁹ It doesn’t matter whether this is an appealing epistemological theory—our goal is to show that it’s *possible* to escape no-lose investigations altogether by denying Closure.

²⁰ To see why the agent’s epistemic process will have to report on some matter other than its own reliability, try to imagine a process that reports directly on its own reliability and then ask yourself under what conditions those reports will count as reliable. (For instance, try to imagine a dashboard light whose job it is to report that all the dashboard lights are functioning reliably.)

Note also that it may be possible to construct a no-lose investigation in which *all* the agent’s inferences are inductive rather than deductive. In that case the Closure-denying reliabilist will also have to deny a parallel principle concerning conclusions inductively supported by justified premises. But I take it the Closure-denier will want to do that anyway.

Another way to maintain Closure but escape no-lose investigations is to endorse

Negative Self-Intimation: If an agent is not justified in believing a proposition, she is justified in believing she is not so justified. ($\sim Jp \rightarrow J[\sim Jp]$)

Negative Self-Intimation is usually accepted as part of a broader position that the justificatory status of a proposition is always accessible to an agent. So adherents of Negative Self-Intimation typically accept Positive Self-Intimation (or the “JJ” principle) as well, according to which an agent has justification to believe she’s justified whenever she is. But in principle Negative Self-Intimation can stand on its own.

One can formally *prove* that Negative Self-Intimation (in the company of Epistemic Reflection and Closure) bars no-lose investigations, but the proof is complicated because it involves reasoning about an agent’s reasoning about what’s justified for her at various times. So I have left the proof to an [appendix](#). Roughly speaking, though, here’s how it works: Suppose for *reductio* that Negative Self-Intimation is true and no-lose investigations are possible. Consider an agent who has arranged an investigation meeting our three conditions for some given p , t_1 , and t_2 .

At t_2 , p is either justified for the agent or isn’t. Suppose for *reductio* that it isn’t. By Negative Self-Intimation, the agent has justification at t_2 to believe that she lacks justification for p . But she also knows that if p is true, she has t_2 justification for p . So by Closure the agent has t_2 justification for $\sim p$. But one of our conditions entails that $\sim p$ is not justified for the agent at t_2 . So we have a contradiction; it must be that p is justified for the agent at t_2 .

At t_1 , the agent can run through all the reasoning in the previous paragraph. So at t_1 the agent has justification to believe that at t_2 she has justification for p . By Epistemic Reflection, the agent then has justification for p at t_1 as well. But one of our conditions for a no-lose investigation was that the agent lacks t_1 justification for p . So we have another contradiction, and we can conclude that given Closure and Epistemic Reflection, Negative Self-Intimation is inconsistent with the possibility of no-lose investigations.

The basic idea of this proof is that if Negative Self-Intimation is true, an agent will always be able to *notice* when she lacks justification for p . In a no-lose investigation, lacking justification for p at t_2 is evidence that p is false. So if Negative-Self Intimation is true the only way to guarantee the agent’s investigation won’t undermine p is to guarantee that that investigation will provide justification for p . But if an investigation is (epistemically) guaranteed to provide justification for p at t_2 , then by Epistemic Reflection the agent already has that justification at t_1 , in violation of our first no-lose condition.²¹

²¹ One might wonder whether the “dogmatist” position defended in Pryor (2000) allows for no-lose investigations. Roughly speaking, Pryor’s proposal is that an agent has automatic *prima facie* justification to believe the deliverances of particular epistemic processes. For example, if it seems to an agent that she sees her hand, she has justification to believe that she has a hand. (The difference with Wright is that Pryor offers no antecedent warrant to believe that perception is reliable.) It seems to me that dogmatism must ascribe to Negative Self-Intimation, at least for epistemologically informed agents and the epistemic processes covered by the view. If an agent fails to seem to see a hand, she is justified in believing that she has failed to seem to see a hand, and therefore justified in believing that she has failed to receive

To summarize the results of this section, an epistemological theory may avoid no-lose investigations by: denying the possibility of justification, denying Closure, allowing agents “warrant for nothing” that particular epistemic processes are reliable, or adopting Negative Self-Intimation. Combinations of these moves will work as well. A number of views, however, will remain in trouble. A Closure-embracing reliabilist, for instance, will grant the possibility of justification but will not give agents either free warrant to accept that their processes are reliable or the ability to detect when such processes are not. So such a view will allow no-lose investigations, as the Roxanne example reveals.

5

We began by asking what’s wrong with bootstrapping, and it’s not clear we’ve answered *that* question. Perhaps many things are wrong with bootstrapping. Perhaps one of them is that bootstrapping generates no-lose investigations. But whatever is wrong with bootstrapping in general, I think we have identified something that goes wrong with any epistemological theory that allows Roxanne to gain justification that her gas gauge is reliable: it creates a no-lose investigation. In general, it is a bad thing if an epistemological theory makes no-lose investigations possible. And while we don’t have necessary and sufficient conditions for a case to qualify as bootstrapping, we have provided such conditions for no-lose investigations. This makes it much easier to identify classes of epistemological theories that avoid no-lose investigations, as we did in the previous section.

It may be objected that we have substituted for the question “What’s wrong with bootstrapping?” the question “What’s wrong with no-lose investigations?” That’s an interesting question as well, but there’s an important contrast: while our intuitive aversion to bootstrapping is a recently-recognized phenomenon, the aversion to all-uptside epistemology that lies behind our rejection of no-lose investigations is much older and better-entrenched. It appeared once in Nozick’s claim (what Nozick (1981) called the “variation condition” and now sometimes goes by “sensitivity”) that an agent knows a proposition only if she wouldn’t believe it were it false; before that it was recognizable in Popper’s (1961) position that a theory can be tested only if it is falsifiable—an idea which in turn has origins as far back as Bacon’s *Novum Organum*.²² Van Cleve can argue that allowing bootstrapping is the price of avoiding skepticism, and being only mildly invested in bootstrapping-avoidance we may entertain that as an acceptable exchange. But something much deeper (in me at least) objects to a view that allows no-lose investigations.²³

Footnote 21 continued

justification in the manner described by the theory. This (at least partial) embrace of Negative Self-Intimation will inoculate dogmatism against no-lose investigations.

²² For references on and a brief discussion of the connection between Popper and Bacon, see Klein (2009). Notice also that Nozick’s tracking theory avoids no-lose investigations by denying Closure.

²³ Compare also our intuitions about observation selection effects—for example, testing the hypothesis that all the fish in a pond are above a certain size by sampling fish from that pond with a net that only catches fish greater than that size.

The possibility of no-lose investigations also has odd consequences. Consider our King again, and imagine that he is someone who values having justified beliefs. He knows that he has no justification for believing the rumor about the jester, but if the true theory of epistemology permits him a no-lose investigation he knows that after talking to the jester he may just (if he gets lucky and the rumor is indeed true) possess such justification. This makes it important to the King to talk to Giacomo, *even though he already knows what information he is going to get*. Even though the King knows what Giacomo is going to say, his face-to-face interaction with the jester has the potential to change the rumor's justificatory status and so is something His Highness will seek out. This strikes me as an odd fetishization of the actual employment of a process whose results are entirely anticipated.²⁴

It will be noted that while our list at the end of the previous section describes *sufficient* conditions for a theory's avoiding no-lose investigations, we haven't shown that making one of the moves on that list is *necessary* if one wants to avoid no-lose investigations. Hopefully our precise definition of a no-lose investigation will some day make a proof of necessity and sufficiency possible. For now let me offer a line of thought that at least suggests that making one of the moves listed is necessary to avoid no-lose investigations. Suppose one allows for the possibility of justification, but doesn't give it away for free in the Wrightian style. Suppose one also accepts Closure. Then it looks like one will also have to accept Negative Self-Intimation to avoid no-lose investigations. If one admits that a lack of justification can be inaccessible to an agent in some cases, then we can use those cases to build an example in which the agent gains justification for p if it's true, but is unable to notice if that justification is lacking. Telling the agent in advance about this arrangement will not prevent it; knowing that a (justificatory) condition is undetectable doesn't give the agent the ability to detect it. So if Negative Self-Intimation fails we can construct a no-lose investigation, taking advantage of the fact that at t_2 the agent can't use her lack of justification for p as a tip-off to its truth-value.

If this line of reasoning succeeds in establishing necessary conditions for avoiding no-lose investigations, we have a complete menu of epistemological options meeting our *desideratum*. What's interesting about this menu is that its options are either overtly skeptical, deny Closure, or have an internalist flavor (here

²⁴ One might object that the King doesn't know *exactly* what information he's going to get from the jester—he doesn't, for instance, know *which* maidens are going to turn out to have fallen victim to Giacomo's wiles. While it's curious to suggest that the particular details have this special significance in justifying the general proposition that the jester is a ladies' man, we can alter the example to remove those details entirely. Suppose that just before the jester begins, the King is called away on urgent matters of state. Later he asks one of his Earls who remained, "Did the jester tell copious tales of seduction?" Granting that the Earl is reliable (and perhaps that the King knows this as well), the Earl's single-word positive response has just as much justificatory force as the King's sitting through an evening with the jester would have.

Now imagine that the King knows in advance that he will be called away after giving Giacomo his orders, knows that he will get a report from the Earl later, and even knows his Earl well enough to know exactly what words will be spoken and how the report will sound. Then even before the jester begins the King knows *exactly* what information he will later receive about the jester's amorous tendencies. Yet actually hearing that information—actually *having* the conversation with his Earl—has positive epistemic value for the King. That's odd.

I'm counting both Wright's approach and Negative-Self Intimation).²⁵ If we want to avoid both overt skepticism and no-lose investigations, we must drop either Closure or epistemological externalism.

Acknowledgments I am grateful to Berit Brogaard, David Chalmers, Juan Comesaña, Kenny Easwaran, Branden Fitelson, Alan Hájek, Michael Lynch, Jonathan Schaffer, Declan Smithies, Elliott Sober, Jonathan Vogel, and an anonymous *Philosophical Studies* referee; also to audiences at the ANU Philosophy Society, the 2009 Australasian Association of Philosophy conference, the University of Wisconsin-Madison first-year philosophy seminar, the University of Wisconsin-Milwaukee, and the University of California, Berkeley.

Appendix

The proof below assumes we are in a situation meeting the two stipulations (concerning context-sensitivity and justification loss) made in Sect. 3 above; it also assumes Closure and Epistemic Reflection. It assumes as well that the agent has justification at t_1 for various propositions about her epistemic situation, including the fact that these stipulations and general epistemological conditions hold. This is certainly in the spirit of a view that maintains Negative Self-Intimation, but more importantly it focuses our attention on the epistemological aspects of interest. We are not concerned with how the agent got to the epistemic position she occupies at t_1 ; we are interested in the features of the investigation that commences at that point. If an epistemological theory allows a no-lose investigation because an agent fails to be informed about the justification conditions for various propositions, that is not a particularly interesting criticism of the theory. If, on the other hand, a theory allows for no-lose investigations even when the agent is epistemologically well-informed, that's a serious problem.

I'll present the proof and then explain some of its steps. The goal is to show that against the background just described, the combination of Negative Self-Intimation with our three no-lose investigation conditions entails a contradiction.

(1)	$J_1(J_2(p \rightarrow J_2p))$	(see below)
(2)	$J_1(\sim J_2p \rightarrow J_2 \sim J_2p)$	(see below)
(3)	$J_1(\sim J_2p \rightarrow J_2 \sim p)$	(see below)
(4)	$K_1(\sim J_2 \sim p)$	Condition 2
(5)	$J_1(\sim J_2 \sim p)$	$K \rightarrow J$
(6)	J_1J_2p	Closure from (3), (5)
(7)	J_1p	Epistemic Reflection
(8)	$\sim J_1p$	Condition 1
(9)	F	

²⁵ This is the first time I have invoked the dreaded internalist/externalist distinction in this article. While I mean to be applying some sort of "access" internalism/externalism distinction here, I won't say more about the distinction except that I trust externalists will be unhappy with both Wright's position and Negative Self-Intimation. (For Wright's description of his position as internalist see his (2004, pp. 209ff.).)

- Step (1) By Condition 3, the agent knows at t_1 that $p \rightarrow J_2p$. By $K \rightarrow J$, the agent is justified in this belief as well. By our no-justification-loss stipulation and the agent's epistemological awareness at t_1 , she has justification at t_1 for the proposition that this proposition will remain justified for her at t_2 .
- Step (2) Follows from the agent's t_1 justification to believe Negative Self-Intimation.
- Step (3) $p \rightarrow J_2p$ and $\sim J_2p$ entail $\sim p$. By Closure, $J_2(p \rightarrow J_2p)$ and $J_2 \sim J_2p$ entail $J_2 \sim p$. Thus if Closure holds, $\sim J_2p \rightarrow J_2 \sim J_2p$ and $J_2(p \rightarrow J_2p)$ entail $\sim J_2p \rightarrow J_2 \sim p$. The agent's t_1 epistemological awareness gives her justification for Closure, and the agent has justification at t_1 for both $\sim J_2p \rightarrow J_2 \sim J_2p$ and $J_2(p \rightarrow J_2p)$ (by Steps (2) and (1)). So by Closure, the agent has justification at t_1 for $\sim J_2p \rightarrow J_2 \sim p$.

References

- Armtenius, F. (2003). Some problems for conditionalization and reflection. *The Journal of Philosophy*, 100, 356–370.
- Cohen, S. (2002). Basic knowledge and the problem of easy knowledge. *Philosophy and Phenomenological Research*, 65, 309–29.
- Conee, E., & Feldman, R. (1998). The generality problem for reliabilism. *Philosophical Studies*, 89, 1–29.
- Feldman, R. (2006). Some epistemological puzzles about disagreement. In S. Hetherington (Ed.), *Epistemology futures*. Oxford: Oxford University Press.
- Fumerton, R. A. (1995). *Metaepistemology and skepticism*. Lanham MD: Rowman & Littlefield.
- Goldman, A. I. (1976). What is justified belief? In G. S. Pappas (Ed.), *Justification and knowledge* (pp. 1–23). Dordrecht: D. Reidel.
- Hájek, A. (2005). The cable guy paradox. *Analysis*, 65, 112–119.
- Klein, J. (2009). Francis Bacon. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Spring 2009).
- Kvanvig, J. L., & Menzel, C. (1990). The basic notion of justification. *Philosophical Studies*, 59, 235–261.
- Nozick, R. (1981). *Philosophical explanations*. Cambridge MA: Harvard University Press.
- Popper, K. R. (1961). *The logic of scientific discovery*. New York: Science Editions.
- Pryor, J. (2000). The skeptic and the dogmatist. *Nouûs*, 34, 517–549.
- Titelbaum, M. G. (2008). The relevance of self-locating beliefs. *Philosophical Review*, 117, 555–605.
- Van Cleve, J. (2003). Is knowledge easy—or impossible? Externalism as the only alternative to skepticism. In S. Luper (Ed.), *The skeptics: Contemporary essays* (pp. 45–60). Ashgate Publishing Limited.
- van Fraassen, B. C. (1995). Belief and the problem of Ulysses and the Sirens. *Philosophical Studies*, 77, 7–37.
- Vogel, J. (2000). Reliabilism leveled. *The Journal of Philosophy*, 97, 602–623.
- Wright, C. (2004). Warrant for nothing (and foundations for free)? In *Supplement to the Proceedings of the Aristotelian Society* (pp. 167–212).