

The Unfortunate Consequences of Progress in Philosophy

Bryan Frances

We tend to think that philosophical progress, to the extent that it exists, is a good thing. I agree. Even so, it has some surprising unfortunate consequences for the rationality of philosophical belief.

1. *Epistemic Upheaval and Progress*

In the bad old days, baseball pitchers were usually evaluated on their number of wins, win-loss percentage, earned run average, and number of strikeouts. Hitters were usually evaluated on their batting average, number of runs batted in, and home runs. No one who has kept up with developments in judging baseball performance thinks this is a valuable way to go.

The reason for the shift is that the field of statistical evaluation of baseball performance has gone through a period of *epistemic upheaval*. Thanks to the work of Bill James and others, we have *far* better statistical tools for evaluating performance in baseball as well as other sports. As a consequence, many old arguments regarding such performance are now viewed as almost primitive. What passed for a good argument forty years ago is often judged naïve today, by experts anyway.

The phenomenon of epistemic upheaval applies to philosophy as well. Much of the work from roughly fifty years ago concerning many philosophical topics is almost primitive compared to that of today. I will indicate what ‘epistemic upheaval’ comes to, in philosophy, through four real-life examples.

A reader even remotely familiar with the state of the art regarding vagueness will find little worth studying from before the 1950s, especially before the 1975 volume 30 issue of *Synthese* on the topic (not nothing: *little*). There was a subsequent explosion of good work that utterly transformed the area, particularly with Williamson’s 1994 *Vagueness*. The same is true of material composition, with van Inwagen’s 1990 *Material Beings* and certain other key works published around the same time. The philosophical study of consciousness was transformed in the 1990s and 2000s starting with Chalmers’ 1996 *The Conscious Mind*. More recently, such transformation happened for the epistemology of disagreement, with initial works by Frances (2005a, 2005b), Kelly (2005), Feldman (2006a, 2006b), and Christensen (2007). Over the last fifty or so years, epistemic upheaval has also occurred with topics in the philosophy of perception, the philosophy of biology, and other areas.

Let’s not be overly sentimental about our personal favorites: not *every* philosophical topic has experienced epistemic upheaval over the last fifty years—comparable to what we see for vagueness, material composition, consciousness, and disagreement. But a good number have.

I have used well-known examples in order to indicate what I mean by ‘epistemic upheaval’. It’s clearly tied to progress: regardless of how one construes ‘philosophical progress’, the four examples above definitely involved *overall* progress (at least epistemic if not alethic) by the philosophical community on philosophical issues (‘overall’ in order to allow for some backsliding in addition to forward movement). I won’t offer a definition of ‘epistemic upheaval’, since it turns out that none is needed for my arguments.

In order to get a hint of the surprising argument of this paper--which concludes with negative news regarding the epistemic consequences of philosophical progress--pretend what it would have been like to be a philosopher of language working on vagueness in the late 1960s but somehow aware that very soon there will be improvements that make many of the arguments and positions being thought of and discussed in the 60s almost primitive in a comparative sense. Perhaps an oracle informs you of these facts about the near future. Naturally, you are excited to learn these things. After all, the investigation into vagueness is about to make significant progress! This is a cause for celebration and happy anticipation.

But then you go back to your study and look at the vagueness papers you have already published or the new ones you have been working on lately. You can't help but ask yourself: what are the odds they contain good arguments and positions, given what the oracle just told you?

You should probably guess that the odds are pretty slim, provided (i) in those works you aren't arguing for logical truths that you can easily symbolize and then derive, (ii) you are arguing for theses more ambitious than 'That argument of so-and-so has such-and-such flaw', and (iii) you recognize that you are not a genius when it comes to philosophy. Sure, your essays on the topic might be publishable and good or even excellent *for the present time*, in the sense they are about as high quality as other essays being published now. The oracle hasn't told you anything that should make you think your work is worse than that of most other philosophers working on the topic now. Neither has she told you anything scandalous about the work being done by others at your time. But you know, from the oracle, that the odds are very high that there are serious flaws in your work on the topic—so serious that your work will be viewed, by near-future experts, as almost *primitive*.

Given that insight, what do you now think about the quality of your current arguments regarding vagueness, or some other topic that is going to go through epistemic upheaval?

2. Major Argumentative Flaws

Just because there is going to be great progress in a field or on a topic doesn't suggest that all or most of the current thinking in that field or topic is mistaken. Even if progress in mathematics is extreme, fifty years from now we are still going to think that $14 \times 14 = 196$ and, more interestingly, there are infinitely many pairs of primes that differ by at most 600. This holds in philosophy as well. Even if over the next decade there is going to be all sorts of revolutionary work done in epistemology, for instance, we can be confident that the odds are quite high that each of the following claims won't be revealed to be seriously problematic.

- Two people can believe the same thing but one person's belief is irrational while the other's is rational.
- Knowledge is objective in this sense: just because someone thinks she knows P doesn't always mean that she really does know P.
- One can have a true belief P without knowing that P.

Even so, it's informative to skim journal articles from fifty to a hundred years ago. Don't go for the classics alone. Instead, open a bunch of journal issues at random and skim the *ordinary* articles that satisfy two conditions: they treat topics you know very well and that have subsequently undergone epistemic upheaval (comparable to the cases indicated in the previous section). If your experience is like mine, it

won't be difficult to find in *the vast majority* of articles *major* argumentative flaws such as the following six, which are worded as if by an honest if brusque journal referee.

Problematic Premise. Look at the second premise, P2. The author thinks that it's unproblematic, as she asserts it with complete confidence, no defense, and not even much comment. [It might not even be stated in the article, if the author is really unaware of its role in her argument.] The author is apparently blind to the fact that there are *excellent* reasons to think P2 is false--reasons that she has no idea how to counter. In fact, we know today that P2 is highly contentious, as the arguments against it are serious. Now, I'm not saying premise P2 *is* false. All I'm saying is that she is relying on P2 as unproblematic and yet it certainly isn't anywhere close to having that status. Sure, she could just respond by saying that she's *assuming* the truth of P2. That's fine, but that move certainly reduces the interest in her argument, as her conclusion switches from C to 'P2 \supset C'.

Ambiguous Premise. There's a serious problem with premise P4. The author seems oblivious to the fact that it's crucially ambiguous in a philosophically relevant way. If you interpret it one way, then I agree with the author that it's unproblematic, or at least highly plausible. But then under that reading it offers virtually zero support for her conclusion. If you interpret it the other way, then it would certainly support her conclusion—if it were true. But of course we know that it's highly contentious when interpreted that way--for reasons that she, like us, simply cannot defeat at all. I'm not saying that P4 really is false. But it certainly stands in need of a defense! Or the author can just assume its truth, but then we have the familiar problem of the significant loss of interest in her argument.

Unsupportive. Perhaps all the premises are true, but the argument isn't even inductively strong, let alone deductively valid. The premises don't do shit to support the conclusion!

Ambiguous Conclusion. The conclusion is ambiguous in a philosophically relevant manner. Interpreted one way, the supporting argument is great, but the conclusion is philosophically boring—even for the author. Interpreted the other way, the conclusion is definitely worthwhile, by just about anyone's lights, but the argument doesn't support it (either validly or inductively), for reasons I already went over elsewhere.

Mess. The argument is a total mess: it's hardly clear what the conclusion or the premises even are. There are multiple ways of developing it but not in any way that remains at all faithful to the original and gives us premises that aren't highly contentious and that offer decent support for a worthwhile conclusion.

Tweak. My criticisms wouldn't mean much if the author could simply tweak her argument so that it was either valid or inductively strong, had premises that weren't highly contentious, and had a worthwhile conclusion. But I see no plausible way to do that--and more importantly, the author certainly gives no indication that she knows of any such way.

The list is not exhaustive! But it captures a good portion of the flaws we actually encounter in philosophical arguments. Call these six plus closely related ones *Standard flaws*.

In saying that the main arguments of those articles suffer from Standard flaws, I am not saying that the articles are bad, relative to the past or even the present. Argumentative quality doesn't fix the worth of a

work of philosophy. For instance, merely articulating an interesting thought experiment can be an enormous advance. And even a *definitely bad* argument can be highly valuable, if, for instance, it is both highly original and it can be repaired by someone else to become much stronger and retain a philosophically interesting conclusion. But often enough, when we publish a paper that contains arguments, we think or at least *hope* the arguments are valid or at least inductively strong, contain premises that aren't highly contentious, and have philosophically worthwhile conclusions. It's fine to have an argument for conclusion C that relies on an *assumption* A that is highly contentious, as when you are not really presenting an argument for conclusion C but for the qualified conclusion ' $A \supset C$ '. But in this case, you don't want your (other) premises to be highly contentious.

With regard to the perusal of old, pre-upheaval articles, I'm saying that they had those flaws *when they were published*. For instance, it's not as though the arguments had unambiguous premises back then but have ambiguous ones now. Decades ago almost no one was in a position to figure out that those arguments had Standard flaws, but the flaws were there anyway, waiting to be revealed.

I restrict myself to what might be called *harmful* ambiguity. One can offer an argument using a predicate with multiple meanings but the argument goes through under any consistent disambiguation. For instance, 'epistemically rational' may have multiple meanings but *under each one* the author's argument has premises that are plausible, the argument is valid or inductively strong, and the conclusion is philosophically worthwhile according to a sizeable block of philosophers. In that scenario the ambiguity need not be harmful.

3. *The Unfortunate Consequences of Epistemic Upheaval*

After learning that one or more of the Standard flaws is true of the large majority of the primary arguments in articles published before a period of epistemic upheaval, I am pretty much forced, by my own rationality and awareness, to reason inductively that one or more of those flaws is probably true of the primary arguments in the papers I'm writing now, provided my current papers concern topics that are going to undergo epistemic upheaval. After all, I'm just an average philosopher, and I'm not going around arguing for logical truths or overly modest theses of the form 'Smith's argument has the following weakness'.

Perhaps there is no current upheaval with respect to topic T that I am writing on today, and there isn't any strong sign of one occurring soon. But given that this is philosophy, isn't it highly probable that there will, in the next fifty years or so, be a period of upheaval for T (again, comparable to the cases mentioned in section 1)? If so, then it's highly probable that that farther-future upheaval will show that it's highly probable that my work on T is flawed in the sense of having one or more of the Standard flaws. Whether the upheaval is going to happen in the next few years or in the next few decades or even century matters not at all: either way, *one or more of the Standard flaws is very probably true of my work now*. We can complain about how progress in philosophy is not nearly as impressive as it is for physics, chemistry, and biology, but what we do know is that we philosophers are simply outstanding at finding serious flaws in arguments and claims. Hence, it's highly probable that in the future the philosophical community could, given a bit of time and effort, show that some of the Standard flaws are true of my work.

If I have appreciated this probability fact just articulated, then what should my attitude be towards my current and past work—or even future work? It may be well worth publishing, given the current state of the art, but should I believe in the soundness or other cogency of its arguments?

Let's make this imprecise argument precise.

- A. I know that for most any pre-epistemic upheaval article from the past, it is highly probable that one or more of the Standard flaws is true of its main arguments.

Epistemic upheaval *can* be entirely positive instead of negative. For a given topic, the upheaval might not find *any* mistake in prior work. It might only add to prior work in a radical way. In that case, (A) could be false. But the kind of upheaval we see in philosophy is partly negative, finding faults in prior work that previous philosophers were largely unaware of.

It's easy to see that many of the arguments of those prior works have false conclusions, since their conclusions are mutually contradictory. But (A) is saying something else: the Standard flaws are there a high percentage of the time. A good way of arguing for this would be to take about 100 pre-epistemic upheaval articles chosen at random and proving that the vast majority of them suffer from Standard flaws. Needless to say, no one is going to go through the trouble of doing that. All I can do here is report that according to my experience looking through pre-epistemic upheaval articles, this was true of them. As mentioned before, this is *not* to say that I discovered or think that those articles were rubbish or that none of them are worth reading today.

As one might expect, certain articles are far less likely to have arguments with Standard flaws. For instance, if you're arguing for a logical truth, and you bother to do the appropriate symbolization and derivation, with due diligence to what the relevant expert community thinks about the reliability of the operative inference rules, your argument is probably pretty solid. More interestingly, if your conclusion is of the form 'That argument suffers from this flaw', then there is a good chance your argument has a true conclusion. But most of us aim higher, trying to establish more potent theses. From now on, only such "aim higher" arguments are intended. There is still a great variety and number of such arguments, with conclusions such as 'Token Identity theory is true', 'Know-how reduces to know-that', 'Propositional content doesn't weakly supervene on internal physical properties and relations', 'Conciliationism about disagreement is false', 'There can be artworks of great aesthetic value that are immoral', 'Presentism is consistent with General Relativity', etc.

Roughly put, premise (B) says that you know that you're not *significantly* better than philosophers in the past when it comes to avoiding Standard flaws in arguments.

- B. I know that if for most any pre-epistemic upheaval article *from the past*, it is highly probable that one or more of the Standard flaws is true of its main arguments, then for most any pre-epistemic upheaval essay *of mine today*, it is highly probable that one or more of the Standard flaws is true of its main arguments.

If you knew you were *vastly* better than most anyone in philosophy, then the conditional in (B) might be false applied to you. Congratulations; for the remainder I will ignore you. For the rest of us, I need to defend (B) in detail.

In my experience, many philosophers respond to claims similar to (B) with a comment that starts with 'Yes, but'.

1. "Yes, but I do philosophy quite differently from how it was done in the past, and the differences make it much less probable that my main arguments suffer the flaws that infect the main arguments of past articles; so, the conditional in (B) is false applied to me".

For example, most of my own published arguments over the past few years are obviously valid and have explicitly formulated premises and conclusions, even when I don't bother to symbolize them in the published versions. My brainstorming documents, scribbled or electronic, are filled with symbolizations and derivations even though I don't publish anything in the more formal areas of philosophy. The vast majority of arguments from fifty or more years ago don't have that kind of rigor, and the vast majority of philosophers didn't use that level of rigor in their personal notes either.

That's nice; I get a loving pat on the head. But it should be obvious that there is a real threat of wishful thinking here. Although I may have *an* advantage over the vast majority of philosophers from the past, that alone hardly suggests that I am significantly more likely to avoid Standard flaws.

2. "Yes, but the topics I am currently writing about are so different from those of the past that the induction in (B) doesn't apply to my work".

If so, then perhaps the conditional in (B) is false applied to you. That's *possible*, provided your current work is really *very* different in kind from past philosophizing. For most of us, however, that simply isn't true.

I formulated (1) and (2) in the first-person because that is how I usually hear them. The argument's conclusion (below) seems a bit insulting, embarrassing, or at least worrying, and people tend to react by defending *their own* case. We focus on ourselves, and ignore the universal quantifications over philosophers in general. One can of course be forgiven for wondering how well the argument applies to oneself. But it's unwise to do this when assessing this argument, since the argument's scope is thousands of times wider than just you or me. (1) and (2) are no good unless they apply to a great many philosophers.

Therefore, in order to evaluate the argument more objectively, ignore how it applies to you (and your favorite philosopher friends) and focus on how it applies to philosophers today in general. The next objection to (B) is similar to (1):

3. "The argumentative standards are much higher now compared to decades ago. So even though there were many Standard flaws back then, things are better now. In effect, most (not all) of us are the epistemic superiors of most (not all) philosophers from the past. Hence, the conditional in (B) is false when applied to most of us today".

I think the first three sentences of (3) are true for many areas of philosophy, and perhaps it means that today's arguments have fewer Standard flaws compared to the old days. Even so, that hardly means the final sentence of (3) is true, for two reasons.

First, there is reason to think that although the frequency of Standard flaws has decreased, it has not done so by much. In particular, the increased rigor will help significantly diminish the frequency of the flaws Unsupportive and Mess but not so much the other four flaws. Second, perhaps there is another way that

philosophy is done differently now that actually *increased* the frequency of Standard Flaws (or a subset of them), thereby offsetting the improvement brought on by the more common attention to making arguments valid and explicit. So I don't think there is much ground for thinking the conditional in (B) is false for most of us today. Perhaps the probability of Standard flaws is lower today, but it's still great enough to be "high", as (B) says.

- C. I know that for most any article I am working on today, it's highly probable that it is a pre-epistemic upheaval article (whether the upheaval is imminent or decades in the future).

For most topics most of us are writing about today, it is highly probable that eventually it will experience negative epistemic upheaval. In fact, it is difficult to think of many philosophical topics that haven't experienced such upheaval over the last century or so. And it would be wildly optimistic to think that typical topics have now, finally, reached a point at which negative upheaval won't happen anymore.

My argument should not be thought to fail if a huge meteor destroys civilization tomorrow. We know that *if* philosophers were to live and work for the few decades in something even approximating the fashion we actually have for the last century, *then* for almost any topic that we are actually writing about now, it *would* go through a period of negative epistemic upheaval in the next few decades. That's the understanding behind premise (C). (A)-(C) support intermediate conclusion (D).

- D. Hence, from (A)-(C) I know that *for most any article I'm working on today, it's highly probable that its main arguments suffer from Standard flaws.*

The italicized claim in (D) will look bad, embarrassing, or at least worrying to many readers. But we should keep mind what (D) is *not* saying.

- Just because an argument has one or more Standard flaws doesn't mean its conclusion is false or that other arguments, philosophical or not, offer excellent evidence for the conclusion of the flawed argument.
- Neither is there any reason here to think that the flawed arguments are worthless, relative to the current or even future state of the art, since flawed arguments can have significant virtues, as pointed out earlier.
- None of these considerations shows that philosophers are irrational in believing various philosophical positions. After all, they might not have put any thought into these unsettling facts about negative epistemic upheaval.
- Neither do they show that such philosophers are irrational if they know that there will be negative epistemic upheaval on their topics but stick with believing their positions anyway: again, they might not have thought things through like we are doing here.
- Neither is there anything epistemically wrong with continuing to articulate and defend your views even after *deeply* appreciating the consequences of negative epistemic upheaval. There is nothing wrong with continuing to publish and professionally endorse those views (e.g., "I will defend thesis T now"), as that can easily be done sans belief (e.g., I have been doing it for many years, and I know that I'm not exceptional in that respect).

However, (D) is unsettling: a quick read of the six Standard flaws shows that if one or more Standard flaws are true of an argument, then it's improbable that it meets this attractive yet modest three-part condition: *under a single interpretation (1) the argument is either deductively valid or inductively strong, (2) none of its premises is highly contentious, and (3) it has a philosophically worthwhile conclusion.* (Recall the qualification from the previous section about employing highly contentious *assumptions*.) Call that three-part condition *Worth*. Most of us want *Worth* to apply to our arguments for "aim higher" theses. If you are forced to conclude that your arguments fail to satisfy *Worth*, then your arguments would most likely be failures by your own lights.

- E. I know that *if* for most any article I'm working on today, it's highly probable that its main arguments suffer from Standard flaws, *then* for most any article I'm working on today, it is improbable that its main arguments satisfy *Worth*.
- F. Conclusion 1: from (D) and (E), I know that for most any article I'm working on today, it is improbable that its main arguments satisfy *Worth*.

Some philosophers will yawn at (F), claiming that they don't really believe their conclusions anyway, or at least were not aiming for *Worth*. Well, okay. But like many other philosophers I would like to articulate arguments that satisfy *Worth*. I work pretty hard on these arguments, and if they can't even meet the *quite modest* standard of (1)-(3), that is pretty bad. In fact, if I think my argument has Standard flaws, then I will probably conclude that I have little reason to accept the conclusion of that argument, at least on the basis of my supporting argument. That leads to the last premise.

- G. If I know that for most any article I'm working on today, it is improbable that its main arguments satisfy *Worth*, then I should suspend judgment on their conclusions unless I know of some other, vastly superior, reason to believe the conclusion.
- H. Conclusion 2: by (F) and (G) I should suspend judgment on the conclusions of the main arguments unless I know of some other, vastly superior, reasons to believe their conclusions.

The 'unless' clause is intended to cover philosophical arguments for conclusions that are commonsensical, such as 'Much of what Hitler did was morally wrong', 'There are trees in North America', ' $2 + 2 = 4$ ', and 'Some tomatoes are at least partially red'. One might mount an argument for those theses (and thus against various radical philosophical theories) even though one thinks that there are other, non-philosophical, reasons for them. I mean to set those aside, merely because they raise special issues that I don't want to address here.

4. *What Are the Consequences of the Unfortunate Consequences?*

There is an obvious self-application issue: why think that my (A)-(H) argument is sound, given that it's a pretty typical philosophical argument, offered by a pretty typical philosopher, and regards a topic that will probably go through negative epistemic upheaval? Set that issue aside and foolishly pretend for a moment that the primary conclusions (F) and (H) are true. What would that mean for the average professional philosopher?

In his book *On Writing: A Memoir of the Craft*, Stephen King repeats the traditional wisdom that when it comes to novel writing, "Kill your darlings, kill your darlings, even when it breaks your egocentric little

scribbler's heart, kill your darlings." I wrote a novel, so I know firsthand how hard it is to do this. Part of what makes following the "kill your darlings" advice difficult is that it's entirely internal since you have to *go against yourself*, and not someone or something external. Many things in life are difficult, like making partner in a law firm, getting a contract as a professional athlete, dealing with a horrible boss, and learning a new language as an adult. Those are largely external challenges, even the last one. But killing one of your own favorite characters or themes in your novel is like killing part of yourself.

Suspending judgment on your own philosophical arguments is somewhat similar to killing part of your story. You work for years on a series of arguments for a certain philosophical view. Hundreds of hours are put into it, and the result seems to be a key part of you: one of your "children". Deciding that the odds are very high that the arguments don't satisfy Worth is quite a blow, a sort of self-inflicted one like when the novelist realizes that despite her tremendous efforts, her novel would be much better without one of her favorite themes or characters. Some novelists can't do the deed, so the theme or character remains in the novel. Some philosophers can't do the deed, so they don't suspend judgment. The good novelists do the deed.

References

- Chalmers, David. 1996. *The Conscious Mind*. New York: Oxford University Press.
- Christensen, David. 2007. "Epistemology of Disagreement: The Good News," *Philosophical Review*, 116: 187–218.
- Feldman, Richard. 2006a. "Epistemological Puzzles about Disagreement," in Steve Hetherington (ed.), *Epistemic Futures*. New York: Oxford University Press, 216–236.
- Feldman, Richard. 2006b. "Reasonable Religious Disagreements," in L. Antony (ed.), *Philosophers without Gods: Meditations on Atheism and the Secular Life*. New York: Oxford University Press.
- Frances, Bryan, 2005a, "When a Skeptical Hypothesis is Live," *Noûs*, 39: 559–95.
- Frances, Bryan. 2005b. *Scepticism Comes Alive*. New York: Oxford University Press.
- Kelly, Thomas. 2005. "The Epistemic Significance of Disagreement," in T. Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology*, vol. 1. Oxford: Oxford University Press.
- van Inwagen, Peter. 1990. *Material Beings*. Ithaca: Cornell University Press.
- Williamson, Timothy. 1994. *Vagueness*. London: Routledge.