

How to Solve the Puzzle of Peer Disagreement

Michele Palmira

University of Barcelona and LOGOS

American Philosophical Quarterly 2019

Special issue “The nature and implications of disagreement”

ABSTRACT

While it seems hard to deny the epistemic significance of a disagreement with our acknowledged epistemic peers, there are certain disagreements, such as philosophical disagreements, which appear to be permissibly sustainable. These two claims, each independently plausible, are jointly puzzling. This paper argues for a solution to this puzzle. The main tenets of the solution are two. First, the peers ought to engage into a deliberative activity of discovering more about their epistemic position vis-à-vis the issue at stake. Secondly, the peers are permitted to do so while entertaining a *sui generis* doxastic attitude of hypothesis.

1 The Puzzle of Peer Disagreement

What ought we do when we discover a disagreement with an acknowledged epistemic peer, that is, somebody who is as well positioned epistemically as we are with respect to the matter at hand?¹ Call this *The Question*.

In this paper I set out to present and solve a specific puzzle which emerges from *The Question*. The puzzle is fuelled by two apparently compelling yet opposite intuitions. The first intuition is elicited by cases such as the following, due to David Christensen (2007):

(RESTAURANT)

Allison is dining with her friend Marc. They’ve been going out to dinner together regularly over the past several years. Most of the time they’ve calculated their respective shares of the bill, they’ve been right equally often. Tonight, after having looked at the bill, Allison asserts with confidence that she has carefully calculated in her head that they each owe \$43, while Marc says with the same degree of confidence that he has calculated in his head that they each owe \$45.

¹ See e.g. Kelly (2005) and Elga (2007) for different definitions of epistemic parity. In this paper I can afford to remain neutral on what the correct definition is.

It seems plausible to say here that Allison and Marc ought not retain their original doxastic attitudes. After all, they cannot simply ignore the fact that somebody who they acknowledge to be as well positioned epistemically as they are has reached an opposite verdict about the bill. Nor can they disregard their opposite number's opinion simply on the basis of the fact that he or she disagrees with them, for that would be exceedingly dogmatic. Call this the *Conformist Intuition*.

Surely there is something attractive about the Conformist Intuition. It's no accident that much of the debate about peer disagreement has revolved around various ways of developing answers to *The Question* which are motivated by the Conformist Intuition. These are called *conciliatory* or *equal weight* views of peer disagreement (see e.g. Christensen 2007, 2011, Elga 2007, Feldman 2006).

However, other peer disagreement cases elicit an opposite intuition. Take the following:

(A PRIORI)

Rudy and Ruth are two philosophers interested in whether there's a meaningful distinction between *a priori* and *a posteriori* justification. They regard themselves as equally thoughtful, sophisticated and careful philosophers accessing the same body of evidence on this question. Rudy believes in the existence of the distinction, whereas Ruth does not.

According to what I shall call the *Non-Conformist Intuition*, Rudy and Ruth are permitted not to revise their respective opinions. The Non-Conformist Intuition can be supported by various considerations. For one thing, one form of philosophical progress is the development of full-fledged competing theories of a given subject matter. Constant revision of our ideas in the face of disagreement with a philosophical peer would put a leash on this form of progress. For another, philosophical discourse is much more tied to understanding than mental calculation is. Plausibly, understanding requires that we figure things out by ourselves by responding directly to the evidence and the philosophically relevant features of a given question. Therefore it seems legitimate not to assign too much extra epistemic weight – that is, the kind of weight which could mandate revision of our philosophical opinions – to a view just because one of our peers defends it.

Despite their initial individual plausibility, however, a joint theoretical articulation of

the Conformist Intuition and the Non-Conformist Intuition seems to give rise to a puzzle. To illustrate.

On a somewhat natural way of expanding on the Conformist Intuition, the fact that one is in disagreement with an acknowledged peer constitutes evidence that somebody has made a mistake; and one has no reason to think that it's one's peer who has made it rather than oneself. This rational doubt seems to require that one change one's view about the matter at issue. And yet, the Non-Conformist Intuition apparently calls for a response at the opposite end of the spectrum, to the effect that peers are permitted not to change their views.

In light of this, if we subscribe to both the Conformist Intuition and the Non-Conformist Intuition, it seems that we are also committed to the following almost paradoxical claims:²

- (a) Whenever you discover a disagreement with an acknowledged peer, you thereby gain evidence that you have made a mistake.
- (b) Whenever you have evidence that you have made a mistake, you ought to change your view about the matter at hand.

Therefore:

- (c) Whenever you discover a disagreement with an acknowledged peer, you ought to change your view about the matter at hand.

And yet:

- (d) Sometimes, when you discover a disagreement with an acknowledged peer, you are permitted not to change your view about the matter at hand.

Call this *The Puzzle*. I submit that any satisfactory answer to *The Question* has to be able to solve *The Puzzle*.

In this paper I offer an answer to *The Question* which provides a solution to *The Puzzle*. I will proceed as follows. In section 2 I articulate my answer to *The Question*. In section 3 I show how it solves *The Puzzle*.

2 The Answer to the Question: Re-Opening and Hypothesising

My answer to *The Question*, which I shall dub *The Answer*, consists of the following claims.

²See also Everett (2015: 275).

First, after the discovery of their disagreement, the peers ought to *re-open the question whether p* .³ Second, while under the requirement to re-open the question, the peers are permitted to revise their attitudes in only two ways: by *suspending judgement* or by *hypothesising*. Third, the peers are *not permitted to retain* their previous doxastic attitudes. Let me articulate these claims in turn.

2.1 Re-Opening the Question

To re-open the question whether a given proposition p is true is to perform certain tasks, such as going over the shared body of evidence by re-evaluating its extension (i.e. whether a given item counts as a piece of evidence or not), carefully re-assessing its probative force, double-checking the reasoning by means of which one put all the evidential items together to get to a conclusion about its support for a specific attitude, and making sure that one's general epistemic and cognitive conditions were normal. Alternatively, re-opening the question whether p may require looking for new evidence and arguments in order to come to a verdict about p 's truth-value.

Re-opening the question of the truth of p is, I contend, a cognitive and deliberative activity whose aim is to find out more about p 's truth-value by re-assessing the whole epistemic situation one is in and, if necessary, improving one's epistemic situation by looking for new evidence, information, or methods of investigation. This reassessment of one's evidence, reasoning, and epistemic credentials puts one in a position to unmask misleading ways of weighing certain evidential items, the insufficiency of the present body of evidence, and/or the need to look for new evidential items, a flaw in the reasoning from evidence to doxastic attitude, a deficient general epistemic status (e.g. the presence of some bias which impaired one's judgment), and so on and so forth.

In light of this, it is plausible to take re-opening the question whether p to be a *truth-promoting* activity. Suppose that during the re-opening activity one discovers that one has mistakenly assessed the probative force of some evidential item by taking it to speak in favour of p 's falsity rather than its truth. Or suppose that one discovers that one has mistakenly taken an item i to be a piece of evidence concerning the matter at hand when it is not. Plausibly, uncovering such mistakes is conducive to improving one's epistemic position vis-à-vis getting p 's truth-value right.

As far as I can see, *The Question* asks us to propose a normative epistemic requirement

³ Some of the terminology I will employ is borrowed from Shah and Velleman (2005).

which tells peers what to do when they discover their disagreement. Since the *doing* I am recommending here – i.e. re-opening the question whether p – is epistemic in character in virtue of its truth-promoting nature, the claim that peers ought to respond to their disagreement by re-opening the question in this way seems to engage with *The Question* at the right epistemic level.

I take this to be a fairly natural way of responding to peer disagreement. Yet natural though it is, it has not previously been defended. Indeed, closer inspection reveals that the requirement that peers re-open the question is inconsistent both with the requirement that they revise and with the permission that they retain their initial doxastic attitudes, *once a narrow yet very common interpretation of doxastic revision is adopted*. If this is right, then the claim that peers ought to respond to their disagreement by re-opening the question of the truth of the contested proposition provides a substantively new way of answering *The Question*.

The widespread yet – I believe – unduly narrow conception of doxastic revision which I will adopt for purposes of this section is the following. One *revises* one’s attitude toward p , in the narrow sense that I intend, if and only if one believes (or—*mutatis mutandis*—disbelieves, or assigns a certain credence to) p at t_0 and one disbelieves (or—*mutatis mutandis*—believes, or assigns a higher or lower credence to) p at t_1 . As for the corresponding narrow notion of doxastic *retention*, let us say that one *retains* one’s attitude toward p if and only if one believes (or disbelieves, or assigns a certain credence to) p at t_0 and one believes (or disbelieves, or assigns the same credence to) p at t_1 . I will label these “*revise_n*” and “*retain_n*” respectively, clarifying later why I take them to represent a particularly “narrow” conception of doxastic revision and retention.

To show that the requirement to re-open the question whether p is inconsistent both with the permission to retain one’s original doxastic attitude and with the requirement to revise it, provided that revision and retention are understood as per the above, I will introduce the notion of *taking the question whether p to be settled*. I will then propose four principles exposing the normative connections among re-opening the question whether p , retaining_n and revising_n one’s doxastic attitude toward p , and taking the question whether p to be settled.

The notion of ‘taking the question whether p to be settled’ does not imply that one knows the answer to the target question, nor is it factive (in the sense of entailing that the question *is* settled). Rather, taking the question whether p to be settled is always relative to one’s own epistemic position: one takes the question whether p to be settled in a given way or in another – e.g. one takes the question whether p to be settled in the affirmative rather

than in the negative – given the evidence, epistemic facts and reasoning powers available to one. Thus, if one’s body of evidence at t is such that – defeasibly and perhaps misleadingly – it makes it the case that there’s a 70% chance that Alice will be eating cheese for lunch, it seems that one can – and is permitted to – close inquiry into the question of what Alice will eat for lunch by being .7 confident that she’ll be eating cheese. Should one acquire new countervailing evidence, one would (if rational) re-open the question and assess the probative force of the new evidence.

Let me first set out the normative connections among re-opening, retaining_n and taking the question to be settled.⁴ Here they are (understanding all *oughts* and permissions as epistemic):

PR1: If one ought to re-open the question whether p , then one ought not take the question whether p to be settled.

PR2: If one is permitted to retain_n one’s initial attitude towards p , then one is permitted to take the question whether p to be settled.

I take both principles to be extremely plausible. The justification for **PR1** is the following. Intuitively, re-opening the question of p ’s truth-value while at the same time taking that question to be settled comes very close to taking a question to be and not be settled at the same time. So, making use of the plausible principle that when you ought to believe that p you ought not disbelieve it, when you ought to re-open the question whether p you ought not take the question whether p to be settled. To be sure, I do not deny that one might be under an obligation to pursue the re-opening activity for prudential or moral reasons, even in the case in which one is certain of the answer to the question (suppose that a billionaire would invest billions of dollars in famine relief if you did so), and it might also well be the case that prudential and moral obligations trump epistemic ones. Yet since **PR1** is about epistemic requirements only, these considerations do not affect its plausibility.

PR2 is formulated in terms of permission, since steadfast views of peer disagreement are usually understood as claiming that one is permitted – as opposed to required – to retain one’s doxastic attitude in the face of disagreement with an epistemic peer. **PR2** is plausible because if one is permitted to retain one’s belief that p , this means that one continues to have epistemic reasons – relative to one’s epistemic position and however defeasible and

⁴ I am indebted to Sven Rosenkranz here.

misleading they might be – to take the question of p 's truth-value to be settled *in the affirmative*. *A fortiori*, then, one is permitted to take the question whether p to be settled.

I should emphasise that **PR1** and **PR2** hold whether we take doxastic attitudes to be full or graded. Within the quantitative framework of degrees of confidence, it does not make sense to speak of taking the question of p 's truth-value to be settled altogether. Such a question can be taken to be settled to a certain (maximal, minimal, or intermediate) degree. So when stated for credences, **PR1** does not require that one take the question whether p to be settled to the maximal degree. This is clearly compatible with how we have characterized the notion of taking the question whether p to be settled, for this notion is not factive and is to be understood as relativised to one's epistemic position vis-à-vis p . Thus if one's evidence is such that it supports – however defeasibly and misleadingly – taking the question to be settled to a given degree, then one is permitted so to take it. It is therefore plausible to say, as **PR2** does, that if you are permitted to retain your .2 degree of confidence toward p , this means that your epistemic situation continues to allow you to take the question whether p to be settled to a certain degree.

Let us now turn to the normative connections among re-opening, revising_n and taking the question to be settled. Let us first adopt the following:

PR3: If one ought to take the question whether p to be settled, then one ought not re-open the question whether p .

PR3 is made plausible by the same line of reasoning that supports **PR1**.

We need a final principle which bears the same relation to conciliatory views that **PR2** bears to steadfast views. Since conciliatory views are usually taken to specify a requirement and not a permission concerning how to respond to peer disagreement, such a principle will take the following form:

PR4: If one ought to revise_n one's attitude toward p , then one ought to take the question whether p to be settled.

Here is the rationale for **PR4**. Suppose that one ought to revise_n one's attitude toward p . One ought to revise one's degree of confidence toward p from .2, say, to .6. Now suppose for *reductio*, however, that (contrary to **PR4**) it is decidedly not the case that one ought to take the question whether p to be settled. Plausibly, though, if one ought not take the question

whether p to be settled to a certain degree, for instance to degree .6, then there is some (no matter how defeasible and misleading) reason against taking the question of p 's truth-value to be settled to that degree. But if one has some such reason, it seems that one has a defeater for revising one's doxastic attitude toward p from .2 to .6. And in the presence of such a defeater, it cannot be said that one is still required to revise from .2 to .6. But this contradicts our original supposition.

Let us now assume the plausible principle that what is required is also permitted and adopt the following abbreviations:

Reo p = re-opening the question whether p

Ret p = retaining_n one's previous attitude towards p

Set p = taking the question whether p to be settled

O = Requirement/ought operator

\mathcal{P} = Permission operator ($\neg O \neg$)

$\mathcal{D} = \neg(O\phi \wedge O\neg\phi)$

We can now run the following arguments:

Inconsistency of being required to re-open and permitted to retain_n:

- | | |
|---|-----------------------------------|
| (1) O (Reo p) | (Ass1) |
| (2) \mathcal{P} (Ret p) | (Ass2) |
| (3) (O Reo $p \rightarrow O \neg$ Set p) | (PR1) |
| (4) (\mathcal{P} Ret $p \rightarrow \mathcal{P}$ Set p) | (PR2) |
| (5) O (\neg Set p) | (1, 3 MP) |
| (6) \mathcal{P} (Set p) | (2, 4 MP) |
| (7) O (\neg Set p) \wedge \mathcal{P} (Set p) | (5, 6, \wedge -intro.) |
| (8) \perp | (7, definition of \mathcal{P}) |

Inconsistency of being required to re-open and required to revise_n:

- | | |
|--|--------------------------|
| (1) O (Reo p) | (Ass1) |
| (2) O (Rev p) | (Ass3) |
| (3) (O Set $p \rightarrow O \neg$ Reo p) | (PR3) |
| (4) (O Rev $p \rightarrow O$ Set p) | (PR4) |
| (5) O (Set p) | (2, 4 MP) |
| (6) O (\neg Reo p) | (3, 5, MP) |
| (7) O (Reo p) \wedge O (\neg Reo p) | (1, 6, \wedge -intro.) |
| (8) \perp | (7, \mathcal{D}) |

The foregoing discussion shows that requiring peers to re-open the question whether p is inconsistent with both conciliatory views (such as the ones defended by Christensen and Elga and steadfast views of peer disagreement (see Wedgwood 2010, Weintraub 2013) and is a substantively new way of answering *The Question*.

However, it is time now to go back to the claim that I have so far been operating with what I have called a “narrow” notion of doxastic revision. Its narrowness stems from two factors. First, it takes into account *only* revising one’s belief into a disbelief (or *vice versa*), or revising one’s credence cr into another credence cr^* . Secondly, it obeys **PR4**. Yet I believe that we can envisage a wider notion of doxastic revision which (unlike revision_n) is normatively compatible with being required to re-open the question whether p .

Given what **PR1** and **PR3** rule out, we are after a kind of revision from a belief (or a disbelief or a degree of confidence) to a type of attitude such that it is permissible to entertain that attitude without taking the question whether p to be settled. In the next (sub)section I argue that there are in fact *two* distinct types of doxastic attitude that meet this constraint.⁵

2.2 *The Attitude of Suspended Judgement and the Attitude of Hypothesis*

The first type of attitude I want to focus on is the attitude of *suspended judgement*. More specifically, I want to focus on a specific notion of suspended judgement recently defended by Jane Friedman (2013, 2017).

According to Friedman, the attitude of suspended judgement cannot be plausibly reduced to more familiar doxastic attitudes (see especially Friedman 2013). Rather, suspended judgement is a *sui generis* non-propositional question-directed attitude of *epistemic neutrality* which one entertains while inquiring into some matter. Suspended judgement, on Friedman’s view, is the attitude we entertain while we aim to find the answer to a question. For this reason, it comes with “a sort of orientation towards or sensitivity to information that bears on the focal question, and perhaps some other related sorts of dispositions to come to know things that will help one close that question” (Friedman 2017: 305).

In light of the kind of role that this attitude of suspended judgement plays within our mental architecture, it is rationally permissible to suspend judgement in a wide range of

⁵ Recognizing these doxastic attitudes will necessitate a broadening of our conception of doxastic *retention* in parallel with that of doxastic *revision*.

circumstances, for instance when one is curious about a question Q , wants to resolve Q , and so on. (see also Friedman 2017: 321).⁶

I submit that there is no inconsistency in having an attitude of suspended judgement about whether p while under a requirement to re-open the question whether p . To see why, let us look at a kind of counterpart of **PR4**, but for suspended judgement and formulated in terms of permission.

PR4*: If one is permitted to revise one's attitude toward p by suspending judgement, then one is required to take the question whether p to be settled.

PR4* is implausible, in that suspended judgement is exactly the type of doxastic attitude one is permitted to entertain while taking the question whether p to be still open. Thus we cannot run the earlier inconsistency argument. It is permissible, then, for one to suspend judgement while at the same time being under a requirement to re-open the question whether p .

That being said, I believe a compelling case can be made in favour of the existence of a doxastic attitude which I shall call the attitude of *hypothesis*. This attitude—unlike suspended judgement, but like belief—is propositional, but, like suspended judgement, it is one which we entertain while we are still inquiring into p 's truth-value.

To begin with, notice that when we commence inquiry into p 's truth-value, it makes sense to say that we are epistemically neutral toward whether it is true or false. This attitude of neutrality is captured by Friedman's notion of suspended judgement. It is also equally plausible to take an inquirer to close her inquiry into whether p is true or false by believing that p (or by being confident to a certain degree in the truth of p). This happens, for instance, when the inquirer manages to find enough evidence and reasons to close the question in given way, for instance in the affirmative. Yet there is also a third possibility, to the effect that the inquirer has gathered some evidence, deliberated for a while about p , and now moves from being suspended about its truth-value to a less neutral state of mind toward p 's

⁶ I should add that there is a different usage of the label 'suspended judgement' (or 'agnosticism') on which it instead picks out our epistemic inability to determine whether a given proposition is true or false. Rosenkranz (2007) offers a characterisation of this kind of epistemic neutrality, which he dubs "True Agnosticism". True Agnosticism is the stance we adopt when we are not in a position to know the truth-value of a proposition and will continue not to know it relative to all states of information we can reach by expanding current methods of inquiry and cognitive powers. This kind of epistemic neutrality is very different from what Friedman means by suspended judgement, for it amounts to having reached a view and closed the inquiry, albeit in neither the affirmative nor the negative, but, rather, in a *neutral* way.

truth-value which does not yet qualify as a belief state. The scenario I am envisaging is one in which, even though the inquirer is not yet ready to believe/disbelieve p , she is nonetheless cognitively inclined toward p , while still however taking the question whether p to be open. This is the attitude we entertain when we currently take p to be the most promising way of closing the question whether p or not- p , without having already taken the question whether p or not- p to be closed. This intermediate attitude between suspended judgement and belief (or credence) is what I call ‘hypothesis’.

I will now theoretically substantiate the existence of such an intermediate type of doxastic attitude by contrasting it with other doxastic attitudes in the vicinity, such as suspended judgement, belief, supposition, and credence. More details are called for than I can provide in this discussion, but I trust that the following remarks will suffice for purposes of this paper.⁷

As has emerged previously, suspended judgement is a doxastic attitude we entertain when we open the inquiry into p ’s truth-value and which manifests our open-mindedness to the possibility that different candidate answers, say that p is true or that p is false, are equally epistemically good answers to the question of what p ’s truth-value is. This attitude involves inquiry-oriented dispositions to collect evidence and information and to be sensitive to anything which can speak in favour either of p or of its falsity (see Friedman 2017).

Hypothesis, unlike suspended judgement, is the doxastic attitude one entertains towards p when one is inclined toward an affirmative stance on the question whether p is true as the most promising way of answering this question.⁸ This attitude involves *focused* inquisitive dispositions, such as being disposed to: make more effort to check whether the probative force of evidence is in favour of the truth rather than the falsity of p , to check whether the truth of p (rather than its falsity) coheres with other already accepted truths in the vicinity, and so on. In contrast to suspended judgment, hypothesis is *asymmetrical* between p and not- p .

To turn now to belief, it is the attitude whereby we close the question whether a given proposition is true or false. By believing that p , we close the question of p ’s truth-value in the affirmative. This involves very different dispositions from those involved in suspending

⁷ In a companion piece (Palmira ms.), I argue at greater length that hypothesis is a *sui generis* doxastic attitude.

⁸ There are in fact various ways of what I might call *holding* p which do not entail taking the question whether p to be settled (nor do they entail hypothesising that p). By “holding” p I simply mean that the relevant doxastic attitudes are favourable, are asymmetrical between p and not- p , and have propositions as contents.

judgement. It is often claimed that when one believes that p , one is disposed to use p as a premise in theoretical and practical reasoning, and to assert that p .

By contrast, when one hypothesises that p , one is not disposed to assert p . By way of illustration, let us assume that assertion is the kind of speech act one performs when one (overtly) undertakes a commitment to authorise the hearer to assert p and all that logically follows from it, to take up the challenge of vindicating the assertion by offering considerations that justify it, and to retract it if certain conditions are met (see MacFarlane 2011 for discussion). It seems that if one is merely inclined to take p to be the most promising answer to a question one will not recommend to others to assert p , for one is still in the process of establishing whether p really is the answer to the question. Moreover, the very idea of vindicating one's assertion that p seems to obey the following pattern: we've made up our mind about the matter at stake, somebody challenges us about p , and we rehearse our reasons for p or come up with new ones. Thus, one vindicates one's assertion after one has made up one's mind about the topic at stake. However, when one hypothesises that p one has not yet made up one's mind about p . So when one hypothesises that p one can be taken to be disposed to perform a different speech act that we may call, following the lead of Williamson (2000: chapter 11), 'conjecturing p '.

Let me now distinguish the attitude of hypothesis from the attitude of *supposition* (or assuming p for the sake of argument). I submit that supposing p differs from hypothesising p in that the latter is a cognitive pro-attitude whereas the former is not. That is to say, if one hypothesises that p one regards p favourably from an epistemic point of view, but this is not the case with supposing. We can suppose that p in a *reductio* argument in order to show that p is false, precisely because we have a cognitive inclination toward not- p . Or we can suppose p just for the sake of conversation, to enable our conversational partner to carry out her reasoning; clearly, doing so doesn't mean that one is inclined toward p in any sense.

Finally, let me outline the main difference between hypothesis and credences. Since I have argued that one can take the question whether p to be settled by assigning a credence to p , I believe that credences can play the role that full belief plays, in that they are attitudes whereby we can (rationally) close inquiries. This contention is reinforced by the idea that credences, just like full belief, can constitute knowledge,⁹ whereas hypothesis, as an attitude we have precisely when the inquiry is open and we do not know whether p is true or false, cannot constitute knowledge.

⁹ See Moss (2013).

Let me also say something about the normative profile of hypothesis. I take this attitude to be governed by the following norm:

(H-Norm)

For all p , one is permitted to hypothesize that p if and only if holding p enables one to make reliable progress with the inquiry into p being the answer to the question of p 's truth-value.

The right-to-left side of (H-Norm) says that if holding p enables one to make reliable progress with the inquiry into the question of whether p 's truth is the answer to the question of whether p is true or false, we may hypothesize p . The left-to-right side of (H-Norm) says that if holding p as true is not conducive to making reliable progress into discovering whether or not p is the answer to the question of p 's truth-value, then we ought not to hypothesize p .

A couple of clarifications are in order. First, (H-Norm) does not require that one believes that progress will be made. Thus, if one thinks that no progress toward the truth of a given question, for instance the question of whether or not there is free will, can ever be made but is still cognitively inclined towards compatibilism, one is not ipso facto an irrational hypothesiser. Secondly, imposing a reliability condition on inquisitive progress yields the intuitively correct verdict of irrationality in cases where hypothesizers make progress with the inquiry into p being the answer to the question of p 's truth-value by sheer luck through motivated reasoning, confirmation bias, and similarly vicious reasoning methods. The foregoing discussion establishes that requiring peers to revise their original doxastic attitudes by turning them into attitudes of hypothesis is compatible with permitting peers to re-open the question whether p . For if they were incompatible the following principle would be plausible:

PR4**: If one is permitted to revise one's attitude toward p by hypothesising that p , then one is required to take the question whether p to be settled.

PR4** is implausible, though, for the same reason that **PR4*** was. I have argued that the attitude of hypothesis is – together with suspended judgement – precisely the attitude one may entertain when one takes the question whether p to still be open. So, since **PR4**** does not hold for hypothesis, we cannot mount an argument for the inconsistency of requiring

peers to re-open the question while permitting them to entertain attitudes of hypothesis.

This completes my defence of *The Answer*. To repeat it. Whenever you disagree with an acknowledged epistemic peer about whether p , you ought to re-open the question whether p , and you are permitted to revise your initial doxastic attitude toward p in one of two ways: by suspending judgement or by hypothesising. These two types of revision differ from more common types of revision in that one may perform them while taking the question whether p not to be settled. We can add that you are *not* permitted to retain your original doxastic attitude in the face of such a disagreement—conciliationists are right at least to this extent.

In the remainder of the paper I explain how *The Answer* solves *The Puzzle*.

3 Solving the Puzzle

I will first examine how *The Answer* deals with the two cases motivating the Conformist and Non-Conformist Intuitions. Secondly, I will show that the Answer affords the means to what I shall call *The Uniform Solution* to *The Puzzle*.

3.1 *The Answer at Work*

It is important to bear in mind that re-opening the question whether p is an activity whose aim is to enable one to find out more about p 's truth-value by re-checking whether its truth (or falsity) is supported by the available body of evidence, whether new evidence is needed to adjudicate the question, whether the reasoning from evidence to a belief can be deductive or not, and so on. While engaged in these tasks, peers are permitted either to be suspended or to entertain an attitude of hypothesis toward either p or not- p . However, if peers achieve the aim of the re-opening activity, e.g. they reach the verdict that p is true and well supported by the available evidence and, say, one of the peers had mistakenly interpreted the probative force of the evidence, they have most reason to bring the re-opening activity to a close. Depending on the outcome of the activity, viz. depending on how the aim of re-opening the question has been achieved, they ought either to believe that p or disbelieve that p (or have the degree of confidence that is epistemically supported by the evidence which bears on the matter at hand).

This points to the fact that *The Answer* recommends a *dynamic* rather than a *static* way of responding to peer disagreement. Peers are permitted to entertain attitudes of hypothesis or an attitude of suspended judgement during the re-opening activity *only*. If they manage to

satisfy the aim of the re-opening activity, they find out whether p is true or has a favourable epistemic status, and they should then pass out of these attitudes. While the way in which the aim of the re-opening activity is achieved might vary across cases, for sometimes one finds out that p is true, sometimes that it is false, sometimes that we cannot say anything about its truth-value, achieving this aim enables peers to determine what to *believe*. To illustrate this point, take (RESTAURANT). Allison and Marc ought to go over the present body of evidence (e.g. by rechecking the price of each course on the menu) and, if they still disagree, look for new evidence (e.g. recalculating the tip carefully with a calculator, checking whether there is a cover charge to pay, and so on) to settle the question. This re-opening activity will enable them to discover who got it right (if either of them did). Assume that Allison got it right. That is to say, the proposition that they each owe \$43 is true and supported by the evidence. If so, then the outcome of the re-opening activity gives her reasons to demote Marc's epistemic credentials and stick to her guns. In light of the re-opening activity, Marc acquires reasons not to believe that they each owe \$45.

However, there might be disagreement cases in which the re-opening activity fails to yield a clear and conclusive outcome such that its aim can be regarded as achieved. In such cases, peers do not have any (epistemic) reason to bring the re-opening activity to a close by determining that they ought to have a certain belief towards the contested proposition.

I contend that (A PRIORI) might plausibly be regarded as an instance of such cases. Even after reconsidering the question whether there is a meaningful distinction between *a priori* and *a posteriori* justification for a long time and with great attention, it will be far from clear who as between Rudy and Ruth is making a mistake about such an issue; what the relevant body of evidence bearing on it is and how to assess the probative force of each alleged piece of evidence; whether Rudy and Ruth rely on incommensurable yet equally objectively valuable general philosophical principles; whether they use the concepts of A PRIORI and A POSTERIORI in slightly different ways; and so on and so forth.

The Answer maintains that peers involved in such complex disagreements ought to keep on inquiring into the question by canvassing and improving their respective epistemic positions, looking for new evidence, reasons, and arguments that may help settle the question. In (A PRIORI) *The Answer* would say that Ruth and Rudy are permitted either to be suspended or to sustain their difference of opinion by holding incompatible hypotheses with respect to p and its negation, while still pursuing inquiry into the question of whether there is a meaningful distinction between *a priori* and *a posteriori* justification.

3.2 *The Uniform Solution of the Puzzle*

The Answer enables us to articulate *The Uniform Solution* to *The Puzzle*. It is a solution in that it enables us to capture what makes the Conformist Intuition and the Non-Conformist Intuition compelling without leading to the puzzling (a)-(d). It is uniform in that it maintains that peers ought to comply with the same requirements in both (RESTAURANT) and (A PRIORI). Let me articulate both points in more detail.

First, by requiring peers to re-open the question whether p , *The Answer* acknowledges that peer disagreement carries *some* epistemic significance. This is what I take to be the deep point conveyed by the Conformist Intuition and cases like (RESTAURANT). One might worry, however, that this does not amount to a strong vindication of the Conformist Intuition. Such a strong vindication would require endorsing the claim that when you discover a disagreement with an acknowledged peer, you have evidence that you have made a mistake. This is claim (a) of *The Puzzle*. I believe that *The Answer* on offer must deny (a). To see why, notice that *The Answer* denies claim (c) of *The Puzzle*, which says that peers ought to change their views.¹⁰ Bear in mind that (c) follows from (a) and (b). Since we can regard claim (b) as fairly plausible, we must conclude that *The Answer* denies (c) in virtue of denying (a).

Closer inspection reveals that a rejection of (a) is not only plausible, but also desirable. To see why, it must be noticed that claim (a) is bolstered by the widespread assumption that in order for peer disagreement to be epistemically significant it has to be regarded as a new piece of evidence that one has made a mistake, to be added to the first-order body of evidence that bears on the question. However, taking disagreement be an additional piece of evidence to the effect that one has made a mistake has proven to have unpalatable consequences for both conciliatory and steadfast views (see e.g. Lasonen-Aarnio 2013 and Rosenkranz and Schulz 2015).

Fortunately, authors such as Christensen (2010) Lasonen-Aarnio (2014) and Rosenkranz and Schulz (2015) articulate at length a better understanding of the epistemic significance of peer disagreement. These authors defend the contention that the epistemic significance of disagreement lies in bringing out the possibility that one had mistakenly assessed the extension or the probative force of the body of evidence to begin with. That is to say, it is not that the discovery of disagreement makes one's belief no longer supported by

¹⁰ This is so provided that "change their views" is to be understood as "revise_n their views." (We have been assuming that the narrow notion of revision is the one people are standardly operating with.)

the evidence and therefore in need of radical revision. Rather, disagreement is *higher-order* evidence that casts doubt on whether one's belief was epistemically supported to begin with.

Under this understanding of the epistemic significance of peer disagreement, it is quite natural to claim that peers have most reason to scrutinise their initial epistemic situation by re-assessing what the original body of evidence is and how to interpret its probative force. Hence, it seems safe to claim that *The Answer* chimes in well with this plausible interpretation of the epistemic significance of disagreement.

I have suggested that a strong vindication of the Conformist Intuition has in the background a problematic account of the epistemic significance of disagreement. Moreover, *The Answer* offers a moderate vindication of what makes the intuition so attractive, namely that peer disagreement is somehow epistemically significant, which is compatible with a less problematic account of such significance. I therefore contend that there is no explanatory loss in rejecting claim (a) of *The Puzzle*, quite the opposite: we have good reason to reject it. *The Answer* is compatible with a better way of understanding the epistemic significance of peer disagreement which does not entail (a).

The Answer is also able to capture the deep point conveyed by the Non-Conformist Intuition. *The Answer* permits peers to entertain attitudes of hypothesis, rather than requiring them to suspend judgement or update their respective degrees of confidence about the matter at issue. *The Answer* thus vindicates the driving thought behind claim (d) of *The Puzzle*, namely that there are disagreement cases – such as (A PRIORI) – in which peers ought not change their view. So Rudy and Ruth are entitled to be cognitively inclined toward, respectively, the truth and falsity of the proposition that there exists a meaningful distinction between *a priori* and *a posteriori* justification while, at the same time, continuing to inquire into it.

4 Conclusion

I have argued that peers ought to respond to their disagreement by re-opening the question of the truth of the relevant matter at hand. This requirement is normatively compatible with two types of doxastic revision: revising by suspending, and revising by hypothesising. This constitutes *The Uniform Solution* to *The Puzzle*: peer disagreement is epistemically significant in all its manifestations, but acknowledging such epistemic significance does not require peers to suspend judgement about the matter at hand or to change their view. *The Uniform Solution* to *The Puzzle* which *The Answer* affords does not fully vindicate the Conformist Intuition,

since it rejects claim (a) of *The Puzzle*. This is however done on the basis of a principled understanding of the epistemic significance of peer disagreement which chimes in with the contention that peers ought to re-open the question whether p . By rejecting claim (a), *The Uniform Solution* also rejects claim (c) of *The Puzzle*. However, since *The Answer* acknowledges that peer disagreement carries some significance which requires peers to re-open the question whether p , *The Uniform Solution* offers a moderate vindication of the Conformist Intuition. Moreover, by permitting peers to retain their respective cognitive inclinations toward the matter at hand, *The Answer* also offers a moderate vindication of the Non-Conformist Intuition.¹¹

References

Christensen, D. (2007) “Epistemology of Disagreement: the Good News”, *The Philosophical Review* 116(2): 187-217.

Christensen, D. (2010) “Higher-Order Evidence”, *Philosophy and Phenomenological Research* 81(1): 85-115.

Christensen, D. (2011) “Disagreement, Question-Begging, and Epistemic Self-Criticism”, *Philosophers’ Imprint* 11(6): 1-22.

Elga, A. (2007) “Reflection and Disagreement”, *Noûs* 41(3): 478-502.

Everett, A. (2015) “Peer Disagreement and Two Principles of Rational Belief”, *Australasian Journal of Philosophy* 93(2): 273-286.

Feldman, R. (2006) “Epistemological Puzzles about Disagreement”, in Hetherington S. (ed.),

¹¹ I gratefully acknowledge the financial support of the Spanish *Ministerio de Economía y Competitividad* (MINECO) under grant agreement #FFI2016–80588-R, the European Commission’s Horizon 2020 programme under grant agreement H2020-MSCA-ITN-2015-675415, the Juan de la Cierva and Beatriu de Pinós postdoctoral fellowship programmes, and a postdoctoral fellowship from the Centre for Research on Ethics of the University of Montreal and McGill University. I wish to thank audiences in Barcelona and Montreal for helpful comments on various parts of this material. Special thanks are due to Sven Rosenkranz, who provided detailed terrific comments on a late version of the paper. Extra special thanks are due to Hichem Naar and Sarah Stroud for their constant and insightful feedback on this project.

Epistemology Futures, Oxford: Oxford University Press: 216–236.

Friedman, J. (2013) “Rational Agnosticism and Degrees of Belief”, in Gendler, T. and Hawthorne, J. (eds.), *Oxford Studies in Epistemology*, Vol. 4, Oxford: Oxford University Press: 57-81.

Friedman, J. (2017) “Why Suspend Judging?”, *Noûs* 51(2): 302-326.

Kelly, T. (2005) “The Epistemic Significance of Disagreement”, in Gendler, T. and Hawthorne, J. (eds.), *Oxford Studies in Epistemology*, Vol. 1, Oxford: Oxford University Press: 167-196.

Kelly, T. (2013) “Disagreement and the Burdens of Judgment”, in D. Christensen and J. Lackey (eds.) *The Epistemology of Disagreement. New Essays*, Oxford: Oxford University Press: 31-53.

Lasonen-Aarnio, M. (2013) “Disagreement and Evidential Attenuation”, *Noûs* 47(4): 767-794.

Lasonen-Aarnio, M. (2014) “Higher-Order Evidence and the Limits of Defeat”, *Philosophy and Phenomenological Research* 88(2): 314-345.

Moss, S. (2013) “Epistemology Formalized”, *The Philosophical Review* 122(1): 1-43.

Palmira, M. (ms.) “The Inquiring State of Mind”.

Rosenkranz, S. (2007) “Agnosticism as a Third Stance”, *Mind* 116(461): 55-104.

Rosenkranz, S. and Schulz, M. (2015) “Peer Disagreement: A Call for the Revision of Prior Probabilities”, *dialectica* 69(4): 551-586.

Shah, N. and Velleman, D. J. (2005) “Doxastic Deliberation”, *The Philosophical Review* 114(4): 497-534.

Wedgwood, R. (2010) “The Moral Evil Demons”, in *Disagreement*, Feldman, R. and Warfield, T. (eds.), Oxford: Oxford University Press: 216-246.

Weintraub, R. (2013) “Can Steadfast Peer Disagreement Be Rational?”, *Philosophical Quarterly* 63(253): 740-759.