

Knowledge, Skepticism, and Defeat ed. Brandon Fitelson, Rodrigo Borges, Cherie Braden, Springer, 2019, pp. 25-38.

Epistemically Useful Falsehoods¹

Catherine Z. Elgin

Pummeled by a relentless barrage of Gettier cases, epistemologists tend to conclude that false beliefs never figure essentially in the justification of a bit of knowledge. Peter Klein demurs. He agrees that Gettier sequences do not result in knowledge. And he agrees that the reason they do not is that a false belief plays an ineliminable role in the agent's inference. But, he urges, not all inferential paths that involve false beliefs essentially lead to Gettier cases. There are, he maintains, epistemically useful falsehoods; they figure essentially in the justification of bits of knowledge. A chain of inference contains a falsehood essentially (or ineliminably) just in case the inference would not be valid if the false statement were eliminated. I am not concerned here to defend or challenge Klein's defeasibility analysis of knowledge (see Klein 1976, 1981). But whether or not a defeasibility theory is correct, it is worth asking whether inferences that lead to empirical knowledge can contain falsehoods essentially.

Here I discuss the position Klein develops in 'Useful False Beliefs' (2008). My goal is not to ask how the position meshes with the rest of his epistemology. It is not even to applaud or object to the fine-grained details of Klein's proposal. It is to ask whether he has supplied good reason to think falsehoods can ever figure ineliminably in the justification needed for inferential empirical knowledge. I think he has.

Klein provides several cases. For expository purposes, I've modified them slightly, while

¹ I am grateful to Samuel Elgin and Robert Shope for useful discussions of the issues raised in this paper.

preserving what I take to be the epistemological features that are relevant to my discussion.

APPOINTMENT: On the basis of an apparent memory, Paul believes that his secretary told him on Friday that he has an appointment on Monday. From that belief he infers that he has an appointment on Monday, which in fact he does. Although his secretary told him about the appointment, she told him on Thursday, not on Friday. Nevertheless, it seems, Paul knows that he has an appointment on Monday.

SANTA CLAUS: On Christmas Eve, Virginia's parents tell her that Santa will put presents under the tree on that night. She believes them and infers that there will be presents under the tree on Christmas morning. Klein contends that Virginia knows it.

AVERAGE RAINFALL: Weatherman believes that the average annual precipitation in northwest Montana is about 13 inches, because he believes that accurate records have been kept for over eighty years and during that period the average comes to 13 inches. The average rainfall is about 13 inches, but accurate records were kept for only seventy-nine years. Still, it seems, he knows the average rainfall.

PTOLEMAIC ASTRONOMER: On September 2, 1203 a Ptolemaic astronomer uses the deferent and epicycle orbits of the Earth, Mars and the Sun and what he takes to be their (then) current positions, to predict that *ceteris paribus* Mars will be visible from Earth on September 2, 2003. (The 'ceteris paribus' controls for the possibility of clouds and the premature annihilation of the relevant celestial objects, but not for the possibility that Ptolemaic astronomy is incorrect.) If the assigned orbits and then-current relative positions of the three bodies allow for sufficiently accurate extrapolations, Klein maintains, the astronomer knows that, *ceteris paribus*, Mars will be visible on September 2, 2003. (Klein, 2008, 37-38)

These cases -- call them Klein cases -- cover a range of grounds for knowledge. APPOINTMENT rests on (apparent) memory; SANTA CLAUS, on testimony; AVERAGE RAINFALL, on recorded empirical information; PTOLEMAIC ASTRONOMER, on observation and calculation.² The range suggests that if Klein is right, epistemic reliance on useful falsehoods is not a local or parochial matter. We should expect to find them in any realm where empirical knowledge stems from inference. You don't have to accept all of Klein's cases. Some may seem more plausible than others. But if any strike us as cases of knowledge, the conviction that ineliminable falsehoods in the inferential path are an anathema to empirical knowledge needs to be reconsidered.

Klein, of course, does not rest his conclusion on the plausibility of these examples. Their role is to illustrate the phenomenon he is interested in and to elicit intuitions. They should, he hopes, suggest that despite the plethora of Gettier cases, it is not so obvious that false beliefs in epistemically ambitious empirical inferences are inevitably pernicious. Some, if not all, of the cases are apt to strike us as instances of knowledge gleaned from inferences that essentially involve false beliefs.

Klein develops an account of how and when a path involving a falsehood yields knowledge. It is, in many respects, a remarkably open-minded account. He takes pains to show that it is compatible with foundationalism and coherentism as well, of course, with infinitism. This is not one of his more radical theories. Still, there are elements of it that give me pause. So after explicating the theory and identifying those elements, I will sketch an alternative way of arriving at the same conclusion. My goal is not to show that he is wrong or that my approach is

2 In Klein's presentation, the Ptolemaic Astronomer has his students make the calculation, based on what they have been taught. If teaching is testimony, this is an instance of useful falsehoods figuring in knowledge based on testimony as well. I simplified the case. Placing the inferential burden on the students is, for my purposes, an unnecessary epicycle.

better. It is to suggest that a variety of disparate epistemological theories will and should have a place for epistemically useful falsehoods. In short, you can't get out of this by simply objecting to the details of Klein's position.

Klein's account

Assuming that propositions are the contents of beliefs (Klein 2008, p. 27), Klein distinguishes between two sorts of justification. *Propositional justification* pertains to the contents of beliefs. 'A proposition h is propositionally justified for S just in case S has an epistemically adequate basis for h ' (Klein, 2008, p. 28). In the cases that concern us, the epistemically adequate basis is either another proposition that S believes or a proposition that is available to S , where to be available to S is either to be one of the contents of S 's mental states or to be appropriately connected to S 's mental states. Ideally one would like to hear more about what it takes to be appropriately connected. For our purposes, however, it suffices that entailment is an appropriate connection. If the content of S 's belief that g entails the content of her belief that h , then g affords her an epistemically adequate basis for h .

Propositional justification pertains belief contents. It is keyed to the epistemic agent in that the resources it draws on are resources available to her. So considerations that are inaccessible to S cannot propositionally justify h for her, even if in an absolute sense they justify h . But a consideration that is accessible to an agent need not actually be accessed by her. An important consequence is that h can be propositionally justified for S without S 's being aware that it is. She might fail to draw the needed inference. Indeed, she might be incapable of drawing the inference. Nevertheless, if h is entailed by S 's belief that g , her belief that h is propositionally justified for S .

Doxastic justification pertains to beliefs themselves, not (or anyway not directly) to their

contents. According to Klein, 'A belief is doxastically justified when it has an appropriate causal pedigree' (Klein 2008, p. 28). How *S* came to harbor the belief that *h* determines whether that belief is doxastically justified. Doxastic justification is the sort of justification that is at issue in determining whether *S* is justified in holding the belief. It is the justification needed for justified true belief.

Propositional justification and doxastic justification are mutually independent. A subject is doxastically justified in believing that *g* if she came to that belief in the right way. But a belief whose causal pedigree is impeccable can still fail to be propositionally justified because there is information available to her that renders the belief content unjustified. On the flip side, she might have a propositionally justified belief *j* that is in fact appropriately connected to other justified beliefs, but if she arrived at it by reading tea leaves and never even attempted to infer it from her other justified beliefs or from her other evidence, she lacks doxastic justification for *j*. Moreover, a proposition that she does not believe can be propositionally justified for her.

In Gettier cases as in Klein cases, a false belief plays an essential role in arriving at a true conclusion. *S*'s belief that someone in the class owns a Ford is caused by her belief that Nogot, who is in the class, owns a Ford. Paul's belief that he has an appointment on Monday is caused by his belief that his secretary told him on Friday that he has a belief on Monday. Virginia's belief that there will be presents under the tree is caused by her belief that Santa Claus will put presents under the tree. Weatherman's belief that the average rainfall is about 13 inches is caused by his belief that the average over the more than eighty year record of annual rainfalls is 13 inches. The astronomer's belief that Mars will be visible on September 2, 2003 is caused by his belief that celestial objects orbit and will continue to orbit the Earth in accordance with the laws of Ptolemaic astronomy. We might conclude then that Klein cases just are Gettier cases. If

so, the intuition that they yield genuine knowledge is wrong.

This does not seem plausible. The difference between standard Gettier cases and Klein cases is that in Klein cases the false beliefs seem to be false in ways that are irrelevant to their epistemic role. Thursday? Friday? Who cares when Paul's secretary told him about the appointment? What matters is that she told him. Seventy-nine years? Eighty years? Who cares exactly how many years they have been keeping rainfall records, given that they have been doing so long enough to establish a reliable record. And so forth. The issue is how to capture this difference.

Aside: even if we want to be hard-nosed and insist that Klein cases are Gettier cases -- that they yield true belief but not knowledge -- there is something epistemologically interesting going on. Unlike standard Gettier cases, they certainly seem to yield knowledge. However we accommodate them, we should acknowledge the fact that they seem quite different from standard Gettier cases.

Klein calls the false beliefs that figure in Klein cases 'useful falsehoods'. Rather than leading us astray, they point us in the right direction. He suggests that:

The belief that *uf* is a useful falsehood to *S* (for acquiring knowledge that *h*) by producing a doxastically justified belief that *h* iff:

1. *uf* is false.
2. The belief that *uf* is doxastically justified for *S*. That is, *S* is justified in believing *uf*.
3. The belief that *uf* is essential in the causal production of the belief that *h*. Without it, there would be no causal chain leading to *S*'s belief that *h*.
4. *uf* propositionally justifies *h*.
5. *uf* entails a true proposition, *t*.

6. t propositionally justifies h .

7. Whatever doxastically justifies the belief that uf for S also propositionally justifies t for S . (Klein 2008, p. 49)

Klein's account depends on the fact that a subject's doxastic justification and her propositional justification are distinct. In the cases of interest, her doxastic justification for h is flawed because she arrives at a truth via a false belief. That false belief would ordinarily be deemed a defeater. But in cases where the falsehoods are useful, Klein suggests that they have a distinctive property: each such falsehood entails a relevant truth that is propositionally justified. 'Paul's secretary told him on Friday that he has an appointment on Monday' entails 'Paul's secretary told him he has an appointment on Monday.' 'Santa Claus will put presents under the tree' entails 'Someone will put presents under the tree.' The subject need not believe or even entertain the truth in question. What is at issue is propositional justification, not doxastic justification. But because, owing to the entailment, the road not taken is close enough to the inferential route actually traversed, the subject knows.³

It is important that the road not taken is not utterly alien to S . Because t is propositionally justified for S , it consists of considerations that are available to her. They are, unfortunately, considerations she did not actually draw on in coming to her belief. Whether or not this shows that useful falsehoods contribute to the production of knowledge, it at least yields a plausible explanation of why Klein cases are apt to strike us as cases of knowledge.

The mere existence of such a road not taken might seem inadequate to secure knowledge. If S does not believe that t , the fact that t is entailed by something she believes may seem too flimsy a basis for claiming that uf provides knowledge. We might seek to shore it up. Perhaps

³ Or is in a position to know. We are here talking only about the justification condition on knowledge.

we can connect uf and t by insisting that because S believes uf , she believes t . Then we could say that t is the real link in the chain of doxastic justification. This is surely false. t might never cross her mind. In that case, it is no part of the doxastic chain. Perhaps we should say, 'Because S believes uf , S would believe t , if she were to entertain it.' As a general matter, this too is false. Were the Ptolemaic astronomer to entertain the idea that, for purposes of prediction, it is just as if Ptolemaic astronomy was true, he might well balk. As Klein notes, he could be incapable of 'just as if' reasoning. Alternatively, he might resist using it in this case. He might think that if Ptolemaic Astronomy is false, there is no reason whatsoever to think that planetary positions will be just as they would be if the theory were true. Virginia could balk as well. Klein stipulates that Virginia is mature enough to know that her parents are generally reliable truth tellers. Her issue here is different. Given the number of days on which her parents have failed to shower her with gifts, she might well think, 'It's either Santa Claus or nothing!' It is then implausible that S has another belief that is the real link in the chain of doxastic justification.

Klein concedes that the relation between uf and t might be looser than entailment. Perhaps probabilifying t or rendering t highly plausible would suffice. We might also worry that entailment is too generous. If t follows from uf by an obscure and convoluted (but logically impeccable) route, we might doubt that t supplies the needed backup. I'm not going to address these concerns. If need be, we could restrict the requirement to obvious entailments and expand it to include other suitably strong relations of support. Here I am interested in seeing whether knowledge can be underwritten by useful falsehoods in anything like the way that Klein suggests. If it can, we can worry about how or whether we need to tweak things.

Concerns

Klein's general account of what it is to be a useful falsehood is plausible if his

assumptions are sound. He has argued for them elsewhere, and they are not unduly radical. Still, I have qualms about the underlying metaphysics. I will raise them and go on to suggest an alternative route to the same conclusion. My goal in doing this is not to maintain that my way is better than Klein's. It is to suggest that there is something deeply right about his conclusion, something that cannot be easily evaded.

Klein maintains that doxastic justification depends on the causal history of a belief. I have a couple of worries about this. The first is this: A subject might come to the true belief that p unjustifiably. She overlooked countervailing evidence, believed an unreliable informant, engaged in sloppy reasoning, or jumped to conclusions. Suppose, however, that, having formed the belief, she continues to amass evidence. The additional evidence enables her to hold that p for good and sufficient reasons. The belief's disreputable causal ancestry is then swamped by trustworthy supports. Tabloid readers came to believe that O. J. Simpson abused his wife because it was so reported in the *National Enquirer*, a notoriously unreliable scandal sheet. Subsequently, more reputable news sources uncovered strong evidence of his abuse. The credulous tabloid readers should not be permanently barred from being doxastically justified simply because they originally formed their belief on the basis of an unreliable source, given that they now take trustworthy news reports as the basis for their conviction. This sort of objection has been raised against causal theories in the past. The obvious and sufficient response is to recognize not just generating causes but also sustaining causes. Then the tabloid readers who now base their belief on reputable news sources are in a position to know.

There is, however, more to worry about than occasional epistemic sloppiness. According to psychology, the etiology of many beliefs is not particularly estimable. In forming beliefs, we overlook, over- or under-emphasize, cut corners, distort what we find in order to bring new

evidence into accord with our presuppositions. We are guided (even blinded) by stereotypes. We are prey to wishful-thinking, self-deception, confirmation bias. We seem to go out of our way to ignore base rates. We are frequently in denial. We suffer from implicit bias, stereotype threat, and all manner of repressions and distortions. And that's if we are normal, mentally healthy cognizers! If doxastic justification turns on how we actually come to our beliefs, much of what we take to be knowledge falls short. Despite what you may think of it, your conviction that Bernie Sanders is more liberal than Hillary Clinton is not caused by your appreciation of their voting records. It is largely due to unresolved issues from your childhood. Your grandfather always brought you presents; your grandmother told you to sit up straight. Your non-costodial father was generous; your mother, with inadequate child support, counted pennies. This is not to deny that we reason inferentially. But one person's *modus ponens* is another's *modus tollens*. If *S* believes that *p* (and wants to be consistent), her recognition that *p* entails *q* presents her with an option. She can either endorse *q* or repudiate *p*. What factors she considers salient, how she weighs the evidence, how she assesses plausibility, and a host of other matters, some of them epistemically quite dubious, influence her choice.

The cases of heuristics and biases that have dominated the literature are of course cases in which, on account of these shortcomings, the agent was not justified in the belief she formed (Gilovich et al., 2002). But, psychologists insist, many of the heuristics and biases are valuable. Although they occasionally engender false beliefs, they often lead quite efficiently to true ones. I'm not going to argue that we should just validate them and consign the failures to the realm of epistemic misfortune. Even if we are tolerant of a measure of epistemic luck, that would be going too far. Rather than depending exclusively on the sequence of considerations that led to the adoption of a belief, or on the full network of considerations that sustain it, I suggest,

doxastic justification depends on those elements of the network that endow the belief with credibility. Which these are is a normative, and not merely a causal matter.

Recently, philosophers have urged that causation is not the only sort of metaphysical dependence (Rosen 2010). Some relations are grounded in others; some states of affairs obtain in virtue of others; some facts hold because others do. If this is so, it opens the door to the idea that a belief's doxastic justification might turn on its standing in an appropriate dependence relation, which need not be (but may be) causal. I am not going to develop this idea here. Much needs to be worked out before we have a clear sense of what the non-causal dependence relations are and of what functions they perform. My point is simply this: Doxastic justification is the agent's justification for believing. It is what she could in principle adduce if asked, 'Why do you believe that?' Klein needs there to be justification for believing (and not just for the proposition believed). He may have concluded too quickly that the cause of the belief is the only thing that could provide that justification. I suggest then that we should say that doxastic justification is that which justifies an agent in holding a belief, and for now remain neutral about whether it consists of causes, grounds, or collateral support of a different kind.

My other metaphysical worry concerns Klein's contention that the contents of beliefs are propositions. Standardly, propositions are held to be that which truth-apt sentences express. And two sentences express the same proposition just in case they are co-intensional. Moreover, it is held, the co-intensionality condition is sometimes satisfied; some pairs of sentences are co-intensional. If 'vixen' and 'female fox' are synonymous, then 'Vixens are fierce' and 'Female foxes are fierce' are co-intensional. They express the same proposition. Any world in which vixens are fierce is one where female foxes are fierce; any world in which female foxes are fierce is one where vixens are fierce. But the human capacity for ignorance and error is legion. Even if

'vixen' and 'female fox' are synonymous, it is entirely possible for Harry to believe that vixens are fierce without believing that female foxes are fierce. Even if 'London is beautiful' and 'Londres est belle' express the same proposition, Pierre can believe one but not the other (Kripke, 2011). Propositions, as standardly characterized, are not sufficiently fine-grained to reflect belief contents. Such contents are evidently hyperintensional.

Some, such as Bealer, reject the co-intensional criterion and maintain that propositions are hyperintensional too (Bealer 1998). But the criterion is so deeply entrenched that it is preferable, I think, to introduce new terminology. Let us call the content of a belief its hyperintension, and say that two beliefs share a content just in case they are co-hyperintensional. It seems plausible that the beliefs of two agents could share a content. They might both believe that vixens are fierce. What is unclear is whether a single agent could harbor distinct mental representations that were co-hyperintensional. In that case she would believe the very same content under different modes of representation. A more plausible alternative is that different modes of representation express different hyperintensions. Then Harry's belief that vixens are fierce and his belief that female foxes are fierce have different contents, even though he realizes that 'vixen' and 'female fox' are synonymous.⁴

Alternative proposal

Initially, Klein's strategy for evading Gettier cases was an anti-defeater strategy. (Klein 1976). Knowledge, he maintained, is justified true belief, where the justification in question is doxastic and there are no genuine defeaters in the doxastically operative inferential chain. A defeater is a false belief that is ineliminable from the inferential chain. A genuine defeater is one

4 Another aside: Although I have put the point in terms of contemporary metaphysics, one can get the same fine grain using Israel Scheffler's inscriptionalism. It is not obvious that we need to inflate our metaphysics to get the resources we need to properly individuate belief contents.(See Scheffler 1955). Nor is it obvious that we need to inflate our metaphysics to get non-causal dependency relations. The resources developed in Goodman's *Structure of Appearance* suffice.(See Goodman 1977).

that is not itself defeated and that is ineliminable from the causal chain. On this account, 'Paul's secretary told him on Friday . . .', 'Santa will bring presents', and 'They have kept accurate records for eighty years' are defeaters.

Thinking that APPOINTMENT, SANTA CLAUS, WEATHERMAN and PTOLEMAIC ASTRONOMER result in knowledge despite involving defeaters, Klein introduced useful falsehoods. Now he holds that:

S knows that *h* only if *S*'s doxastic justification contains no genuine defeaters essentially
or

S's doxastic justification contains essentially only defeaters that entail relevant propositionally justified truths.

Disjunctive analyses are not dreadful, but they are unattractive. We would prefer to avoid them. Moreover, as we saw earlier, a falsehood's entailing a propositionally justified truth does not obviously compensate for its defectiveness as a link in a justificatory doxastic chain.

I suggest then that rather than focusing on what they entail, we take a closer look at the useful falsehoods themselves. The discussion of hyperintensionality highlights that beliefs and other hyperintensional contents have a fine-grained internal structure. Perhaps we can evade the difficulties by exploiting that structure. In 'Truth about Jones', Ullian and Goodman argued that a false statement (or proposition, although they would never use the term) can, despite its falsity, be true about something (Ullian and Goodman, 1977).

'Sam is in New Haven and Gareth is in Delhi'
is false, since Gareth is in Mumbai. But it is true about Sam, for he is in New Haven. Insofar as the truth value of the sentence depends on where Sam is, it is true. This, I suggest, is what makes the falsehoods in Klein cases useful. They are true about the factors that bear on or figure in the

agent's inference. They owe their overall falsity to factors that do not matter.

Philosophers frequently focus on truth-makers. Here we are interested in their counterparts -- falsity-makers. Let's look at Klein's cases. What makes Paul's relevant belief in APPOINTMENT false is that his secretary told him about the appointment on Thursday rather than on Friday. What makes Virginia's relevant belief false in SANTA CLAUS is that it is her parents rather than Santa Claus who will put the presents under the tree. Although they make the propositions in question false, these factors are inferentially idle. They do no work in the agents' inferences.

That 'on Friday' plays no inferential role in APPOINTMENT seems obvious. The phrase could simply be eliminated from the statement of Paul's belief, and the inference would still go through. SANTA CLAUS is a bit trickier. To just drop the term 'Santa Claus' leaves an open sentence. For Virginia to have a content capable of being believed, a suitable noun phrase must replace it. But beyond the requirement that the replacing term putatively denote an agent or entity seemingly capable putting presents under a tree, it doesn't really matter what the replacement is. To see this, imagine that Virginia's parents mumbled. Virginia didn't quite catch who exactly they said would put presents under the tree. For Virginia's inference, this doesn't matter. Both 'on Friday' and 'Santa Claus' are, in Strevens terms, not difference-makers (Strevens, 2008).

It is less clear, however, that the falsity-making factor is inferentially idle in AVERAGE RAINFALL. Weatherman's inference requires that the average rainfall was recorded for a long time. That the duration was greater than eighty years doesn't matter, but that it was a considerable number of years does. Just how many years is not obvious. It depends on how variable rainfall is from year to year. If it is quite variable, then a good many years will be

required to get an average we can trust; if it is pretty constant, relatively few years will do. Luckily, that is a problem for meteorology. We don't have to solve it. We have a different problem -- the fact that the falsity-maker seems to be doing some work. It could not simply be eliminated from Weatherman's inference in the way that the falsity-makers in APPOINTMENT could. What work is it doing? I suggest that it is facilitating the inference because 'They've been recording the weather for eighty years' is true enough (See Elgin 2017). Eighty years, the duration mentioned in the sentence, is a good many years. In Weatherman's statement, the term 'eighty years' not only denotes number in the extension of the predicate 'a good many years', it exemplifies that it does so. The precise number of years mentioned in Weatherman's statement is inferentially idle. That it belongs to the extension of 'a good many years' is not.

PTOLEMAIC ASTRONOMER seems to present an even harder problem. Here we have not an infelicitous phrase that could be suitably paraphrased or eliminated from an otherwise estimable sentence. The Ptolemaic astronomer's driving assumption -- the geocentric model of the heavens -- is wildly incorrect. How could the poor man possibly know what planets would be visible in 800 years when he is dead wrong about what is circling what?

The amazing thing about Ptolemaic astronomy is that it works. Even today, celestial navigators both on the sea and in the desert adopt a Ptolemaic perspective because for getting around in vast, trackless realms, it is practical to treat the earth as motionless. And because the Ptolemaic astronomers were very good at their job, to a surprisingly precise degree Ptolemaic astronomy preserves the relevant appearances. Contemporary celestial navigators get where they want to go. To determine what will be seen in the future, it is appearances that need to be preserved. Hence here too we can resort to the true enough strategy. If all we care about is the appearances, then what is true enough is what correctly describes and predicts the appearances. I

suggest then that the astronomer knows whether Mars will be visible from Earth 800 years in the future because, in his reasoning, Ptolemaic astronomy functions as no more than a calculational device. He may use it because he believes the constitutive claims of the theory. But in his reasoning the constitutive Ptolemaic assumptions do no more than supply an inference ticket that brings brings him from accurate inputs to accurate outputs. Despite its being very far from the truth, it is true enough to underwrite the inference.

If this is right, a constructive empiricist like van Fraassen could maintain that he knows the consequences of scientific inferences. We do not, the constructive empiricist maintains, know which empirically adequate theory is true. But if useful falsehoods lead to knowledge, it simply doesn't matter whether he arrived at his belief that p via a true theory or via a usefully false, empirically adequate alternative (see van Fraassen 1980).

We can test my suggestion by seeing what happens when we vary some of Klein's examples. Begin with APPOINTMENT. Suppose, that Paul's secretary (although excellent in other respects) is a bit of a scatter brain. She is, in general, less than reliable about telling him what is in his calendar. But recognizing that with the weekend impending it is important to make sure Paul knows his schedule for Monday, she tends to take more care and be far more reliable in what she tells him on Fridays than she is about what she tells him on other days of the week. If that were so (and Paul was aware of this fact), it is plausible that he would know that he has an appointment if she told him on Friday, but not if she told him on Thursday. In that case 'on Friday' would not be not idle in his inference. It would be a difference-maker.

Consider now AVERAGE RAINFALL. Suppose the annual rainfall in Northwest Montana varies a lot from year to year. There are trends, and in the long run they smooth out. But an average taken over only a few years would likely be misleading. Meteorologists have

concluded that they need a minimum of 80 years to get numbers they can trust. In that case, Weatherman's conclusion, based on only 79 years, would not qualify as knowledge. That 79 years is and exemplifies that it is a considerable number of years would not make Weatherman's grounds true enough in this context. It may seem a mite churlish to insist that evidence from merely 79 years is too sparse if 80 years suffices. But a slippery slope threatens. If meteorologists were justified in setting the threshold where they did, then failing to meet that threshold is good reason to think that Weatherman does not (quite) know.

The account I offer does not rely on a road not taken. It focuses on the beliefs and belief contents the agent actually has and uses in her inferences.⁵ It exploits the fact that only some aspects of the belief contents figure in her inferences. If she is right about them (where this includes their being true enough, where that is all that is required), she is on solid epistemic ground. Moreover, my account is not disjunctive. I do not give one condition for standard justified true beliefs and another for useful falsehoods. Rather, the position I have sketched enables us to say, as Klein originally did, that *S* knows that *h* only if *S*'s doxastic justification contains no genuine defeaters essentially. This is so because not every ineliminable falsehood is a defeater. If the falsehood is true enough -- if it is either true or it is false in ways that make no difference to the function it is performing in the agent's doxastic economy -- the agent knows.

Klein considers relying on useful falsehoods an epistemic defect, albeit not a fatal one. It would be preferable all things considered if Paul had simply reasoned from 'My secretary told me I have an appointment on Monday' to 'I have an appointment on Monday.' It would be preferable if Weatherman had reasoned from 'They've been keeping accurate records for 79

5 Normally I rely on L. Jonathan. Cohen's notion of acceptance rather than belief. An agent accepts that *p* just in case she is willing to use *p* as a basis for assertoric inference or action when her ends are cognitive. (Cohen 1992). Here I speak of beliefs to bring my position into contact with Klein's. But belief seems not to be the relevant attitude when we ask how thought experiments, and idealizations function epistemically. So talk of belief is rather out of character for me.

years' (or 'for a considerable number of years') to 'The average rainfall is about 13 inches.' He may be right. His assessment handles his cases and comports with the veritistic spirit of most epistemology. But I'm not so sure. The sort of cases I normally focus on concern idealizations and simplifying assumptions. In these cases, the utility of the falsehood lies precisely that it prescind from features of actual situations that do not matter.

Consider Galileo's famous thought experiment about falling bodies. According to Aristotelian physics (and common sense) heavier bodies fall more quickly than lighter bodies. So, Galileo says, consider a cannon ball and a musket ball. The cannon ball is heavier and should fall faster than the musket ball. Tie the two together with a rope. The new, composite object, is heavier than the cannon ball, for we've added the weight of the rope and the musket ball. Imagine dropping an untethered cannon ball (of the same size, shape, and material as the tethered cannon ball) and the composite object. Which will fall more more quickly? Aristotelians answer: the composite object. But, Galileo notes, part of the composite object is the musket ball, which Aristotelians say, must fall more slowly. As it falls, it serves as a brake, slowing the rate of fall of the composite object. That being so, the composite object should fall more slowly than the cannon ball alone. But it cannot fall both more quickly and more slowly than the cannon ball. So we have a contradiction. It follows that the rate of fall of an object must be independent of its weight.

This is a familiar, beautiful thought experiment. It is conclusive. It would not be better if rather than constructing the thought experiment, Galileo did a bunch of tests, dropping things from various towers and measuring their rate of fall. Had he done so, he would have needed to show how he accommodated potentially confounding factors -- air resistance, limitations on the accuracy of his measuring devices, and so forth. He controlled for them in the thought

experiment simply by leaving them out. Their introduction would only mislead.

However lovely, it might seem that this example is not about inferential empirical knowledge. It is an imaginary exercise about the relation between two magnitudes. It seems rather a priori. Indeed, its conclusion is a modal claim. Rate of fall and weight *must be* independent. But from the necessity claim a claim about actuality immediately follows. If rate of fall and weight must be independent, then rate of fall and weight are independent. The inference yields knowledge about a matter of fact.

The thought experiment is what I have called a felicitous falsehood (See Elgin 2017). The difference between my felicitous falsehoods and Klein's useful falsehoods may be just terminological. Perhaps not. In Klein's discussion, useful falsehoods seem to be second-class epistemic citizens. We grudgingly admit that they yield knowledge; but to be honest, we'd rather not. I suggest that there's nothing second class about felicitous falsehoods. They are typically epistemically preferable to the unvarnished truth, for they prescind from factors that threaten to mislead. In the cases Klein considers, useful false beliefs involve some sort of mistake. I'm suggesting here, that some falsehoods -- the felicitous ones -- involve no mistake. Moreover, their utility lies not in their facilitating inferences despite their falsity, but in their facilitating inferences because of their falsity. The issues are complicated and take us into neighborhoods where Klein may not want to venture. But once you open the door to false beliefs playing a legitimate justificatory role, the question 'How large a role?' emerges.

Conclusions

I've made a number of suggestions that Klein and his followers may be loathe to accept. Following Klein, I allowed that doxastic justification is what justifies the epistemic agent in holding her belief. It is about her believing that p , not (directly) about p itself. Contrary to

Klein, I suggested that the dependence relation needed for doxastic justification need not be causal. Another sort of metaphysical dependence might suffice.

I argued that belief contents are exceedingly fine-grained. I suggested that the fine-grainedness figures not just in the propositional justification (or perhaps we should say hyperintensional justification), but also in the affordances a belief provides for doxastic justification. In making an inference, the agent draws on only some features of her belief content. Some features are not difference makers. They perform no function in her inference, so can be ignored. If her belief is false only on account of those features, it is a useful falsehood. As I see it, we need not detour through a road not taken. Rather, we need only attend to the operative elements of the inferential road we are actually on.

One final point: It may seem that Klein has replaced his defeasibility theory with a hybrid account. I disagree. Rather, I suggest, he has discovered that not all false beliefs that lead to inferential empirical knowledge are defeaters.

References

- Bealer, George (1998). 'Propositions' *Mind* 107: 1-32.
- Cohen, L. Jonathan (1992). *An Essay on Belief and Acceptance*. Oxford: Clarendon Press.
- Elgin, Catherine (2017). *True Enough*. Cambridge: MIT Press.
- Gilovich, Thomas, Dale Griffin, Daniel Kahneman (eds.) (2002) *Heuristics and Biases*. Cambridge UK: Cambridge University Press.
- Goodman, Nelson (1977). *The Structure of Appearance*. Dordrecht: Reidel.
- Klein, Peter (2008). 'Useful False Beliefs' in *Epistemology: New Essays* ed. Quentin Smith. Oxford: Oxford University Press, pp. 25-62..
- Klein, Peter (1981). *Certainty*. Minneapolis: University of Minnesota Press.
- Klein, Peter (1976). 'Knowledge, Causality, and Defeasibility' *Journal of Philosophy* 73: 792-812.
- Kripke, Saul (2011). 'A Puzzle About Belief' in his *Philosophical Troubles*. Oxford: Oxford University Press, pp.125-161.
- Rosen, Gideon (2010). 'Metaphysical Dependence: Grounding and Reduction' in *Modality: Metaphysics, Logic and Epistemology* (ed.) Bob Hale and Aviv Hoffman. Oxford: Oxford University Press, pp. 109-134.
- Scheffler, Israel (1955). 'On Synonymy and Indirect Discourse' *Philosophy of Science* 22: 39-44.
- Strevens, Michael (2008). *Depth*. Cambridge MA: Harvard University Press.
- Ullian, Joseph and Nelson Goodman (1977). 'Truth About Jones' *Journal of Philosophy* 74:317-338.
- Van Fraassen, Bas (1980). *The Scientific Image*. Oxford: Oxford University Press.