How to Argue for Pragmatic Encroachment[*]

BLAKE ROEBER
University of Notre Dame

Abstract: Purists think that changes in our practical interests can't affect our knowledge unless those changes are truth-relevant with respect to the propositions in question. Impurists disagree. They think changes in our practical interests *can* affect our knowledge even if those changes *aren't* truth-relevant with respect to the propositions in question. I argue that impurists are right, but for the wrong reasons, since impurists haven't appreciated the best argument for their own view. As I show, there is an argument for impurism sitting in plain sight that is considerably more plausible than any extant argument for impurism.

How, if at all, do our practical interests affect our knowledge? According to the thesis that I will call 'purism,' changes in our practical interests can't affect what we know unless those changes are *truth-relevant* with respect to the propositions in question. According to the negation of this thesis, which I will call 'impurism,' changes in our practical interests *can* affect what we know even if those changes *aren't* truth-relevant with respect to the propositions in question. If impurism is true, then changes in our practical interests might affect our knowledge without affecting our evidence for the relevant proposition, the reliability of the cognitive faculties responsible for our belief in that proposition, the safety of our belief in that proposition, and so on, for any other truth-relevant property that we might care about.[1]

The literature contains two kinds of arguments for impurism: what I will call 'principle-based arguments' and 'intuition-based arguments' ('PBAs' and 'IBAs' for short).[2] The former attempt to motivate impurism by motivating some principle like KA, below, and then deducing impurism from this principle.

(KA)   *S* knows that *p* only if she can rationally act as if *p*.

Prominent PBAs include Fantl and McGrath's (2002, 2009), Ganson's (2008), Weatherson's (2012), Schroeder's (2012), Ross and Schroeder's (2014), and on one reading Hawthorne and Stanley's (2008). In contrast to PBAs, IBAs offer impurism as the best explanation of our intuitive reactions to pairs of cases that differ with respect to the subject's practical interests, but do not differ with respect to the strength of the

---

[1] By 'truth-irrelevant factors,' I mean exactly what DeRose means (2009: 25): factors that don't affect the probability that the belief in question is true, either from the point of view of the person who holds that belief or from any more objective point of view.
[2] Cf. Roeber 2018a, p. 1.

subject's epistemic position.[3] Stanley's (2005) IBA is the central example. His argument involves a person named 'Hannah' who seems to say something true by uttering the sentence 'I know that the bank will be open' in the case that Stanley calls 'Low Stakes,' and then seems to say something true by uttering the sentence 'I don't know that the bank will be open' in the case he calls 'High Stakes.'[4] Since Hannah's practical interests change between Low Stakes and High Stakes while the strength of her epistemic position with respect to the proposition that the bank will be open does not, Stanley offers impurism as part of the best explanation of our intuition that Hannah says something true in both Low Stakes *and* High Stakes.

Purists have raised important objections to both kinds of arguments for impurism. Stanley's IBA faces two challenges. First, recent results in experimental philosophy suggest that it's *not* intuitive (in the relevant sense) that Hannah says something true in High Stakes.[5] But if this isn't intuitive, then there's nothing here for impurism to explain. Second, even if it *is* intuitive that Hannah says something true in High Stakes, this intuition is arguably best explained in purist-friendly terms. The contextualist explanations forwarded by DeRose (1992, 2009), Cohen (1999, 2005), and others, the warranted assertability maneuvers advanced by Rysiew (2001), Brown (2006), Reed (2013), Lutz (2014), Locke (2017), and others, and the explanations in terms of missing belief developed by Bach (2005) and Nagel (2008, 2010) all provide purist-friendly explanations of our intuition that Hannah says something true in High Stakes. And as I point out elsewhere (Roeber 2018a: 18-19), by changing just what Hannah says in High Stakes, we can elicit the intuition that she *does* know that the bank will be open in High Stakes, which suggests that our intuitions about High Stakes track what Hannah tells us about her knowledge rather than the relevant details of her practical interests. Since analogous challenges would face any other IBA, the case against purism arguably rests on the best PBAs in the literature. But PBAs depend on principles like KA, and these principles seem susceptible to the counterexamples forwarded by Brown (2008), Reed (2010), Cohen (2012), Anderson (2015), myself (2018a), and many others. With all of the relevant considerations in view, there seems no serious threat to purism—or so many purists think.

---

[3] As I'm using the word 'intuitive,' our intuitive reactions to cases report how these cases *initially strike us*. They needn't report our settled opinions about these cases, and may even conflict with our settled opinions. The skeptical paradox helps illustrate the distinction. In the relevant sense of 'intuitive,' I find each of the following simultaneously intuitive: (a) that I know that I have hands, (b) that I don't know that I'm not a disembodied brain in a vat, and (c) that *if* I know that I have hands, then I *do* know that I'm not a disembodied brain in a vat. My settled opinion is that, since (a) and (c) are true, (b) must be false. But while my settled opinion is that (b) is false, I still find (b) intuitive, in the relevant sense. Throughout this paper, whenever I use the words 'intuition,' 'intuitive,' *etc.*, I will be using them in this sense.

[4] These cases come originally from DeRose 1992.

[5] See, for example, Buckwalter 2010, May, Sinnott-Armstrong, Hull, and Zimmerman 2010, Feltz and Zarpentine 2010, Schaffer and Knobe 2012, Phelan 2014, Buckwalter and Schaffer 2015, Turri 2017, and especially Rose *et al* 2017.

This, in fact, is exactly what I used to think. But as I will now argue, impurists haven't given the best argument for their own view, and purists haven't produced any adequate response to this argument. In certain respects, this argument has been sitting in plain sight. It simply develops an idea suggested by Fantl and McGrath's passing comment that "we care more about getting the best results than getting the best expected results" (2009: 81), which they use to block an objection to their argument for the principle they call "Safe Reasons," which they employ in their argument for the principle they call "KJ," which is the main premise in their argument against purism. But as I hope will become clear, the idea suggested by Fantl and McGrath's passing comment supports a much more direct and conservative argument for impurism than any extant argument for impurism, and this new argument for impurism isn't susceptible to the apparent counterexamples that plague Safe Reasons and the rest of the principles in the pragmatic encroachment literature. In fact, as I hope to show, this argument gives impurists a way to motivate their view even if the best PBAs and IBAs in the literature all fail.

1. Preliminaries

Purism is the thesis that truth-irrelevant changes in our practical interests can't affect our knowledge. Thus, purism is true iff no pair of possible cases satisfies the following description: *S* knows that *p* in Case A but not Case B, even though Case B is identical to Case A except for some truth-irrelevant difference in *S*'s practical interests. Call pairs of possible cases that *do* satisfy this description 'encroachment cases.'[6] Since impurism is the negation of purism, impurism is true iff there is at least one pair of encroachment cases. Thus, impurism is logically equivalent to an existential claim. Now, of course, *skepticism* trivially entails that encroachment cases aren't possible, so it trivially entails that purism is true. And skepticism also trivially entails that principles like KA are true, since it entails that there is no possible scenario where one of these principles has a true antecedent. Since the interesting question is whether purism is true *even if skepticism is false*, and since the literature contains virtually no debate between impurists and *skeptical* purists, I will assume throughout that skepticism is false. Indeed, whenever I say 'purists,' I will mean *non-skeptical* purists. With this in mind, consider the following cases (adapted from Reed 2010), and assume that they are identical except for the stipulated difference in your practical interests.

> **Case A**: You are participating in a psychological study where the researcher asks you questions about ornithology—a subject with which you are very well acquainted. For every correct answer you give, the researcher will reward you

---

[6] Encroachment cases are a species of what I call 'encroachment scenarios' in Roeber 2018b. Specifically, encroachment cases are just encroachment scenarios where *practical interests* are the truth-irrelevant factors in question.

with a jellybean; for every incorrect answer, and every question left unanswered, you will get nothing. The first question is whether hawks are raptors. You know that the answer is 'yes.'

Case B: You are participating in a study exactly like the one in Case A, except in this study you know that you will be punished by an extremely painful electric shock each time you give an incorrect answer. Again, the first question is whether hawks are raptors. Though you remain convinced that the answerer is 'yes,' and while your epistemic position with respect to the proposition that hawks are raptors is just as strong in this case as it was in Case A, you shouldn't answer the question, since you know that the shock would be unbearable, and you know that, even if you are right, you will only get a jellybean.[7]

These cases don't support impurism by evoking anything analogous to the intuition that Hannah says something true by uttering the sentence 'I do not know that that bank will be open' in High Stakes, the intuition that Hannah does not know that the bank will be open in High Stakes, or any other knowledge denying intuition—as would be required if impurists wanted to use Case A and Case B in an IBA. Indeed, as Reed points out (p. 229), it's actually intuitive that you *do* know that hawks are raptors in Case B, and for this reason it's tempting to think that Case B is a counterexample to principles like KA. By adding further details to Case B, however, impurists can make it increasingly plausible that our intuitions about Case B are actually mistaken, and that, contrary to the way Case B initially strikes us, you *do not* know that hawks are raptors in Case B. Since they can add these details while honoring our assumption that Case B is identical to Case A except for the stipulated differences in your practical interests, they can make it increasingly plausible that Case A and Case B give us a pair of encroachment cases. And since impurism is true if there is a single pair of encroachment cases, they can make it increasingly plausible that impurism is true.

In Case A and Case B, you understand the setup of the experiment. In both cases, you know what your options are, you know what consequences might follow from these options, and you know how to rank these consequences. Specifically, in both cases, you know that you have exactly three options: answering 'yes,' answering 'no,' and refraining from answering. In Case A, you know that, if hawks are raptors, then you will get a jellybean if you answer 'yes,' and you will get nothing if you answer 'no' or refrain from answering. In Case B, you know that, if hawks are raptors, then you will get a jellybean if you answer 'yes,' you will get a severe electric shock if you answer 'no,' and you will get nothing if you refrain from answering. So, in *both* cases, you know that, if hawks are raptors, then you will get a jellybean if you answer 'yes,' and you will

---

[7] Throughout, when I talk about answering the question whether hawks are raptors, I won't mean forming or holding any *belief* about the answer to this question. Instead, I will mean performing some *action* (saying 'yes,' pressing a button marked 'yes,' or something like that).

get either nothing or an electric shock if you do not answer 'yes.' Finally, in both cases, you know that it would be better to get a jellybean than to get either nothing or an electric shock, and you know that these are the only relevant consequences of your options. So, in both Case A and Case B, you know that, if hawks are raptors, then answering 'yes' will have the best consequences of your options.

This is all part of my intended reading of Case A and Case B. But we should be careful here. As Reed points out (ibid: 230), we might say that an option is best in order to convey that it has the highest *expected utility* of someone's options. This is important because, from the fact that hawks are raptors, it doesn't follow that answering 'yes' has the highest expected utility of anybody's options, including yours. The expected utilities of your options depend (among other things) on the relevant probability that hawks are raptors, not the *fact* that hawks are raptors.[8] If this probability is too low (if, for example, you are insufficiently confident that hawks are raptors, or you have insufficient evidence that hawks are raptors), then answering 'yes' will not have the highest expected utility of your options even though hawks *are* raptors. Since you can't know what isn't true, it can't be part of the correct understanding of Case A and Case B that you know that, if hawks are raptors, then answering 'yes' has the highest expected utility of your options.[9]

What *is* the correct understanding of Case A and Case B, then? Here we must distinguish between the expected utility of an action and its actual utility, where the actual utility of φ-ing will be the value of the consequences that will actually result from φ-ing. Suppose you toss a coin and bet me $1 that it landed tails. If I accept the bet and the coin landed tails, I lose $1 and neither gain nor lose anything else. If I accept the bet and the coin landed heads, I gain $1 and neither gain nor lose anything else. Finally, if I decline the bet, I neither gain nor lose anything at all. Now suppose the coin landed heads, but neither of us knows this yet. In this case, the expected utility of accepting the bet equals the probability of heads times $1, plus the probability of tails times –$1, which equals $0, assuming that heads and tails are equally probable. In contrast, because the coin *did* land heads, the actual utility of accepting the bet is exactly $1. In Case B, since hawks are raptors, the actual utility of answering 'yes' equals the value of receiving the jellybean, the actual utility of answering 'no' equals the value of receiving the electric shock, and the actual utility of refraining from answering equals the value of maintaining *status quo*. The conditional according to which, if hawks are raptors, then answering 'yes' will have the best consequences of your options is meant to have the

---

[8] The relevant probability will presumably be either your subjective probability (credence) that hawks are raptors or some epistemic probability that they are. It will not be the objective probability that hawks are raptors.

[9] I'm not here denying the material conditional *if you know that hawks are raptors, then answering 'yes' has the highest expected utility of your options* (cf. Weatherson 2012). Instead, at this point, I'm merely denying the strict conditional *if hawks are raptors, then answering 'yes' has the highest expected utility of your options* and pointing out that, because this conditional is *false*, you can't know that it's true.

same truth-conditions as the following material conditional: if hawks are raptors, then answering 'yes' will have the highest *actual* utility of your options. It's part of my intended reading of Case A and Case B that you know that *this* conditional is true.

So far so good. Now let's imagine impurists adding a detail to both Case A and Case B. Let's imagine that, in the versions of these cases that impurists want us to consider, the following claim is true:

> (C)   After reading the first question and considering the setup of the experiment, you reason as follows: "If hawks are raptors, then answering 'yes' will have the best consequences of my options. Hawks *are* raptors. So, answering 'yes' will have the best consequences of my options." On the basis of this reasoning, you form the belief that answering 'yes' will have the best consequences of your options.

Impurists can ask us to consider whatever cases they want us to consider. So, however we are inclined to interpret Reed's original jellybean case, impurists can stipulate that (C) is true in Case A and Case B. With this in mind, consider (1), below, and read it as a material conditional.

> (1)   If you know that hawks are raptors in Case B, then you know that answering 'yes' will have the best consequences of your options in Case B.

Should purists accept (1)? Some purists will accept it because it follows from some closure principle that they accept, but the deeper reason every purist should accept (1) is that, even if every extant closure principle is has counterexamples (as I'm inclined to think), purists will have no principled reason for resisting the bald stipulation that (1) is true. Why? Consider the following principle.

> **Deduction Guarantees Ignorance (DGI)**: For every value of '*S*,' every value of '*p*,' and every value of '*q*,' necessarily, if *S* deduces *q* from her belief that *p* and her belief that ($p \supset q$), then *S* does not know that *q*.

DGI seems absurd on its face, but in case it isn't already obvious why every purist should reject DGI, note that it entails KGI, below.

> **Knowledge Guarantees Ignorance (KGI)**: For every value of '*S*,' every value of '*p*,' and every value of '*q*,' necessarily, if *S* knows both that *p* and that ($p \supset q$), and *S* competently deduces *q* from her belief that *p* and her belief that ($p \supset q$), then *S* does not know that *q*.

Some purists might think that, if you deduce *q* from your belief that *p* and your belief that ($p \supset q$), then your epistemic position with respect to *q* will be weaker than both

your epistemic position with respect to p and your epistemic position with respect to $(p \supset q)$. Even so, no purist will think that, if you competently deduce q from your belief that p and your belief that $(p \supset q)$, then your epistemic position with respect to q must be *so much weaker* than your epistemic position with respect to both p and $(p \supset q)$ that your belief that q must fall short of knowledge *even if you know both that p and that $(p \supset q)$*. Competent deduction from knowledge may not guarantee more knowledge, but it certainly doesn't guarantee *new ignorance*. So purists should reject both KGI and DGI. They should conclude that it's possible to know each of p, $(p \supset q)$, and q after competently deducing q from p and $(p \supset q)$).

Once purists accept this conclusion, however, they will have no principled means of resisting the bald stipulation that (1) is true. Since the truth of the consequent of a material conditional suffices for the truth of the conditional itself, q entails $(p \supset q)$, for any values of 'p' and 'q.' Thus, where 'Kr' names the proposition that you know that hawks are raptors in Case B, and 'Kb' names the proposition that you know that answering 'yes' will have the best consequences of your options in Case B, Kb entails $(Kr \supset Kb)$. And since $(Kr \supset Kb)$ is identical to (1), Kb entails (1). So take any strength of epistemic position that a given purist thinks is necessary for knowledge. Then there seems no reason why impurists couldn't fill in the details of Case A and Case B so that your epistemic position with respect to the proposition that answering 'yes' will have the best consequences of your options is at least that strong. But this means that, for any given purist, there should be some way to fill in the details of Case A and Case B so that, *by the lights of that purist*, Kb is true. Thus, since Kb entails (1), there seems no reason why impurists couldn't fill in the details of Case A and Case B so that, by the lights of any given purist, (1) is true. This is why, if impurists just *stipulate* that (1) is true, purists will have no principled reason for rejecting the stipulation.[10]

2. Expected Utility *versus* Known Actual Utility

Let's take it for granted, then, that (1) is true in the versions of Case A and Case B that impurists are asking us to consider. Given this assumption, what follows? Case A stipulates that you know that hawks are raptors. By hypothesis, Case B differs from Case A only insofar as the stipulated difference in your practical interests requires that they differ. Thus, either Case A and Case B give us a pair of encroachment cases, or you know that hawks are raptors in Case B. If you know that hawks are raptors in Case B, however, then (1) entails that you know that answering 'yes' will have the best consequences of your options in Case B. So, if purism is true, you know that answering

---

[10] Again, I'm assuming that skepticism is false. (Skepticism entails that DGI and KGI are both *true*, since it entails that the consequent of DGI is true in every possible world while the antecedent of KGI is false in every possible world.)

'yes' will have the best consequences of your options in Case B. This follows from our assumption that (1) is true. With this in mind, consider (2), below, and read it as a material conditional like (1).

> (2)   If you know that answering 'yes' will have the best consequences of your options in Case B, then you may answer 'yes' in Case B.[11]

I find (2) plausible, at least on the intended reading of Case A and Case B. And as we'll see below, (1) and (2) give impurists a straightforward argument against purism. Before impurists can rely on (2) in an argument against purism, however, they must head off an objection suggested by Brown (2012), Reed (2010, 2013), and other purists. According to this objection, impurism is false unless fallibilism is true, but fallibilist knowledge doesn't require certainty. Thus, just as you might know that hawks are raptors even though you can't be certain that hawks are raptors, you might know that answering 'yes' will have the best consequences of your options even though you can't be certain that answering 'yes' will have the best consequences of your options. Because the potential costs of a mistake in Case B are so high, however, answering 'yes' won't maximize expected utility unless you *can* be certain that answering 'yes' will have the best consequences of your options. Since rationality requires that you maximize expected utility, you are rationally required to refrain from answering, even though you know that answering 'yes' will have the best consequences of your options. Thus, you shouldn't answer the question, even though you know that answering 'yes' will have the best consequences of your options. Given the truth of fallibilism, it's therefore clear how you might know that answering 'yes' will have the best consequences of your options in Case B, even though you shouldn't answer 'yes' in Case B. Since, intuitively, you know that hawks are raptors in Case B, and since, together with (1), your knowing that hawks are raptors in Case B entails your knowing that answering 'yes' will have the best consequences of your options in Case B, we have excellent reason to reject (2).

This is essentially Reed's take on Case B (2010: 229). In reply, impurists should emphasize that it's still unclear why you shouldn't answer 'yes' in Case B, if you know that answering 'yes' will have the best consequences of your options. As I noted above, when I say that φ-ing will have the best consequences of a person's options, I mean that φ-ing will have the highest *actual utility* of her options. Thus, (2) tells us that you may answer 'yes' in Case B if you know that answering 'yes' will have the highest actual utility of your options. So, to resist (2), purists must explain why you shouldn't answer 'yes' in Case B even though you know that answering 'yes' will have the highest actual utility of your options. And to explain this, purists cannot merely say that answering 'yes' doesn't have the highest *expected* utility of your options. First, answering 'yes' *does* have the highest expected utility of your options if the probability that hawks are

---

[11] Here and throughout, I am using 'you may' as shorthand for 'it is false that you should not.'

raptors conditional on your relevant evidence is 1, so purists must argue that, even though you know that hawks are raptors, the probability that hawks are raptors conditional on your relevant evidence is lower than 1. This won't be easy, but I won't press the point, since I think purists can successfully defend a version of fallibilism on which, even though you know that hawks are raptors, the probability that hawks are raptors conditional on your relevant evidence is lower than 1 (see Brown 2012: 57-9, for example). Indeed, from here forward, I will assume that exactly this sort of fallibilism is correct, and I will assume that the probability of a proposition conditional on your relevant evidence might be lower than 1 even *if* you know that the proposition is true. A point I *will* press, however, is that, even if this version of fallibilism is correct and answering 'yes' *does not* have the highest expected utility of your options, it's still hard to see why you should pick the option with the highest expected utility, if you know that answering 'yes' will have higher *actual* utility.

Why should a person care about, or be mindful of, or act in accordance with the expected utilities of her options? Why should expected utilities guide a person's actions, or constrain how she should act? According to one standard answer, the expected utilities of a person's options are her best guide to maximizing actual utility. This answer is clearly often right. It's easy to see why the expected utilities of a person's options would be her best guide to maximizing actual utility in a case where she doesn't know which of her options will have the highest actual utility. It's also arguable that the expected utilities of a person's options would be her best guide to maximizing actual utility in a case where she *does* know which of her options will have the highest actual utility and this option is also her option with the highest *expected* utility. But why would the expected utilities of a person's options be her best guide to maximizing actual utility in a case where she knows which of her options will have the highest actual utility and this option *does not* have the highest expected utility? More to the point, if purism is true, why would the expected utilities of your options be your best guide to maximizing actual utility *in Case B*? If purism is true, refraining from answering will have the highest expected utility in Case B, even though, in Case B, you know that you will maximize actual utility by answering 'yes.' But why should we think that, even though your option with the highest expected utility in Case B will not have the highest actual utility in Case B, and even though you know which of your options *will* have the highest actual utility in Case B (namely, answering 'yes'), the expected utilities of your options are still your best guide to maximizing actual utility—not just in general or on average, but in Case B?

Here purists might invoke what we might call the 'traditional view' of expected utilities, tracing back to Ramsey (1926). On this view, we never have direct access to the world, so, when we decide what to do, we're always in effect "guessing" what the world is like—as Ramsey puts it (p. 183). We *do* have direct access to our relevant values and credences, however. Or, at least, we have much better access to our relevant values

and credences than we have to the external world. Since expected utilities weight our values by our credences, the expected utilities of our options are thus *always* our best guides to maximizing actual utility, and this means they're your best guide to maximizing actual utility in Case B.

The traditional view has its virtues, but it's hard to see how *purists* can invoke it to explain why the expected utilities of your options would be your best guide to maximizing actual utility in Case B. After all, if you know that hawks are raptors, if you know that you will get a jellybean if you answer the first question 'yes,' and so on—as purists must say—then you're *not* just guessing what the world is like, and this leaves it unclear how purists can agree with the traditional view that, even in Case B, where you *know* all of the relevant facts about the world, you have better access to your relevant values and credences than you have to the world.[12]

A much better response—one that purists are nearly certain to give—goes as follows: your *evidence* for and against the proposition that answering 'yes' will have the highest actual utility of your options is what matters for your choice between answering the question and leaving it unanswered, not whether your belief in this proposition counts as an item of knowledge. In Case B, however, your evidence for and against the proposition that answering 'yes' will have the highest actual utility of your options does not eliminate every possibility of error. With respect to the question whether you can answer 'yes' in Case B, there is therefore no relevant difference between your knowing that answering 'yes' will have the highest actual utility of your options and your having equally strong but *misleading* evidence that answering 'yes' will have the highest actual utility of your options. So now recall that the researcher has a long list of questions that she intends to ask you. If you have enough evidence to answer the first question 'yes,' there is no reason why you cannot have enough evidence to answer *every* question that she asks you—assuming that they are all roughly as hard for you as the question whether hawks are raptors. But even given this assumption, there *is* a reason why you shouldn't answer every question she asks you: namely, you should see in advance that you would eventually make a mistake and receive an unbearable electric shock. So, even though you know that answering 'yes' will have the highest actual utility of your options, you shouldn't answer 'yes.'

According to this line of thought, the reason you shouldn't answer the first question 'yes,' even though you know that answering 'yes' will have the highest actual utility of your options, is that it should be clear to you in advance that relevantly analogous behavior with respect to the rest of the questions would not maximize actual utility. In response, impurists might argue that relevantly analogous behavior would consist in

---

[12] While I know that hawks are raptors (at least as I sit here typing at my computer), I don't know what credence I have in this proposition. I know it's pretty high, and I know it's also lower than (say) my credence that $1 = 1$, but this is about all I know. There's no value of 'x' for which I can be anywhere near certain that my credence in this proposition is exactly x.

answering only questions for which you know the answer (cf. Williamson 2014). A more conservative reply, however, consists in simply denying that the foreseeable consequences of answering every question adequately explain why you shouldn't answer the first question. Compare Case B to Case C, below.

> **Case C**: You are participating in a study exactly like the one in Case B, except in this study you know that there is an important difference between the first question and all of the rest of the questions on the test. Specifically, you know that the researcher will disconnect the machine after the first question, so that there's no risk of an electric shock on any question *except* the first question.

Can you answer the first question in *this* case? Again, you definitely shouldn't. But the foreseeable consequences of answering the rest of the questions do not explain why, since you know that you will just accumulate a bunch of jellybeans by answering the rest of the questions. Yet the reason you shouldn't answer the first question in Case C is presumably identical (whatever it is) to the reason you shouldn't answer the first question in Case B. So it's still unclear why you shouldn't answer the first question in Case B, if you know that answering this question 'yes' will have the highest actual utility of your options.

Here purists might say that you shouldn't answer the first question in Case B and Case C because, in the long run, not just on the test but over the course of your whole life, you will gain more actual utility by following the rule *maximize expected utility* than by following the rule *if you know which option will have the highest actual utility, pick that option; otherwise, maximize expected utility*. But this reply faces exactly the same challenge. Consider Case D, below.

> **Case D**: You are participating in a study exactly like the one in Case B, except that in this study you know that you are about to die. There is only one question: whether hawks are raptors. If you answer correctly, you will get a jellybean and then die. If you answer incorrectly, you will get an extremely painful electric shock and then die. If you refrain from answering, you will get nothing and then die. Your aversion to being tortured by an electric shock the last moment of your life is much stronger than your desire to enjoy a single jellybean the last moment of your life.

Once again, you shouldn't answer the question whether hawks are raptors. But in this case it's *false* that, over the course of the rest of your life, you will gain more actual utility by following the rule *maximize expected utility* than by following the rule *if you know which option will have the highest actual utility, pick that option; otherwise, maximize expected utility*. So, what else can purists say? They might be tempted to say that you shouldn't answer the first question because, if you *weren't* going to die, then you *would* gain more actual utility over the rest of your life by following the rule *maximize expected utility* than by

following the rule *if you know which option will have the highest actual utility, pick that option; otherwise, maximize expected utility*. But even if the truth of this counterfactual explains why you shouldn't answer 'yes' in Case D (which I think is doubtful), impurists can simply swap out Case D for a case in which this counterfactual is false. Simply imagine a case where (for whatever reason) you have a guardian angel that would make you completely infallible immediately *after* answering the first question, *if* you weren't going to die. Even if you knew for certain that, in every possible world where you don't die immediately after answering the first question, your guardian angel makes you completely infallible from that moment on, you *still* shouldn't answer the first question. But now the explanation can't be that, if you weren't going to die, then you would gain more actual utility over the rest of your life by following the rule *maximize expected utility* than by following the rule *if you know which option will have the highest actual utility, pick that option; otherwise, maximize expected utility*. It seems that the correct explanation (whatever it is) must focus on features of *just* the first question. So, it's still hard to see why you shouldn't answer the first question if you know that answering 'yes' will have the highest actual utility of your options.

Perhaps, at this point, purists will argue that you should do whatever you are rationally required to do, and then argue (on the basis of representation theorems, or something like that) that you are rationally required to refrain from answering, even though you know that answering 'yes' will have the highest actual utility of your options. But this argument won't convince. After all, if you know that answering 'yes' will have the highest actual utility of your options, then you can deduce that you would simply make yourself worse off, and make literally nothing better off, by doing what this argument says you are rationally required to do. Since it is unclear why rationality (thus conceived) should have inviolable normative force, we still have a puzzle. And since this puzzle arises from rejecting (2), purists cannot reject (2) without incurring a significant explanatory burden.[13]

With (2), however, impurists have an easy argument for their view. Together, (1) and (2) jointly entail the following material conditional.

> (3) If you know that hawks are raptors in Case B, then you may answer 'yes' in Case B.

But together with the stipulated details of Case A and Case B, (3) entails impurism. By hypothesis, you shouldn't answer the question in Case B. So, if (3) is true, you don't know that hawks are raptors in Case B. But by hypothesis, you *do* know that hawks are raptors in Case A. Since Case A and Case B are identical except for the consequences of an incorrect answer, and since these consequences aren't *truth-relevant* with respect

---

[13] Note that, because purists think you *know* that you will only make things worse by doing what this argument says you are rationally required to do, purists cannot respond to this puzzle by saying (along with Broome 2007) that the requirements of rationality take wide scope.

to your belief that hawks are raptors, it follows that, if (3) is true, then Case A and Case B give us a pair of encroachment cases. Since impurism is true if there is a single pair of encroachment cases, impurism is true if (3) is true. And since impurists can motivate (3) by simply adding details to Case A and Case B and asking us to reflect on the normative significance of expected *versus* known actual utility, impurists can modify a case forwarded by purists as a problem for impurism and get a pair of cases that seems to *entail* impurism.

3. Problems for Purism

As we noted in the opening paragraphs of this paper, the pragmatic encroachment literature contains two kinds of arguments against purism: PBAs and IBAs. As we also noted in the opening paragraphs of this paper, purists have produced powerful responses to both kinds of arguments. In defense against PBAs, purists have produced cases that look like counterexamples to the various principles upon which PBAs rely, and, in defense against IBAs, they have produced purist-friendly explanations of the intuition that Hannah says something true by uttering the sentence 'I don't know that the bank will be open' in High Stakes, the intuition that Mary says something true by uttering the sentence 'Smith doesn't know that the flight will stop in Chicago' in Cohen's (1999) airport case, and other relevant knowledge-denying intuitions, plus a mountain of data suggesting that the folk don't have these intuitions in the first place. Assume, then, that these purist responses to PBAs and IBAs all succeed. Given this assumption, does the argument in §§1-2 *still* threaten purism?

I think it does. The argument in §§1-2 doesn't rely on anything like the *prima facie* intuition that Hannah says something true by uttering the sentence 'I don't know that the bank will be open' in High Stakes. Instead of asserting that intuitively you don't know that hawks are raptors in Case B and then offering impurism as the best explanation of this intuition, it concedes that intuitively you *do* know that hawks are raptors in Case B and then shows how this intuition leads to trouble. Thus, neither the purist-friendly explanations of the relevant knowledge-denying intuitions nor the data suggesting that the folk don't have these intuitions have any straightforward application to the argument in §§1-2. Now perhaps the empirical data or one of these explanations puts purists in a position where, by making certain moves, they can give an adequate response to the argument in §§1-2. But if so, let's see the moves. If, for example, purists can respond to the argument in §§1-2 with a modification of some WAM, let's see the modification of the relevant WAM. How's it going to go, given that Case B doesn't contain any utterance of any knowledge-ascribing or knowledge-denying sentence?[14] The argument in §§1-2 doesn't rehash any IBA, so purists can't respond to it by simply rehashing one of their objections to an IBA.

---

[14] Cf. Roeber 2014, §7.

Consider the purported counterexamples in the literature, then—Brown's (2008) surgeon case, Anderson's (2015) syringe case, my (2018a) survey case, *etc*. Do any of these cases threaten the argument in §§1-2? The argument in §§1-2 is clearly valid: (1) and (2) jointly entail (3), and, together with the stipulated details of Case A and Case B, (3) entails that purism is false. Impurists can just stipulate that (1) is true, so, to resist (3), purists must produce arguments against (2). But it's hard to see how any of the purported counterexamples in the literature could threaten (2). After all, (2) says just that, *in Case B*, either you shouldn't answer 'yes,' or you don't know that answering 'yes' will have the highest actual utility of your options. Even if Brown's surgeon case is a counterexample to Hawthorne and Stanley's Knowledge-Reason Principle (2008: 578), or Anderson's syringe case is a counterexample to Fantl and McGrath's KJ (2009: 66), or my survey case is a counterexample to Ross and Schroeder's Knowledge-Action Principle (2014: 262), or some other case is a counterexample to some other principle, it doesn't follow that, in Case B, you shouldn't answer 'yes' even though you know that answering 'yes' will have the highest actual utility of your options. This means that there is no *direct* conflict between (2) and the purported counterexamples in the literature.

Now, of course, (2) follows from KU (below), and some version of one the purported counterexamples in the literature might be a counterexample to KU. (Read it as a strict conditional.)

> (KU)   If *S* knows that φ-ing will have the highest actual utility of her options in *c*, then she may φ in *c*.

But of course, (2) might be true while KU is false. To illustrate, suppose Williamson (2014) is correct that you might know that *p* even though you are rationally required to believe that you *do not* know that *p*, and suppose we get a counterexample to KU from a version of Brown's surgeon case where (a) you know that performing the surgery without double-checking the patient's records will have the highest actual utility of your options, but (b) your evidence is "radically misleading" about your epistemic position and (c) you are consequently rationally required to believe that you *do not* know that performing the surgery without double-checking the patient's records will have the highest actual utility of your options (ibid: p. 973). Even if this case is a counterexample to KU, it doesn't follow that (2) is false, since (2) concerns *just the version of Case B that we have been considering*, and impurists can stipulate that, in *this* version of Case B, your evidence isn't radically misleading. To cast doubt on (2), purists must provide specific reasons for thinking not just that KU or some other principle that entails (2) is false, but that (2) is *itself* false. As we have already pointed out, however, a successful argument against (2) will require explaining why, *in the version of Case B that we have been considering*, you shouldn't answer 'yes' even though you know that answering 'yes' will have the highest actual utility of your options.

Thus, as far as I can tell, even assuming that the various objections to IBAs and PBAs in the literature all succeed, the argument in §§1-2 still poses a serious challenge for purism. Now of course, if purism is false, then truth-irrelevant factors can affect our knowledge. Since the argument in §§1-2 relies on the stipulation that you know that hawks are raptors in Case A, perhaps purists will want to say that, because it's wildly implausible that truth-irrelevant factors can affect our knowledge, we should insist that you *can't* (after all) know that hawks are raptors in Case A. But as Grimm (2011) emphasizes, there are ways of characterizing the dependence of knowledge on truth-irrelevant factors where the conclusion that truth-irrelevant factors can affect our knowledge doesn't seem at all implausible, much less *wildly* implausible. And as I argue elsewhere (Roeber 2018b), knowledge seems to depend on truth-irrelevant factors even if *purism* is true, so purists can hardly object that *impurism* makes knowledge depend on truth-irrelevant factors. This means that purists can't resist the stipulation that you know that hawks are raptors in Case A by simply pointing out that, if this stipulation is permissible, then truth-irrelevant factors can affect our knowledge.

Now perhaps, at this point, purists will search for a theory of knowledge that entails that, while it's possible to know *some* things, it's *not* possible to know that hawks are raptors. If purists could motivate a theory of knowledge like this, they could deduce that you don't know that hawks are raptors even in Case A, and do so without trivializing the debate by retreating to full-blown skepticism. But it's unlikely that purists would find what they were looking for. For first, it's hard to imagine what *independent* motivation they could produce for accepting a theory of knowledge like this. And second, even if purists could find some independent motivation for accepting a theory like this, it's hard to see why impurists couldn't just swap out Case A and Case B for a pair of cases built around a proposition that *is* knowable by the lights of the non-skeptical theory that we are imagining purists defending. So it's doubtful that purists could successfully avoid the stipulation that you know that hawks are raptors in Case A by retreating to some non-skeptical theory of knowledge that entails that it's impossible to know that hawks are raptors.

The problem for purists, then, is this. The best strategy for resisting the argument in §§1-2 requires rejecting (2), which requires saying that you shouldn't answer the first question 'yes' in Case B, even though you know that answering 'yes' will maximize *actual* utility in Case B—which seems absurd. In order to dissolve this apparent absurdity, purists must explain why (for example) the expected utilities of your options must govern your decisions, not just usually or often, *but in Case B*, even though you know that maximizing expected utility would just prevent you from maximizing what you know matters more: actual utility. But as the preceding discussion of (2) shows, it's not clear how this explanation would go. So, at the very least, the argument in §§1-2 saddles purists with a significant explanatory burden. And since impurists can simply deny both that you know that hawks are raptors in Case B and that you know that

answering 'yes' will have the highest actual utility of your options in Case B, purists have a significant explanatory burden that *impurists* lack. So even if purists have given adequate responses to all of the PBAs and IBAs in the literature—as I think they have—purists haven't given any adequate response to the *best* objection to their view.[15]

References

Anderson, Charity (2015), "On the Intimate Relationship of Knowledge and Action," *Episteme*, Vol. 12, No. 3: pp. 343-53.

Bach, Kent (2005), "The Emperor's New 'Knows'," in *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, eds. G Preyer and G Peter (Oxford: Oxford University Press).

Briggs, Rachel (2014), "Normative Theories of Rational Choice: Expected Utility," *Stanford Encyclopaedia of Philosophy*.

Broome, John (2007), "Wide or Narrow Scope?," *Mind*, Vol. 116, No. 462, pp. 359-70.

Brown, Jessica (2006), "Contextualism and Warranted Assertability Manoeuvers," *Philosophical Studies*, Vol. 130: pp. 407-35.

——— (2008), "Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning," *Noûs*, Vol. 42, No. 2: pp. 167-89.

——— (2012), "Practical Reasoning, Decision Theory, and Anti-Intellectualism," *Episteme*, Vol. 9, No. 1: pp. 43-62.

Buckwalter, Wesley (2010), "Knowledge is not Closed on Saturday: A Study in Ordinary Language," *Review of Philosophy and Psychology*, Vol. 1, No. 3: pp. 395–406.

Buckwalter, Wesley and Jonathan Schaffer (2015), "Knowledge, Stakes and Mistakes," *Noûs*, Vol. 49, No. 2: pp. 201-34.

Cohen, Stewart (1999), "Contextualism, Skepticism, and the Structure of Reasons," *Philosophical Perspectives*, Vol. 13: pp. 57-89.

——— (2005), "Knowledge, Speaker and Subject," *The Philosophical Quarterly*, Vol. 55, No. 219: pp. 199-212.

———— (2012), "Does Practical Rationality Constrain Epistemic Rationality?," *Philosophy and Phenomenological Research*, Vol. 85, No. 2: pp. 447-55.

DeRose, Keith (1992). "Contextualism and Knowledge Attributions," *Philosophy and Phenomenological Research*, Vol. 52: pp. 913-29.

———— (2009). *The Case for Contextualism* (Oxford: Oxford University Press).

Fantl, Jeremy and Matthew McGrath (2002), "Evidence, Pragmatics and Justification, *The Philosophical Review*, Vol. 11, No. 1: pp. 67-94.

———— (2009), *Knowledge in an Uncertain World* (Oxford: Oxford University Press).

Feltz, Adam, and Chris Zarpentine (2010), "Do you Know More When it Matters Less?," *Philosophical Psychology*, Vol. 23, No. 5: pp. 683-706.

Grimm, Stephen (2011). "On Intellectualism in Epistemology," *Mind*, Vol. 120, No. 479: pp. 705-33.

Hawthorne, John and Jason Stanley (2008), "Knowledge and Action," *The Journal of Philosophy*, Vol. 105, No. 10.

Locke, Dustin (2017), "Implicature and Non-Local Pragmatic Encroachment," *Synthese*, Vol. 194: pp. 631-54.

Lutz, Matt (2014), "The Pragmatics of Pragmatic Encroachment," *Synthese*, Vol. 191: pp. 1717-40.

May, Joshua, Walter Sinnott-Armstrong, Jay Hull and Aaron Zimmerman (2010), "Practical Interests, Relevant Alternatives, and Knowledge Attributions: An Empirical Study," *Review of Philosophy and Psychology*, Vol. 1, No. 2: pp. 265–73.

Nagel, Jennifer (2008), "Knowledge Ascriptions and the Psychological Consequences of Changing Stakes," *Australasian Journal of Philosophy*, Vol. 86, No. 2: pp. 279-94.

———— (2010), "Epistemic Anxiety and Adaptive Invariantism," *Philosophical Perspectives*, Vol. 24: pp. 407-435.

Phelan, Mark (2014), "Evidence that Stakes Don't Matter for Evidence," *Philosophical Psychology*, Vol. 27, No. 4: pp. 488-512.

Reed, Baron (2010), "A Defense of Stable Invariantism," *Noûs*, Vol. 44, No. 2: pp. 224-44.

——— (2013), "Fallibilism, Epistemic Possibility, and Epistemic Agency," *Philosophical Issues*, Vol. 23, No. 1: pp. 40-69.

Roeber, Blake (2014), "Minimalism and the Limits of Warranted Assertability Maneuvers," *Episteme*, Vol. 11, No. 3: pp. 245-60.

——— (2018a), "The Pragmatic Encroachment Debate," *Noûs*, Vol. 52, No. 1: pp. 171-95.

——— (2018b), "Anti-Intellectualism," *Mind*, Vol. 127, Issue 506, pp. 437-66.

Rose, David *et al* (2017), "Nothing at Stake in Knowledge," *Noûs*, DOI: 10.1111/nous.12211.

Ross, Jacob and Mark Schroeder (2014), "Belief, Credence and Pragmatic Encroachment," *Philosophy and Phenomenological Research*, Vol. 88, No. 2: pp. 259-88.

Rysiew, Patrick (2001), "The Context Sensitivity of Knowledge Attributions," *Noûs*, Vol. 35: pp. 477-514.

Schaffer, Jonathon and Joshua Knobe (2013), "Contrastive Knowledge Surveyed," *Noûs*, Vol. 46, No. 4: pp. 675-708.

Stanley, Jason (2005), *Knowledge and Practical Interests* (Oxford: Oxford University Press).

Turri, John (2017), "Epistemic Contextualism: An Idle Hypothesis," *Australasian Journal of Philosophy*, Vol. 95, No. 1: pp. 141-56.

Weatherson, Brian (2012), "Knowledge, Bets and Interests," in *New Essays on Knowledge Ascriptions*, eds. J. Brown and M. Gerken (Oxford: Oxford University Press).

Williamson, Timothy (2014), "Very Improbable Knowing," *Erkenntnis*, Vol. 79: pp. 971-99.