For *Well-Founded Belief: New Essays on the Epistemic Basing Relation*,
(eds.) P. Bondy & J.A. Carter, (Routledge).

The Superstitious Lawyer's Inference

Patrick Bondy (Wichita State) and J. Adam Carter (Glasgow)

Abstract: In Lehrer's case of the superstitious lawyer, a lawyer possesses conclusive evidence for his client's innocence, and he appreciates that the evidence is conclusive, but the evidence is causally inert with respect to his belief in his client's innocence. This case has divided epistemologists ever since Lehrer originally proposed it in his argument against causal analyses of knowledge. Some have taken the claim that the lawyer bases his belief on the evidence as a data point for our theories to accommodate, while others have denied that the lawyer has knowledge, or that he bases his belief on the evidence.

In this paper, we move the dialectic forward by way of arguing that the superstitious lawyer genuinely infers his client's innocence from the evidence. To show that the lawyer's inference is genuine, we argue in defense of a version of a doxastic construal of the 'taking' condition on inference. We also provide a pared-down superstitious lawyer-style case, which displays the key features of the original case without including its complicated and distracting features. But interestingly, although we argue that the lawyer's belief is based on his good evidence, and is also plausibly doxastically justified, we do not argue that the lawyer *knows* that his client is innocent.

I. Introduction

The epistemic basing relation, which is the relation that obtains between a belief and the reason(s) *for which* it is held, is of central importance in epistemology. This is because epistemologists care about *epistemic justification*, and the epistemic basing relation is often taken to be necessary for distinguishing between the two principal kinds of epistemic justification: propositional and doxastic. The core idea is roughly as follows: a proposition p is *propositionally justified* for a subject S just in case S possesses good epistemic reasons for believing p; whereas S's belief that p is *doxastically justified* (i.e. S's belief is held in an epistemically permissible fashion) if and only if S believes p *on the basis of* S's good epistemic reasons for believing p, viz., those reasons that propositionally justify p for S.

Many epistemologists are happy to accept that doxastic and propositional justification are 'connected' in this way by the epistemic basing relation.[1] However, it is highly contentious how to spell out the conditions that must be satisfied for a belief to count as being based on a reason. Although it is not perfect match, externalists about epistemic justification have generally aligned themselves with some form of a *causal account* of the basing relation. To a first approximation, this account maintains that for S, p, reason R, S's belief that p is based on R iff S's belief that p is (non-deviantly) caused by R[2], where the matter of how to spell out 'non-deviance' has divided proponents of the

---

[1] See, however, Turri (2010) and Lord and Sylvan (this volume) for notable lines of resistance to this orthodoxy.

[2] Champions of this kind of approach include, e.g., Harman (1970), Swain (1978), Moser (1989), 02/08/2019 00:03:00Turri (2011), and McCain (2012).

causal account.[3] By contrast, *doxastic theories* of the basing relation—accounts that are often found to be attractive by epistemic internalists—take basing to be fixed (in short) principally by beliefs rather than causes; that is, typical doxastic accounts hold that for S, p, reason R, S's belief that p is based on R only if S has a meta-belief to the effect that R is a good reason to believe p (e.g., Audi 1982).

Interestingly, the most famous and controversial counterexample to any version of the causal account was the very first counterexample ever raised to it: Keith Lehrer's (1971) case of the superstitious lawyer, a case initially posed to Gilbert Harman's (1970) early causal account. The matter of whether Lehrer's counterexample worked against causal accounts was a lively topic in 1970s and 1980s,[4] though little consensus was reached—perhaps at least in part because of the sheer complexity of the case. And then, for whatever reason, discussion of the case waned, while the causal account has since gained popularity as the default view.

We take exception to this orthodoxy, which we think is largely unearned. As we see it, Lehrer's counterexample is far worse news for causal accounts than the contemporary thinking about the causal account would suggest. Yet, this is for reasons that have thus far not been fully appreciated, including by proponents (early and present) of doxastic accounts.

The time, we believe, is ripe to revisit Lehrer's famous example with a critical eye, both to the example and to the causal theory it is meant to challenge. In doing so, we

---

[3] See Korcz (2015) for an overview. See also Bondy (2016).

[4] See, for example, Kvanvig (1985, 1987); Lemke (1986); Tolliver (1982); Swain (1981); Goldman (1979).

will show *why* it is—as Lehrer initially thought—a case where one bases a belief on the basis of a reason that is not a cause of the belief. More specifically, as we shall argue, the case is best understood as a case of normal inference from premises to a conclusion, one that is sufficient both for the lawyer to base his belief in the conclusion at least in part on the good evidence he has, and, plausibly, for the lawyer's belief to be doxastically justified by that good evidence. If this diagnosis is correct, then the problem that the case poses for causal accounts can be expressed much more simply, and we think forcefully, than previously thought.

Here is the plan. In §2, we discuss basing, illustrating how causal and doxastic theories handle a wide range of ordinary cases of basing and non-basing correctly. In §3, we discuss inference, in particular arguing in support of a type of doxastic 'taking' condition on inference. In §4, we explain the central features of the original superstitious lawyer case, and we provide a new, much simpler case to illustrate the point of the original one. In both the original case as well as in our new one, the lawyer bases his belief on the evidence, and the evidence doxastically justifies his belief, even though the evidence is causally inert. Interestingly, however, because the lawyer's belief is *unsafe*, it might very well be that critics were right to say that the superstitious lawyer lacks *knowledge*.

II. Basing

Although giving accounts of inference and the basing relation has proven to be a difficult task, identifying clear cases of basing and non-basing is comparatively easy. To orient ourselves in reflecting on the basing relation, then, it's useful to begin with some clear

cases of beliefs held on the basis or not held on the basis of reasons. Identifying features

present in the cases of basing and absent in the cases of non-basing will help us in

thinking about more difficult cases, such as the case of the superstitious lawyer.

Case 1: No Reason, No Basing

Sam believes that it will rain today. Sam has not seen a weather forecast, no one

has told her anything about the weather, and she has not even looked outside. Sam

simply woke up feeling sure that it would rain.

In No Reason, No Basing, there just is no reason on the basis of which Sam holds her

belief.[5] Plausibly, Sam's belief is doxastically unjustified,[6] but the point here isn't about

justification: it's about whether her belief is held on the basis of a reason. There is

nothing that Sam recognizes or would recognize on reflection as her reason for believing

that it will rain. Intuitively, it appears that Sam doesn't hold her belief on the basis of a

reason.[7]

---

[5] When we say "reason" without qualification in what follows, we mean what are often called *motivating* reasons: reasons which *prompt* belief or action, or which the subject *takes* or would take upon reflection to recommend belief or action. Motivating reasons for belief can be good or bad. When they are good, they are normative reasons, and they can doxastically justify belief. When they are bad, the subject might still hold beliefs on the basis of them, but her beliefs will not thereby be doxastically justified.

[6] It is worth noting that in 'no reason' cases, such as this, the subject is sometimes claimed to possess what is called a 'no reason defeater'. For discussions of no-reason defeaters, see Plantinga (1994) and Bergmann (1997).

[7] Perhaps it will appear to some as though Sam's belief is based on something, viz. her hunch. If so, that is no problem; Sam's feeling upon waking up is a natural but dispensable piece of the example. Instead, we

Case 2: Good Reason, No Basing

Sam believes that it will rain today. Sam saw a weather forecast from a reliable source indicating a 95% chance of rain today. But Sam places no trust in weather reports, and the weather report that she saw plays neither an initiating nor a sustaining causal role with respect to her belief. She simply woke up feeling sure that it would rain.

This case is just as in No Reason, No Basing, except that this time there is a fact or feature of the case that could act as Sam's motivating reason for belief, and if she were to hold her belief on the basis of the weather report, her belief could be doxastically justified. But she does not see the weather report as providing good reasons for belief, and it is causally inert with respect to Sam's being in the mental state of believing that it will rain. Her belief, just as in No Reason, No Basing, is based on nothing at all.

Case 3: Bad Basis

Sam believes that it will rain today. She believes that because her weekly horoscope from six days ago indicated that she would see rain this week, and she has not seen rain yet this week.

Unlike in the previous cases, in Bad Basis there is something that Sam regards as a good reason for thinking that it will rain today: her weekly horoscope. Horoscopes do not generally provide good reasons for believing their predictions, so Sam's reason is not a good reason her belief, and all else equal her belief is doxastically unjustified. Still, in this case there is clearly a (motivating) reason on which her belief is based.

---

can imagine that Sam has been hypnotized into believing that it will rain, say, as part of a course of treatment for a debilitating case of heliophobia. Then there is clearly nothing on which her belief is based.

Case 4: Well-Founded

Sam believes that it will rain today. She believes that because she saw a weather forecast from a reliable source, which indicated a 95% probability of rain. Sam trusts weather reports in general, and believes that this report provides a good reason for thinking that there is a high probability of rain today.

As in Bad Basis, here Sam has a reason which she sees as making it the case that she should believe that it will rain today. In this case, the reason is a good one, and all else equal, it doxastically justifies her belief. Nonetheless, there is a clear sense in which the relation between Sam's belief that it will rain today and the good reason that she possesses in Well-Founded is the same kind of relation which obtains between Sam's belief that it will rain today and the bad reason she possesses in Bad Basis.[8] That relation is the basing relation. By contrast, in the first two cases, there is no basing relation instantiated.

The intuitive responses to the four cases above are compatible with a range of views about the nature of the basing relation. For instance, they are compatible with a purely causal account according to which a belief B is based on a reason R iff R non-deviantly causes B.[9] In No Reason, No Basing, and in Good Reason, No Basing, there is no reason which causes Sam's belief, so the causal account of basing yields the result that

---

[8] At least, this is the prevailing view among epistemologists, but see Sylvan and Lord (this volume) for an argument to the contrary.

[9] Non-deviance needs to be spelled out in order for the causal account to be satisfactory, of course, but we can set that aside for now. Our aim in this paper is not to raise causal deviance problems. See Ye (this volume) and Turri (2011) for attempts to handle causal deviance in the account of the basing relation.

Sam does not hold her belief on the basis of a reason in these cases. That is the intuitively correct result.

In Bad Basis and Well-Founded, by contrast, there is a reason which appears to non-deviantly cause Sam's belief, and so the casual account yields the result that Sam's belief is held on the basis of these reasons. Again, that is the intuitively correct result.

But these results do not by themselves settle the issue in favour of causal accounts of the basing relation. For example, we can correctly categorize these cases if we opt for a doxastic account of basing according to which a belief B is based on a reason R iff the subject S has an appropriate meta-belief, to the effect that an available reason is a *good* reason for belief. This sort of doxastic account also correctly categorizes the cases we've seen so far: in the first two cases, Sam has no such meta-belief, so the doxastic account entails that her belief is not based on a reason, while in the latter two cases, plausibly, she does have such a belief.

So, in a wide range of clear cases where subjects hold beliefs on the basis of reasons, the reasons are causally efficacious in sustaining the subject's belief, *and* the subject believes that the reason is a good one for holding the target belief. And in a range of clear cases where beliefs are not held on the basis of reasons, there is no reason which the subject takes to be a good reason for the belief, *and* there is no reason that is non-deviantly effective in causing or sustaining the belief.

That there is a wide range of cases of basing in which (i) subjects take there to be good reasons for their beliefs and (ii) those reasons causally support the target beliefs should be unsurprising. After all, we are creatures capable of rationally regulating our beliefs in accord with the reasons that appear to us to be good ones, and when things are

going well, we do causally regulate our beliefs in that way. We should therefore expect that in normal cases, where our cognitive faculties are functioning normally, our recognition of a reason as a good reason for belief will prompt us to hold the target belief. And in such cases, when a reason causes us to hold a target belief, it does so only because we take the reason to recommend holding the target belief.

So it is unsurprising that casual and doxastic accounts yield the same verdicts in in ordinary cases. Each type of account identifies a relation that clearly appears to obtain in ordinary cases where beliefs are held on the basis of reasons, and that fails to obtain in clear cases of non-basing. To help us decide whether the basing relation is really a causal relation or a doxastic relation (or both[10], or neither[11]), it will help to consider cases where causal and non-causal accounts of the basing relation diverge, such as superstitious-lawyer types of cases. In these cases, the intuitive response regarding whether the belief is held on the basis of a particular reason is not so clear—or at least, intuitions have historically been divided.

In the superstitious-lawyer case, a lawyer (in short) infers his client's innocence from some good evidence; the evidence is causally inert with respect to the lawyer's belief in its conclusion; but he still draws the inference.[12] Lehrer used the example to argue that the lawyer can have knowledge of the conclusion of the inference, crucially, *on the basis of the causally inert inference*. If Lehrer is right about that, then—*contra* the

---

[10] See Korcz (2000) for an example of a hybrid account.

[11] For instance, according to Lord and Sylvan (this volume), the basing relation is prime.

[12] See below, section IV, for a full description and discussion of the case.

causal account— beliefs can be held on the basis of reasons which have no causal efficacy.

In Section, 4 we will argue that because the lawyer draws the relevant inference, we should see him as holding the belief in his client's innocence on the basis of his appraisal of the evidence. In other words, we argue, inference from believed premises to a believed conclusion is sufficient for basing belief in the conclusion on belief in the premises. That is just what we see in the case of the superstitious lawyer.

Before turning to the full case of the superstitious lawyer, it will therefore be profitable first to clarify what we mean by "inference," and why we think inference is sufficient for basing.

III. Inference

The term "inference" has several related uses. Sometimes we talk about inferences as conscious mental events; sometimes we talk about implicit inferences that we can be said to draw, when we unconsciously or subconsciously base beliefs on perceptual input. Sometimes we talk about inferences as actual events or processes that have taken place in a person's mind, while other times we are interested in potential inferences that a person might draw; the former are concrete events, while the latter are abstract objects, consisting of propositional contents and inference-warrants which permit the acceptance of some propositional contents on the basis of some other accepted propositions. The relevant sense of "inference" here is that of a concrete mental act, involving the

consideration of some premises and a conclusion, where the subject takes the premises to support the conclusion.[13]

The key question in articulating an account of inference is how that "taking" condition should be understood: what does it mean to say that a subject takes a set of premises to support a conclusion? For example, does it mean that the subject *believes* that the premises support the conclusion? Or perhaps it means that the subject is disposed to *follow a rule* according to which one ought to form a belief in the conclusion, given that one accepts the premises?

One way to approach the taking condition is to consider cases where a subject S believes that the premises do *not* support the conclusion. In such cases, clearly, S does not take the premises to support the conclusion. Similarly, we can consider cases where a subject has no opinion on whether some premises support a conclusion. Again, cases like this are likely to strike us as clearly cases where the subject does not take the premises to support the conclusion. And so the most natural interpretation of the taking condition is in doxastic terms:

> **Taking (Doxastic):** S takes a premise P to support a conclusion C iff S believes that P supports C.

---

[13] Some philosophers also distinguish theoretical inference from practical inference, in virtue of the type of attitude featuring in the conclusion: theoretical inferences conclude with beliefs, and practical inferences conclude with intentions. For the purpose of this chapter, we are only addressing theoretical inferences, and we do not take a stand on this distinction. Even if practical inferences are distinguished from theoretical inferences in that way, the doxastic construal of the taking condition that we defend here will apply to practical inferences in the very same way as it applies to theoretical inferences.

Because this is the most natural interpretation of the taking condition, it is the one philosophers tend to consider first, before moving on to other interpretations when problems come up for the doxastic construal.[14] In our view, the objections that have been raised against the doxastic construal are far from decisive. They do show that Taking (Doxastic) is too crude as it stands, but what the objections point to is the need for a more refined understanding of the doxastic construal. And, with the refined doxastic construal in hand, we will be in position to see that superstitious-lawyer-style cases are genuine cases of inferential basing of beliefs on reasons.

The two most important objections to Taking (Doxastic), in our view, are the over-intellectualization objection and the causal deviance objection. We will address these in turn.

III (i) *Over-intellectualization*

Paul Boghossian (2014) writes:

> A child, we are inclined to think, can reason. Luke and Drew are playing hide-and-seek. Seeing Drew's bicycle leaning against the tree, Luke thinks: ''If he were hiding behind that tree, he would not have left his bicycle there. So, he must be behind the hedge.'' That looks like reasoning. But do children have meta-beliefs about the relations between their premise judgments and their conclusions? Do children have the concepts of premises and conclusions? Do they have the normative concept of one belief justifying another? (2014, pp.6-7)

---

[14] E.g., Boghossian (2014); Neta (2013).

The worry is that (surely) children engage in reasoning. But they are also (surely) too unsophisticated to have the concept of a premise, or a reason. So they cannot have the meta-belief required to satisfy Taking (Doxastic).

Examples like this one, we think, give children too little credit. From a very young age, children are able to ask and answer "why"-questions; and to ask why one ought to hold a belief, or why (for example) one must go to school, is to ask for normative reasons. Young children—and probably most adults—no doubt have never heard the phrase "normative reason," or even "premise" and "conclusion," but they do know what it is to hold beliefs or perform actions for good reasons, or for bad reasons. That is all the conceptual apparatus required to satisfy Taking (Doxastic). And, in Boghossian's example, it looks as though Luke *does* have the required conceptual sophistication to form the necessary meta-belief. It is a dispositional belief, to be sure; Luke likely does not consciously dwell on it. But it is a belief of his nonetheless. So that is not a counterexample to Taking (Doxastic).

On the other hand, if we turn our attention to children who are so young and unsophisticated that they do not even know what it is to hold beliefs for good or bad reasons, then it is not clear that such children are really capable of drawing inferences, in the reflective, consciously aware sense of considering premises and drawing conclusions on the basis thereof.[15]

---

[15] Of course, even these very young children must make *some kinds* of mental moves, transitions in thought that fall short of basing. A natural way to think about such transitions in thought is in terms of *cognitive heuristics*, mental shortcuts which do not correspond with any kind of rational regulation, but which

III (ii) *Causal deviance*

Ram Neta (2013)[16] argues that if the Taking Condition is construed as the condition that the subject *represents* (doxastically or non-doxasically) that the premises support the conclusion, then there will be counterexamples where a subject S satisfies the relevant taking condition, and S's belief in the premises causes S to believe the "conclusion," but where the belief in the conclusion is caused in a deviant manner by the belief in the premises and the meta-belief. In such a case, S does not draw the inference.[17] Causation

---

nonetheless can be practically useful. For a recent discussion of cognitive heuristics – described in the language of Type-1 thinking – see Kahneman (2011).

[16] Neta is following Boghossian (2014) here, but Neta makes this point much more explicitly than Boghossian does.

[17] Another kind of objection in this neighbourhood is highlighted by Wallbridge (2018), who writes that doxastic accounts fail:

> to capture the sense in which the basis of a belief is what *explains why* one believes as one does. To the extent that the kinds of meta-beliefs proposed by doxastic accounts are able to provide an explanation, they do so in virtue of their covariance with causal or counterfactual factors (2018, 1619).

This strategy of objection, we think, problematically presupposes that the only kind of explanation that would be captured by the pretheoretical intuition that bases *explain* why one believes what they do is causal. If it were, then a causal theory would be the only kind suited to satisfying the pretheoretical intuition. However, the doxastic account proponent can reasonably resist Wallbridge's preferred way of reading the pretheoretical intuition while maintaining that the doxastic account satisfies a more ecumenical version of it.

of belief in a conclusion by belief in premises, in conjunction with possessing a meta-belief that would satisfy the Taking Condition, is thus insufficient for drawing an inference. For example:

> Roderick has done nothing of value throughout his life except for offering a solution to the Gettier problem. Now, on his deathbed, Roderick thinks back on his otherwise worthless life and says "well, at least I solved the Gettier problem". But as he says this, Timothy bursts into his room and proves to Roderick that his solution to the Gettier problem is unsuccessful. Roderick now believes that his solution to the Gettier problem is unsuccessful, and also believes that this very fact supports the conclusion that his life was worthless. Believing these two things, Roderick falls into a state of despair, and, out of despair, believes that his life was worthless. Of course, that his life was worthless is a conclusion that Roderick could have inferred from two beliefs that he has, viz., that his solution to the Gettier prolem fails, and that the failure of his solution to the Gettier problem supports the conclusion that his life was worthless. But Roderick does not make this inference, and it is not on the strength of any such inference that Roderick believes that his life was worthless. Rather, it is out of despair that he believes that his life was worthless. (Neta 2013, p.390)

The case of Roderick appears to be a counterexample to Taking (Doxastic) because Roderick has two first-order beliefs (P: he has failed to solve the Gettier problem; C: his life was worthless), and he has a meta-belief about those first-order beliefs (that belief in P supports belief in C), but he does not draw the inference from P to C. This counterexample shows that in formulating a doxastic construal of the taking condition,

we need to be explicit that the subject does not simply have a mental representation of separate propositions P and C, and a meta-belief that belief in P would support belief in C. S must represent *her own* belief in P as a good reason for *her own* belief in C.

> **Taking (Reflexive):** S takes P to support C iff S believes that *her own* belief in P supports (/is a good reason for) *her own* belief in C.[18]

Taking (Reflexive) explains what is going wrong in the case of Roderick: Roderick has two beliefs, P and C, and he has a belief that P supports (the content of) C, but he has not drawn the connection between them. He does not believe that his very belief in P supports his very belief in C. That is why he does not inferentially base his belief in C on his belief in P. But whenever S believes P, and believes C, and believes that her very belief in P is a good reason for holding her very belief in C, S inferentially bases her belief in P on her belief in C.

One potential worry about Taking (Reflexive) is that it brings the overintellectualization problem back into focus. Above, we considered the problem that the Taking (Doxastic) rules out young children and higher animals from drawing inferences because they appear to lack the concept of a normative reason. We responded that they do after all have a rudimentary concept of a normative reason. And now the

---

[18] We treat Taking (Reflexive) as a condition requiring full belief here, but note that it can be weakened by giving it a "de re awareness" reading, according to which S need not form the full-fledged belief that her belief in P supports her belief in C; S only needs to be aware of her own belief in P as supporting her own belief in C. (See Korcz (2015, section 3), and Moser (1989) for more on the de re awareness condition.) What we are most concerned to emphasize in Taking (Relfexive) is not what kind of belief-like representation or awareness the subject has, but simply that the representation is reflexive.

overintellectualization worry comes back as a problem for Taking (Reflexive) because young children and higher animals might have beliefs that are based on reasons (e.g., based on visual or olfactory perception), but they might also lack the concept of belief, in which case they cannot have beliefs about their own beliefs. If that is right, then young children and higher animals cannot satisfy Taking (Reflexive), and so according to our account, they cannot draw inferences after all.

We can respond to the overintellectualization worry for Taking (Reflexive) as we did in the context of Taking (Doxastic). For one thing, just as young children and higher animals plausibly have a rudeimentary concept of good reasons, so too they plausibly have at least a rudimentary concept of belief. For example, when a gazelle hides from a lioness, the gazelle is trying to ensure that the lioness doesn't notice that the gazelle is there. Plausibly, the gazelle represents its own perspective on the world as distinct from the lioness's perspective. It tries to ensure that the lioness's perspective does not include a correct representation of the gazelle's location, and so the gazelle stays downwind so that the lioness doesn't discover the gazelle by smelling it. If that is right, then the gazelle does have at least a rudimentary conception of the lioness's beliefs.

Further, as we note above in section III (i), if we turn our attention to children so young or animals so intellectually undeveloped that they don't have even a rudimentary conception of belief, then it just is not clear that such children and animals are capable of drawing inferences, and we should not want our account of inference to include them.[19]

---

[19] Of course such children and animals would make some mental transitions, but those transitions need not count as inferences. See above, note 13.

Might one, at this stage, attempt to challenge Taking (Reflexive) by appealing to a 'Lewis Carroll'-style regress argument? In his famous regress argument, Carroll (1895) envisaged a student who was told P, and told that P entails Q, and told that if P & (P entails Q), then Q. But yet, the student (we may imagine) still does not *know how* to draw the inference; the student, despite this propositional knowledge, failed to *see that* her belief that Q follows from the premises the student believes. And this would remain the case for any further iteration of a belief *about* the premise and conclusion (or their relationship to one another) that the teacher could possibly tell the student. Carroll's argument—borrowed to similar effect by Ryle (1945)—was meant to cast doubt upon the idea that the kind of 'know-how' involved in skillfully drawing an inference is something that will ever be secured by simply possessing propositional knowledge (and *ipso facto*, by simply possessing beliefs) about how the premise and conclusion stand in relation to each other.

By parity of reasoning, we might imagine Carroll challenging Taking (Reflexive) along broadly analogous lines: S might (as the argument would go) believe that *her own* belief in P supports (/is a good reason for) *her own* belief in C, but (like Carroll's student – *failing to see the connection* between P and C) nonetheless fail to *take* P to support C.

The problem with this objection is that the taking relation that is relevant to characterizing inference is *weaker* than the kind of relation Carroll purported to show could not be secured by the possession of propositional attitudes. *That* relation was a success-oriented relation: one that involved a kind of factive *seeing*. We are happy to grant Carroll (and Ryle) that a thinker's believing that *her own* belief in P supports (/is a good reason for) *her own* belief in C does not imply that that thinker *sees how* P supports

C. But this is perfectly compatible with countenancing that that a thinker's believing that *her own* belief in P supports (/is a good reason for) *her own* belief C suffices for that thinker to *take* her P belief to support her C belief.

The above reply can be helpfully appreciated against the background of a more general point: *not all inferences must be either good inferences or for that matter skillful inferences*. Novice (and intermediate, and even expert) logicians sometimes mistakenly draw inferences in accord with faulty inference-rules. And what goes for inference goes for basing: sometimes, a subject S will hold a belief in C on the basis of a belief in a good reason P, but S misunderstands the support-relation at hand. In such cases, S's holding a belief on the basis of the available good reason is insufficient for doxastic justification.[20] What is required for doxastic justification is not just *basing* of one's belief on good reasons, but *correct*, or perhaps *skillful*, basing of one's belief on good reasons.

IV. The Superstitious Lawyer, Reconsidered

---

[20] Turri offers a number of examples to illustrate the problem of beliefs held badly on the basis of good reasons. For example: in a case where P1-P4 make it very likely that Mansour is Guilty, two members of the jury form the belief that he is guilty.

> Miss Proper reasons like so:
>> (Proper Reasoning) (P1–P4) make it overwhelmingly likely that Mansour is guilty. (P1–P4) are true. Therefore, Mansour isguilty.
>
> Miss Improper, by contrast, reasons like this:
>> (Improper Reasoning) The tea leaves say that (P1–P4) make itoverwhelmingly likely that Mansour is guilty. (P1–P4) are true. Therefore, Mansour is guilty. (2010, p. 316)

With the foregoing considerations in hand, let's now look again at the Superstitious

Lawyer case, which is summarised pithily by Kvanvig (1985) as follows:

> The counterexample concerns a [sic. superstitious] lawyer who, like the rest of his
>
> contemporaries, takes his client to be guilty. However, because of his [sic.
>
> superstitious] nature, the lawyer is inclined to trust what the tarot cards say, and upon
>
> learning that the tarot cards say that his client is innocent, comes to believe that his
>
> client is innocent. What the tarot cards say also prompts the lawyer to re-examine the
>
> evidence, which the lawyer comes to recognize conclusively establishes that his client
>
> is innocent. However, given his rather impressionable character, the lawyer also
>
> realizes that were the sustaining power of the tarot cards removed, the sway of public
>
> opinion would cause him to be unable to see that the evidence establishes his client's
>
> innocence. Nonetheless, the lawyer now justifiably believes that his client is innocent
>
> on the basis of his examination of the evidence. But this examination of the evidence
>
> neither prompts his belief that his client is innocent nor does it sustain his belief that
>
> his client is innocent—his belief in what the tarot cards say holds the dubious
>
> distinction of being responsible for both. (1985, 153-4)

Despite the complexity of the case, a simple point can be put on the table. Once the

lawyer's superstition causes him to re-examine the evidence, it's clear that the lawyer

believes that his own belief in the evidence (call this 'P') supports (/is a good reason for)

his own belief in the conclusion that the client is innocent (call this 'C'). According to

Taking (Reflexive), it follows that the lawyer thereby *takes* P to support C.

This point, in conjunction with our argument in §III, makes for a simple diagnosis

of the superstitious lawyer case. (i) An inference from believed premises to a believed

conclusion is sufficient for basing belief in the conclusion on belief in the premises; and (ii) the superstitious lawyer makes an inference from believed premise(s) to a believed conclusion. And so the case is straightforwardly a case of inferential basing. Finally, because the evidence is a *good* reason to believe the conclusion—and, moreover, we have no good reason to think the inference itself is not a skilful or competent one[21]—the superstitious lawyer's belief is plausibly doxastically justified. And this is the case *despite* the evidence not *causing* the belief.

---

[21] One might object here that the inference is not skillful because, were the sustaining power of the tarot cards removed, the sway of public opinion would have prevented the inference from being drawn successfully. This reasoning, though, is problematic; it conflicts with ordinary thinking about skill in epistemology. To borrow an example from Sosa (2015), suppose a skilled basketball player makes a free throw, and that the free throw manifests the shooter's skill. Now, add to the story that the electricity to the gymnasium could have suddenly failed, and had it failed, the shot would have missed due to darkness. But it did not suddenly fail. Does the fact that the electricity could have failed—even if it could *easily* have failed—count against the shooter's possessing and manifesting free-throw shooting skill in making the shot while the lights are on? Very plausibly, the answer here is 'no'; and a similar point we think goes for the superstitious lawyer. True, he would be unreliable *if* the initial conditions were altered. But, situated as he is, he is, we may presume, in an analogous position to that of a skilled basketball shooter whose reliability requires that the lights stay on and would be thwarted if they weren't. Put another way, the lawyer is in a very different position from a hopeless counterpart lawyer who, (like the original lawyer) is also a follower of tarot who is easily swayed by public opinion, but who is, in addition, *inferentially challenged*. The lawyer, but not the counterpart, is situated so as to infer reliably.

The sheer complexity of the case, we think, is at least partly responsible for confusions shared on both sides of the debate.[22] These complexities, we believe, are ultimately unnecessary. The thrust of the example can be made much more simply. Consider the following 'pared down' version of the case, one that does what the original case *aims* to do, but without unnecessary distractions:

Brain Scientists

A lawyer is working for a cabal of brain scientists. She has good evidence for believing her client is innocent, and through careful and skilled reasoning recognizes it as good evidence, but the scientists have rendered evidential considerations with respect to her client's guilt or innocence causally inert for her. The lawyer would make the evidence causally effective, if she could, but she cannot. Nevertheless, she fully believes that the evidence supports her belief in her client's innocence.

Intuitively, the lawyer in Brain Scientist has drawn an inference to the conclusion that her client is innocent. Further, because the lawyer has drawn this inference and done so skilfully, her belief in her client's innocence appears to be doxastically justified. At least, this is clearly not a textbook case of propositional justification without doxastic

---

[22] Consider, for example that the following counterfactual is true of the lawyer in the original case: if they did not base the belief on the good reason they have, then they would have (rather than withheld judgment) based the belief on a *bad* reason they have. We can of course grant this conditional while stressing that the lawyer nonetheless *does* in fact base their belief on a good reason.

Korcz (2000) presents other useful variations on the superstitious lawyer case, but they are at least as complex as the original. Nowhere in the literature do we find straightforward examples that illustrate the relatively simple idea that Lehrer was driving at with his complicated case.

justification. In standard cases of that sort, a subject possesses good evidence but disregards it, or fails to properly appreciate it. By contrast, the lawyer in Brain Scientist properly appreciates the evidence she possesses; the evidence is just incapable of causally supporting her belief.

This case controls for the kind of unnecessarily loaded aspects of the explanation for *why* the evidential considerations are causally inert, which we find in Lehrer's original case. By simply appreciating (i) that the evidential considerations are causally inert; along with the fact that (ii) normal inference takes place, the force of the strategy the superstitious lawyer case initially used can be better appreciated as seriously problematic for proponents of the causal theory of the basing relation. By focusing our attention to the simple structure of the strategy, the burden on the causal theorist sharpens: the causal theorist can dodge this *kind* of case only if they can give a good explanation for why *in the presence of normal inference* the causal inefficacy of a reason should *matter* for basing.

Thinking about the case in this way will accordingly make certain *kinds* of responses less relevant. This includes Goldman's response (on behalf of the proponent of the causal theory) that the lawyer (contrary to what Lehrer originally suggests) lacks *knowledge*. Our reading of the landscape commits us to viewing the situation of the superstitious lawyer as analogous to the situation in the Well Founded case we discuss in section II, in that it is a case of inferential basing that leads to doxastic justification.

Importantly, though, inferential basing that leads to doxastic justification can be

unsafe in a way that leads the target belief to fall short of knowledge.[23] (Compare: subjects in barn façade cases have paradigmatically well-founded beliefs, and yet their beliefs are unsafe because very easily they could have been false given how they are formed.) And, of course, the superstitious lawyer could easily have drawn a different tarot card than the one he did.

What this means is that the structure of the superstitious lawyer case—which is highlighted in our revised version—is one that suffices to make the case against the causal theory *regardless* of whether the belief one ends up with counts as knowledge— and thus, regardless of whether one's evidence lacking causal efficacy is enough to undermine knowledge of the target proposition.

V. Conclusion

Superstitious lawyer-style cases are surely unusual—indeed, for all we have said here, they might even be necessarily exceptional. For example, perhaps the capacity to draw inferences depends on a subject's having the capacity to causally regulate her beliefs in accord with her appraisal of the available evidence in normal cases. We do not argue that that is so, but neither have we said anything to rule out that sort of possibility. And it also bears emphasizing that we have only been concerned here with beliefs involved in conscious inferential processes. For all we have said here, non-deviant causation of beliefs by reasons might suffice for basing in cases of non-inferential belief formation.

---

[23] On one canonical way of thinking about safety and knowledge due to Pritchard (2005), a belief is known only if it is safe in the sense that it couldn't easily have been false given the conditions of its formation.

Although we have not argued that the superstitious lawyer has knowledge of his client's innocence, we find it intuitively very plausible that he bases his belief on the available evidence, and that—given that he does this normally and plausibly skilfully—his belief is consequently doxastically justified. But this verdict about the case does not simply rest on intuition. For if a refined doxastic construal of the taking condition on inference of the kind we have proposed in section III is correct, then in superstitious lawyer-style cases, the subject genuinely draws an inference from the available evidence to the target belief, and consequently holds the target belief on the basis of the evidence. Because the lawyer's evidence is good evidence for the target belief, and he competently appreciates its force, the lawyer's belief is doxastically justified, even though the evidence for it is causally inert.[24]

References

Audi, Robert. 1982. "Believing and Affirming." *Mind* 91 (361): 115–20.

Bergmann, M. A. (1997). *Internalism, Externalism and Epistemic Defeat.* (Doctoral dissertation).

Boghossian, Paul. 2014. "What is Inference?" *Philosophical Studies* 169: 1-18.

Bondy, Patrick. 2016. "Counterfactuals and Epistemic Basing Relations." *Pacific Philosophical Quarterly* 97 (4): 542–69.

Carroll, Lewis. 1995. "What the Tortoise Said to Achilles." *Mind*, 104 (416): 691-693.

Goldman, Alvin I. 1979. "What Is Justified Belief?" In *Justification and Knowledge*, 1–23. Springer.

---

[24]

Harman, Gilbert H. 1970. "Knowledge, Reasons, and Causes." *Journal of Philosophy* 67 (21): 841–55.

Kahneman, Daniel. 2011. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.

Korcz, Keith Allen. 2000. "The Causal-Doxastic Theory of the Basing Relation." *Canadian Journal of Philosophy*, 30 (4): 525-550.

———. 2015. "The Epistemic Basing Relation," in *Stanford Encyclopedia of Philosophy* (Fall ed.), ed. Edward N. Zalta.

https://plato.stanford.edu/archives/fall2015/entries/basing-epistemic/.

Kvanvig, Jonathan L. 1985. "Swain on the Basing Relation." *Analysis* 45 (3): 153.

———. 1987. "On Lemke's Defence of a Causal Basing Requirement." *Analysis* 47 (3): 162–67.

Lehrer, Keith. 1971. "How Reasons Give Us Knowledge, or the Case of the Gypsy Lawyer." *Journal of Philosophy* 68 (10): 311–13.

Lemke, Lory. 1986. "Kvanvig and Swain on the Basing Relation." *Analysis* 46 (3): 138–44.

Moser, Paul. 1989. *Knowledge and Evidence* (Cambridge, UK: Cambridge University Press).

Neta, Ram. 2013. "What Is an Inference?" *Philosophical Issues* 23: 388-407.

Plantinga, Alvin. 1994. "Naturalism defeated." *Unpublished manuscript*.

Pritchard, Duncan. 2005. *Epistemic Luck.* Oxford: Oxford University Press.

Ryle, Gilbert. 1945. "Knowing How and Knowing That: The Presidential Address." *Proceedings of the Aristotelian society* (46): 1-16.

Sosa, Ernest. 2015. *Judgment and Agency.* Oxford: Oxford University Press.

Swain, Marshall. 1978. "Reasons, Causes, and Knowledge." *Journal of Philosophy* 75
(5): 229–249.

———. 1981. *Reasons and Knowledge*. Vol. 92. Cornell University Press.

Tolliver, Joseph. 1982. "Basing Beliefs on Reasons." *Grazer Philosophische Studien* 15:
149–61.

Turri, John. 2010. "On the Relationship between Propositional and Doxastic
Justification." *Philosophy and Phenomenological Research* 80 (2): 312–26.

Wallbridge, Kevin. 2018. "The Peculiar Case of Lehrer's Lawyer." *Synthese* 195 (4):
1615-1630.