Competent inference and the (ir)rationality of level-splitting

Inferência competente e a (ir)racionalidade do level-splitting

10.21680/1983-2109.2025v32n68ID38611

Veronica de Souza Campos

Faculdade Jesuíta de Filosofia e Teologia (FAJE)

182vkai@gmail.com

Resumo: Em um artigo de 2012, Ralph Wedgwood defendeu a visão de que inferências competentes são sempre inferências racionais. Sua teoria da significância de inferências competentes faz um acréscimo a um debate proeminente na epistemologia analítica contemporânea, o debate em torno da racionalidade de respostas level-splitting. Grosso modo, os casos de inferência competente de Wedgwood foram apontados como exemplificando uma situação em que tais respostas seriam racionais. Neste artigo, argumentarei contra Wedgwood desafiando uma das suposições que sustentam sua visão, a saber, a suposição de que fazer uma inferência competentemente é suficiente para tornar racional, para um agente, endossar sua conclusão de uma forma que seja independente das circunstâncias. Como mostrarei, as circunstâncias importam. Para que level-splitting seja racional em casos de inferência competente da forma aludida, ela teria que ser a melhor resposta disponível, ou seja, melhor do que respostas alternativas que o agente poderia ter, mas não está claro que seja.

Palavras-chave: Level-splitting; Evidência de ordem superior; Requisitos de racionalidade; Inferência competente.

Abstract: In a 2012 paper, Ralph Wedgwood has argued for the view that competent inferences are always rational inferences. His theory of the significance of competent inferences jots down at a prominent debate in contemporary analytic epistemology, the debate around the rationality of level-splitting. Roughly, Wedgwood's cases of competent inference have been pointed to as exemplifying a situation whereby level-splitting is rational. In this paper, I'll argue against Wedgwood by challenging one of the assumptions underpinning his view, namely, the assumption that drawing an inference competently suffices to make it rational for an agent to endorse its conclusion in a way that is independent of the circumstances. As I'll show, circumstances matter. For level-splitting to be rational in cases of competent inference in the way alluded to, it would have to be the best response available, i.e., better than alternative responses the agent could have, but it is not at all clear that it is.

Keywords: Level-splitting; Higher-order evidence; Requirements of rationality; Competent inference.

1. The rationality of level-splitting

A prominent debate in contemporary analytic epistemology concerns the rationality of the so-called level-splitting response to one's total evidence. One's response to one's total evidence is said to be a level-splitting response whenever one's response to one's first and higher-order evidence come apart. This response will, thus, encompass a conjunction of two conflicting doxastic attitudes (beliefs, or credences), that might take the form "***p***, but the evidence does not support ***p***"; or even "***p***, but I shouldn't be confident that ***p***".

A good way to get clear on the type of situation that will typically give rise to a discussion about whether level-splitting could be rational is by considering the case of hypoxia,

presented by David Christensen (2010). Hypoxia is a condition in which the body is deprived of adequate oxygen supply. It commonly affects mountain climbers and aircraft pilots, who are exposed to the low pressures encountered at high altitudes (10,000 feet is often mentioned as the mark from which it becomes a concern). Hypoxia causes impaired judgment, but it is very rarely recognized by the sufferer at its initial onset. This is Christensen's original case:

> **PILOT**
>
> You're alone, flying a small plane to a wilderness airstrip. You're considering whether you have enough fuel to make it safely to an airstrip 50 miles further away than your original destination. Checking the relevant factors, you become extremely confident that you do have sufficient fuel — you figure that you've got a fair bit more than the safety-mandated minimum. So you begin your turn toward the more distant strip. But then you notice that your altimeter reads 10,825 feet. You feel completely clearheaded and normal; however you're fully aware of the insidious effects hypoxia can have. Should you trust your recently formed confident judgment about having sufficient fuel, and continue on your path toward the more distant airstrip? (Christensen, 2010, p. 126-127)

PILOT can be given a more technical description as follows. Subject ***S*** wants to know whether "fuel is sufficient to get to the more distant airport" (***p***). He does the math, taking all the relevant parameters and data into account, and gets the result ***m***, whereby ***m*** indicates that fuel is sufficient. In this scenario, ***m*** is first-order evidence that ***p***. But then, the altimeter reminds ***S*** of the phenomenon of hypoxia. The information displayed by the altimeter, m_2, suggests that ***S***'s math might have been flawed (for instance, he might have miscalculated, or failed to take all the data that happens to be relevant into account). So m_2 is ***S***'s high-order evidence: it doesn't concern ***p*** directly, but rather the support relation

between $\boldsymbol{m}$ and $\boldsymbol{p}$. Specifically, $\boldsymbol{m_2}$ indicates that $\boldsymbol{m}$ might not support $\boldsymbol{p}$. In this situation, the pilot could respond in different ways. He could:

1. Trust $\boldsymbol{p}$ as well as his judgment that $\boldsymbol{m}$ supports $\boldsymbol{p}$.

2. No longer trust $\boldsymbol{p}$, neither his judgment that $\boldsymbol{m}$ supports $\boldsymbol{p}$.

3. Trust $\boldsymbol{p}$, but no longer trust his judgment that $\boldsymbol{m}$ supports $\boldsymbol{p}$.

4. No longer trust $\boldsymbol{p}$, but trust his judgment that $\boldsymbol{m}$ supports $\boldsymbol{p}$.

In the face of dilemmas such as the one depicted in pilot, scholars have taken varied positions. As others have observed (e.g. Roush, 2017, §3), there is something inherently problematic about response 4, for even if one can't actually bring himself to trust $\boldsymbol{p}$, if it is rational for him to trust that $\boldsymbol{m}$ supports $\boldsymbol{p}$, then this by itself should suffice, prima facie, to render it rational for him to trust $\boldsymbol{p}$, even if $\boldsymbol{p}$ is false. We're thus left with scenarios 1-3. Scenario 1 is what came to be known in this literature as "the steadfast response". It has been defended by scholars such as Thomas Kelly (2005) and Miriam Schoenfield (2018), who argue that, since it is not certain whether $\boldsymbol{m_2}$ defeats the support relation between $\boldsymbol{m}$ and $\boldsymbol{p}$, the agent should stick to his guns: maintain that the support relation holds, and maintain his recently formed judgment. Scenario 2, in turn, is what came to be known as "the conciliatory response". Scholars such as Richard Feldman (2005), Adam Elga (2007, 2013) and Christensen himself (2007, 2010) have argued that this is the most rational response. According to them, the agent should take $\boldsymbol{m_2}$ as as having some impact on the support relation between $\boldsymbol{m}$ and $\boldsymbol{p}$, so that the recently formed judgment ($\boldsymbol{p}$) is no longer supported and, thus, should no longer be trusted to the same

extent[1]. And last, scenario 3 is what came to be known as the level-splitting response. It has been argued by scholars such as Timothy Williamson (2011, 2014), Sophie Horowitz (2014), Hawthorne, Isaacs & Lasonen-Aarnio (2015) and Christensen (2024) that this response could be rational under some circumstances. The rationale underpinning their view is that the level-splitting response is the only response that actually respects all the evidence.

The cases that have been fuelling defences of level-splitting are considerably limited. They are basically variations of pilot-like cases that we can sort in three sets. The first set comprises the so-called cases of improbable knowing, whereby an agent knows something but the evidence suggests that it is very unlikely that he does so. Those cases have been discussed in Williamson (2000, 2011, 2014); and by others, such as Elga (2013). The second set is composed by analogies with moral cases, whereby the most rational response an agent can give to a moral conundrum involves inconsistency between what they believe they should do and what they end up doing. Those have been discussed in Brian Weatherson (2010) and Allen Coates (2012)[2]. And the third set, that will

[1] Some scholars purport that whether the most rational response is steadfast or conciliatory will vary across different PILOT-like situations. For a discussion of this position, see Kelly (2010).

[2] Weatherson recreates epistemic analogs of the classical kantian yarn of the murderer at the door, whereby the most rational response involves lying to the murderer, even though one still believes the kantian maxim that lying is wrong. Coates, in turn, recreates an epistemic analog of Mark Twain's story of Huckleberry Finn, whereby an agent refrains from handing a runaway slave out to the authorities, in spite of believing that that would be the right thing to do. The common idea underpinning their arguments is that just like in some moral conundrums the most rational response an

be my focus here, is comprised by the so-called cases of competent inference followed by higher-order counter-evidence, whereby an agent, much like the pilot in pilot, reaches a conclusion by means of a competent inference, but then encounters evidence indicating that his inference wasn't actually competent. Those cases have been presented and discussed by Ralph Wedgwood (2012), who embraces the view that level-splitting is the most rational response.

In this way, one thing that is important for the debate about the rationality of level-splitting is examining if, or to which extent, these sets of examples actually support the case for the rationality of level-splitting *more than* they support the alternative responses, i.e., the steadfast and/or the conciliatory response. Do cases of improbable knowledge, analogies with moral conundrums and competent inference followed by higher-order counter-evidence actually give us compelling grounds for the claim that level-splitting is the best response an agent can give? In this paper, I'll focus exclusively on the latter, the cases construed by Wedgwood, and I'll argue that they do not warrant a pro-level-splitting verdict, at least not more so than alternative verdicts.

This being said, here is the plan for the paper. In section 2, I'll present Wedgwood's conception of competent inference. In sections 3 and 4, I'll present two arguments against the claim that level-splitting is rational in situations of competent inference followed by higher-order counter-evidence. These are not intended as arguments against the rationality of level-

---

agent could give might involve inconsistency between what one does and what one believes that he should do, in some epistemic conundrums the same sort of situation comes by. For an excellent discussion of this set of cases, see Joshi (2016) and Lo Guercio (2018).

splitting per se, but rather against the rationality of level splitting under those specific circumstances (in situations whereby a competent inference is followed by higher-order counter-evidence). The takeaway of my discussion will be that level-splitting is not a rational response in those cases, since an analysis of the cases doesn't make it clear how, or why, it would be more appropriate than alternative responses an agent could give. I'll conclude, in section 5, by summarizing and briefly discussing the implications of this analysis.

2. Wedgwood's view of competent inference

In a 2012 paper titled "Justified Inference", Ralph Wedgwood advances a theory of the epistemic significance of inferences. According to this theory, inferences are rational whenever they are justified, and that inferences are justified whenever they have been drawn competently. He states that:

> (…) it is the real nature of this internal process—the fact that it is a process of competent inference—and not the higher-order beliefs that the thinker has, or even the beliefs that the thinker is justified in having, about the nature of that process, that is crucial to the rationality of the mental event that results from that process. (Wedgwood, 2012, p. 292)

Thus he purports that "(…) competent inferences are, without exception, always rational" (Wedgwood, 2012, p. 285). I take him to be stating here not that competent inferences are rational by a matter of stipulation, or that it is analytic to say that competent inferences are rational inferences, but rather that having been drawn competently suffices to render an inference rational.

It is not completely clear, from Wedgwood's words alone, what competently drawing an inference is a matter of. He says that "on this conception, a competent inference consists in the manifestation of a disposition that can only be

manifested in the coherent acceptance of a rationally accessible argument", so that "(…) whenever one manifests such a disposition, one's drawing the inference in question can be called a case of *competent inference*" (Wedgwood, 2012, p. 285, italics original). What exact disposition this is, and whether it is necessary or sufficient to render the inference competent, or both, remains unclear. One plausible hypothesis is that the disposition in question is a matter of possessing an embedded skill, for instance, being able to apply the rules of inference consistently throughout[3].

Also, according to Wedgwood, once a rational inference has been drawn, accepting the inference is the same as conditionally believing its conclusion (Wedgwood, 2012, p. 278)[4]. In this way, beliefs formed through justified inferences will always be justified, and any belief one comes to have as a result of rational inferences will be rational to have. He says:

[3] This hypothesis is not related to the Russellian notion of knowledge by acquaintance, or to the notions of know-how and/or tacit knowledge put forward by Gilbert Ryle in *The Concept of Mind* (1949). Rather, I'm thinking here in terms of Alvin Goldman's (1979) process reliabilism. Goldman's theory was aimed at explaining why, or when, a belief is justified (it is justified whenever it has been produced by a reliable cognitive process). Here, I'm borrowing this general idea as the basis for an hypothesis aimed at explaining why, or when, an inference is competent (it is competent whenever it has been produced by a reliable cognitive process).

[4] He says: "Accepting an argument is (...) not an attitude towards a single proposition, but towards the whole argument—the attitude of conditionally believing the conclusion, on the assumption(s) of the premise(s)" (Wedgwood, 2012, p. 278).

> Suppose that you also have the following disposition: in every normal case in which you rationally draw an inference from a set of premises each of which you have some kind of rational unconditional belief in, you respond to your drawing this inference by incorporating the conclusion into your belief-system, in such a way that the new belief-system is just like the old one, except for the fact that (i) the new belief-system incorporates the conclusion into its overall ranking of propositions, and (ii) the new belief-system also meets all conditions of rational coherence. My final proposal about the epistemological significance of inference is the following: The mental event of incorporating a proposition into your belief system by manifesting this sort of disposition is always a rational mental event. (Wedgwood, 2012, p. 286)

This proposal is what Wedgwood refers to as "Proposal 3", since it is the third in a series of proposals that make up his theory of the epistemic significance of inferences. It has already been observed by others – e.g., Horowitz (2013) – that Proposal 3 has a counter-intuitive consequence: it entails that the position known in contemporary epistemology as level-splitting will be rational in cases of competent inference. Level-splitting, as discussed in the previous section, is the view according to which it is sometimes rational for an agent to endorse an inconsistent combination of doxastic attitudes, such as "***p***, but my evidence does not support ***p***". According to Wedgwood, if one's belief that ***p*** has been acquired as a result of a competent inference, then ***p*** will be justified and holding it will always be rational, hence one will be compelled to hold on to ***p*** *even if* one comes to encounter further evidence against ***p***.

Following Wedgwood's line of reasoning, one's endorsing the conclusion one takes his or her reasoning to entail will continue to be rational insofar as the inference leading to that conclusion was indeed performed competently. New evidence won't make it irrational, unless this new

evidence casts doubt on the nature of the reasoning itself, and the reasoning itself was flawed. That is to say, new evidence makes no difference to the rationality of endorsing the conclusion of an inference unless it indicates that the inference was not drawn competently, and the inference, in fact, was not drawn competently. If the new evidence indicates that the inference was not drawn competently but it turns out that it was drawn competently (i.e., if what we have is misleading higher-order evidence), then the agent will be justified in holding on to its conclusion, even if they are led to believe, on the basis of the new evidence, that that conclusion is wrong and that holding on to it is irrational.

One might write this consequence of Proposal 3 down in a more technical fashion in the following way. Say subject S competently draws the conclusion that c (from, say, evidence E) and then comes to access higher-order evidence E_2, whereby E_2 indicates that E does not support c. S might acknowledge that E_2 casts doubt on the outcome of his reasoning, and he may even come to believe that c is false – but this not for a moment undermines the rationality of the reasoning leading to c, nor the rationality of endorsing c, provided that the reasoning leading to c was indeed performed competently. In such cases, S might find himself in the weird situation of believing he is being irrational in endorsing c (by virtue of accommodating E_2) but endorsing c will be rational nonetheless (by virtue of c having been drawn competently).

By committing himself to the claim that being competent suffices to render an inference rational, Wedgwood is compelled to accept that level-splitting in cases of competent inference followed by misleading higher-order evidence is acceptable, or even desirable. This is because, under such circumstances, the agent has reasons to withdraw the conclusion and to maintain it. As a rational agent, he is

under pressure to accommodate any new evidence concerning that matter, since it would be odd to say that it could be reasonable of a rational agent to remain entirely oblivious to new evidence regarding matters he has been reasoning about, especially whenever this new evidence counters his former conclusions. But, as a competent agent, he is also under pressure to trust the results of his own inferences. So it seems that, if Wedgwood is right, in this sort of situation the agent ought to stick to the conclusion of his former reasoning, even if he is convinced that, in so doing, he is acting irrationally.

Two arguments might be presented against the claim that level-splitting is rational in cases of this type. The first argument, that will be presented in section 3, is a sufficiency argument, in that it tackles Wedgwood's seemingly equating the rationality of an inferential process with the rationality of endorsing its conclusion. If this objection works out, the fact that an inference was drawn competently should not suffice to render the endorsing of its conclusion a rational act, since sometimes the reaching of a conclusion calls us to suspend judgment and revise the inference, rather than to endorse it, even if it was, in fact, drawn competently; and this, in turn, is a variety of conciliatory response to total evidence. The second argument, in section 4, is a design argument, in that it tackles Wedgwood's makeup of his view of competent inferences, a view that accounts for the way in which a very specific sort of epistemic agent is expected to reason. If this objection works out, it should make it clear that the epistemic situations Wedgwood's view was designed to account for actually require from the agent the so-called steadfast response, not a level-splitting response.

3. The sufficiency claim

Wedgwood's view brings out a few problems. First, circumstances matter. The idea that drawing an inference competently suffices to render the endorsement of its conclusion a rational act regardless of the circumstances is remarkably problematic. While you can competently derive a conclusion from a set of reasonable premises, and while endorsing that conclusion might seem like the obvious thing to do, it might not be the most rational thing to do, or at least not independently of the circumstances. The rationality of endorsing a conclusion does seem to depend strongly on the premises from which the conclusion has been drawn being true, for one can draw an inference pretty competently and yet arrive at a false conclusion, because one or more premises were false. Likewise, one can draw an inference competently, and yet arrive at a false conclusion because of missing information one didn't have. Here are two simple low level examples of this.

> **MERCURY**
>
> Sam finds a glass bottle in the chemistry lab filled with a shiny silver-white metal. He observes that the thermostat shows it is 25°C in the room. He then reasons in the following way: The substance inside of this glass bottle is liquid at 25 °C; and, in general, metals are solid at 25 °C. Therefore, the substance inside of this glass bottle is not a metal.
>
> **TWIN PLANET**
>
> Kim comes to know that Venus has a similar size and mass to Earth. He also comes to know, through telescopic observations, that Venus has a cloud-covered atmosphere. Kim then infers that Venus has an Earth-like environment, possibly suitable for life.

In MERCURY, Sam draws the conclusion that the substance inside of a glass bottle is not a metal. This

conclusion is false, because it stems from the false premise that metals are solid at 25 °C; but the inference is competent, insofar as the conclusion indeed follows from the premises. In TWIN PLANET, likewise, Kim draws the conclusion that Venus has an Earth-like environment. This conclusion is false, since, unbeknownst to Kim at the time, the extreme greenhouse effect on Venus leads to surface temperatures over 450°C and an atmosphere dominated by carbon dioxide and sulfuric acid clouds, which renders Venus unsuitable to life as we know it. But the inference, again, is competent, since the conclusion is reasonable given the premises. If this is so, and if Wedgwood is right in what he says about every competent inference being rational, then it is rational for Sam and Kim to endorse their false conclusions, provided that they have been drawn competently, which they have.

While strictly speaking this is not unacceptable (that it can be rational for subjects to endorse false conclusions), the problem begins when higher-order counter-evidence enters the stage. When, say, Sam learns that mercury is a metal that is liquid at 25°C, or when Kim learns about extreme greenhouse effects in Venus. This sets up circumstances under which to simply continue to endorse the conclusions in question might no longer be the most rational attitude. That's because in both cases the subject, as a rational epistemic agent, is now under pressure to accommodate the new evidence. Accommodating the new evidence is not just a matter of computing it; it involves taking into account what the new evidence means. If Sam hears of mercury from his colleague, it would be odd to say that it could be reasonable for him to just incorporate this evidence (in the form of an extra belief) into his doxastic mesh without reflecting on its meaning, as much as it would be unreasonable for him to remain entirely oblivious to new evidence, especially since the new evidence indicates that his former conclusion is possibly false. If this

new evidence was misleading, it could perhaps be fine to ignore what it means, but Sam – and this is crucial – doesn't know, after all, whether the new evidence is misleading or not. The intuition the example exposes is clear: if the information comes to Sam's awareness that there is one metal that is liquid at room temperature, he should not maintain the same level of confidence in the original conclusion of his reasoning, that the shiny substance inside of the glass bottle is not a metal.

What this means is that circumstances matter for the rationality of endorsements far more than Wedgwood seems inclined to accept. For higher-order evidence countering what is in fact a false conclusion is actually evidence pointing to the truth; and it is hard to see how sticking to the false when the truth has just been exposed or at least pointed towards could possibly be the most rational thing to do, *even if* the inference leading to that conclusion has been drawn competently. Especially, it seems as though knowing that the new evidence is misleading (or that it is not) plays a role. The problem is, when confronted with new evidence, we don't typically know that.

MERCURY and TWIN PLANET are "easy cases", since they are cases whereby the conclusions of the inferences are clearly false, and the higher-order evidence discussed in relation to them has been stipulated as being clearly not misleading. Wedgwood could claim that these are not cases of competent inference, after all, to the extent that they contain false premises. In other words, he could contend that Sam and Kim are not competent to make the inferences that they made, not because they are ignorant of the rules of inference (which they are not), but because they are not knowledgeable enough to be reliable in those matters, and that competent inferences are necessarily those that are drawn by reliable agents, in this sense. Those are more demanding standards for competent inference, but they are still plausible. So, how would my

objection fare in a not-so-easy case, i.e., one that does not contain premises that are obviously false, and in which the reasoning subject is qualified to reason about the matter? Consider NEPTUNE, below.

> **NEPTUNE**
>
> In the beginning of the nineteenth century, it was discovered that the orbit of Uranus, one of the seven planets known at the time, departed from the orbit as predicted on the basis of Isaac Newton's theory of universal gravitation. From this, Jim, a young astronomer in the 1830's infers that there must be an error within Newton's theory of universal gravitation. As it turned out, however, two astronomers, J. C. Adams and U. Leverrier, instead suggested that there was an eighth, yet undiscovered, planet in the solar system; which, they thought, explained Uranus' deviating orbit within a Newtonian framework. Not much later, this planet, which is now known as Neptune, was discovered.[5]

Here is a simplified and hypothetical outline of Jim's inference:

> **(1)** The orbit of Uranus displays a series of irregularities that cannot be explained by Newton's theory of universal gravitation.
>
> **(2)** If Newton's theory of universal gravitation was empirically adequate, it would be capable of explaining the orbit of Uranus.
>
> Thus,
>
> **(3)** Newton's theory of universal gravitation is not empirically adequate.

Jim's inference was drawn competently, since it obeys the rules of inference, and since Jim is an astronomer (a

[5] Adapted from Douven (2017).

reliable epistemic agent, when it comes to matters of astronomy). Thus, even in Wedgwood's more demanding standards, this is a rational inference. Notwithstanding, by means of this inference, Jim has reached what we can take to be a false conclusion, since, unbeknownst to him at the time, premise 1 is false. As a consequence, Jim has acquired a false belief. According to Wedgwood, the fact that the inference is competent suffices to make it rational for Jim to hold this belief.

To be sure, Jim is reasoning from a scientific finding (that the orbit of Uranus presents certain irregularities) and a reasonable scientific expectation (that Newton's theory of universal gravitation was supposed to explain those irregularities). And he is taking, as an auxiliary (and hidden) premise, what was, at his time, the overall accepted proposition that the number of planets in the Solar System was seven. Someone could object that since this hidden premise is false, despite being widely accepted at the time, Jim's inference is actually not competent. But this objection doesn't go a long way. The one who objects along these lines would be saying, in other words, that no astronomer reasoned competently prior to the discovery of the correct number of planets in the Solar System, and perhaps that no scientist of our time or of any time has ever reasoned competently, since virtually every scientific proposition experts deem true today (as well as at any time) can end up being proven to be false at some point in the future. Thus, holding that Jim's inference is not competent because one of its premises is false, despite its being deemed true at the time, has its costs – it comes at the expense of radical skepticism. He who accepts it is saying that, strictly speaking, no one can ever know, of any particular inference, whether it is competent or not; and Wedgwood's theory of competent inference would become pointless. To avoid the unnecessary complications brought about by radical

skepticism, let's take it, *prima facie*, that Jim's inference is competent, since there is no self-evident reason why it wouldn't be.

Now, given the great empirical success of Newton's theory for what was then more than two centuries, doesn't it seem that things would have turned out better for Jim (that is, doesn't it seem that he would have been more rational) had the outcome of this inference, 3, motivated him to doubt his premises, rather than to come to believe the conclusion? A good reason why this would have been better is that, had Jim questioned his premises instead of endorsing the conclusion, he would have ended up closer to the truth than he actually did. This is especially so if we add to this tale the consideration, which is actually a fact, that what motivated astronomers U. Leverrier and J. Adams to make the telescopic observations which ultimately disclosed that planet Neptune existed was precisely their questioning the assumption that the number of planets in the solar system was seven, having seen that, if it were seven, then Newton's theory of universal gravitation would be empirically inadequate. This is why Neptune was mathematically predicted before it was directly observed, and also why its observation was regarded as a conspicuous confirmation of Newtonian theory of universal gravitation by the scientific community at the time.

Here is a simplified and hypothetical outline of U. Leverrier and J. Adams' reasoning:

> **(1a)** The orbit of Uranus displays irregularities that cannot be explained by Newton's theory,
>
> unless
>
> **(1b)** there is another heavy body, the size of a planet, disturbing its orbit, that we don't know about.
>
> So, either

(2a) Newton's theory of universal gravitation is not empirically adequate,

or

(2b) there must be another heavy body, the size of a planet, disturbing the orbit of Uranus.

Not-(2a): Newton's theory of universal gravitation is a good theory and has achieved great empirical success for more than two centuries.

Thus,

(2b) There must be another heavy body, the size of a planet, disturbing the orbit of Uranus. Let's find it out.

Both Jim's *modus tollens* and Leverrier & Adams' inference to the best explanation have been drawn according to the rules of inference, and by competent agents. Both are, therefore, competent inferences, even under Wedgwood's more demanding standards. If Wedgwood's claim that an inference being competent suffices to make it rational for an agent to believe its conclusion is correct, then he would have to say that believing the conclusion of both inferences, Jim's and Leverrier & Adams', is equally rational, to the extent that both inferences are equally competent. This is the same as saying that it is rational for Jim to believe the conclusion of his *modus tollens*, as much as it is rational for Leverrier & Adams to believe the conclusion of their inference to the best explanation. Up to a point whereby Jim is not aware of Leverrier & Adams' rationale, this prognosis does not sound so bad. But would that continue to hold from the moment Jim finds out about the discovery of planet Neptune?

Say Jim opens the latest scientific newspaper one day and comes across a full coverage of U. Leverrier and J. Adams' feat. He accesses the reasoning leading these astronomers to posit that Neptune existed, and the experiment through which

they confirmed that it did and, thus, that Newton's theory of universal gravitation could, in fact, explain the orbit of Uranus. What is reported in the newspaper functions, in this scenario, as higher-order evidence straightforwardly countering the conclusion drawn by Jim, that there was an error within Newton's theory. That's because what is reported in the newspaper presents Jim with a better line of reasoning, or (some might say) with the proper way of reasoning about that matter. It presents him with a relevant possibility he didn't engage with before, and, in this way, it indicates that his former reasoning might have been inadequate.

The question now is: is it still rational for Jim to be confident of this conclusion to the same extent as before? Nothing in his own technical skills or expertise has changed (he is as competent as before), and nothing in the *modus tollens* itself has changed (it is still in accordance with the rules of inference). If Wedgwood was right in that an inference being competent sufficed for the rationality of endorsing its conclusion, then it would still be rational for Jim to stick to the conclusion of his reasoning, that Newton's theory of universal gravitation is not correct, as much as before. If, however, he also believes what he reads in the scientific newspaper, then he ends up with a belief of the form "Newton's theory of universal gravitation is not empirically adequate, but the evidence shows it is", which is a typical level-splitting response. It is not difficult to envisage how much the progress of science would be hindered if everyone in the scientific community reasoned in this way. No one would ever revise their beliefs – everyone would basically just add new beliefs to their belief systems, but they would never abandon false beliefs formed before. Would that not undermine the very competency of those agents, in the long run?

As I've just mentioned, this response – "Newton's theory of universal gravitation is not empirically adequate, but the evidence shows it is" – is a typical level-splitting response to the total evidence. For Wedgwood's prognosis to stand, this response has to be better than (meaning, more rational than) alternative responses to total evidence, such as the conciliatory response (to suspend judgment, or to decrease confidence in one's former conclusion, in response to higher-order evidence indicating that that conclusion might be false). For the question of whether an attitude is rational is not so much the question of whether it is rational hypothetically, in isolation, or abstractly, but rather whether it is the most rational attitude available to the agent in the situation. But it is not at all clear how believing that "Newton's theory is not empirically adequate, but the evidence shows it is" could possibly be more rational for an astronomer in Jim's position, than to no longer trust (or no longer trust to the same extent as before) his conclusion that Newton's theory is not empirically adequate, like we expect scientists to do in this type of situation. In addition to the level-splitting response being boldly counter-intuitive, one who thinks that this is the most rational response a scientist could have is, again, one who thinks the entire history of science is the history of irrational subjects who opted to revise their beliefs in light of new evidence, instead of sticking to them *and* at the same time to the conflicting idea that those beliefs were not what the evidence supports[6].

[6] Some might have wondered whether our analysis of NEPTUNE hinges on the assumption that Newton's theory of universal gravitation *is* empirically adequate. It does not. Contrast NEPTUNE with the case of Mercury's perihelion: for a long time, astronomers believed it to be evidence of a planet closer to the Sun, Vulcan, causing disruptions in Mercury's orbit, as in Neptune's. However,

4. The type-of-agent claim

In fact, Wedgwood anticipated this type of objection I've been discussing – that a level-splitting response to total evidence might be acceptable, of a rational agent, when taken hypothetically, in isolation, or abstractly; but might not be the best response available. He considered that upon addressing cases of computational intractability and cases of arguments with implausible conclusions[7].. But what he argues does not seem to fully solve the issue. He says, in reply to this objection:

> I am trying to characterize what counts as *perfectly rational* thinking. If all your beliefs are *perfectly* rational, then the degrees of confidence that you attached to the premises, prior to your drawing the inference in question, must already be utterly unimpeachable (Wedgwood, 2012, p. 290).

This is the reason why, for him, "(…) a perfectly rational thinker would continue to draw the inference even if she had (misleading) evidence that she was reasoning incompetently, and even if she entertained serious doubts about whether or not she really was perfectly rational" (Wedgwood, 2012, p. 294).

What Wedgwood is failing to acknowledge here is that no degree of confidence in whatever non-analytical proposition can lie beyond the possibility of revision and, hence, of defeat. That there is no such thing as a "perfectly rational thinker" among us, at least not in the sense just

---

over time, no such planet was observed. In this case, a conclusion "Newton's universal law of gravitation is not empirically adequate" would be true. I'm grateful to an anonymous referee from *Princípios* for pointing this out to me.

[7] Cases akin to NEPTUNE have been presented by Harman (1986); and the sort of objection they help building are briefly discussed by Wedgwood in the final pages of his article.

sketched by him: a person whose degrees of confidence in the premises prior to making an inference is "utterly unimpeachable". Given that no particular epistemic agent is perfectly rational in this sense, it is possible for any epistemic agent to have high degrees of confidence in false propositions, as well as to rationally arrive at false conclusions; and in these situations more than in any other we expect agents to revise the beliefs in question in light of counter evidence of various kinds, including high-orders, and eventually acknowledge that their former inference has a false conclusion, in spite of having been drawn competently, from what they knew at the time.

To sum up what I have argued for so far: the existence of situations in which the reaching of a conclusion strongly appears to call for a revision of the inference (in spite of its having been drawn competently) rather than for its own endorsement exposes a weakness in Wedgwood's assumption, the assumption that drawing an inference competently suffices to render the endorsement of its conclusion a rational attitude regardless of the circumstances. This assumption might hold for perfectly rational thinkers, but, ironically as it seems, perfectly rational thinkers could never find themselves in a situation like Jim's, from neptune, since their assessment of the premises would be "utterly unimpeachable" – meaning, beyond revision – from the beginning. That is to say, a perfectly rational thinker would reject premise 1 right off the bat, which means they wouldn't adduce it in reasoning, like Jim did. Less-than-perfect rational thinkers, in turn, whose assessments of premises are not utterly unimpeachable, cannot take it for granted that the premises they adduce in reasoning are beyond revision. To those thinkers, therefore, the requirements of rationality are more demanding. Specifically, they may be required to revise their own lines of reasoning, in spite of their being confident that the inferences

underpinning their lines of reasoning were performed competently.

Notwithstanding, it is this very same assumption – the assumption that drawing an inference competently suffices to render the endorsement of its conclusion a rational attitude regardless of the circumstances – that ultimately grounds Wedgwood's embracing of the idea that level-splitting can be rational insofar as the inference in question has been drawn competently, and insofar as the agent in question is perfectly rational. That is, it is due to this assumption that Wedgwood technically has to recommend level-splitting to someone in Jim's position as the most rational response, if he is to remain faithful to his own Proposal 3. This assumption being absent, or rightfully circumscribed to the epistemic lives of perfectly rational thinkers (of which someone like Jim is not an exemplar), Wedgwood loses his prime motivation to embrace level-splitting, or the idea that level-splitting can be the most rational response to one's total evidence in cases of the like – for the set of cases to which such recommendation is meant to apply will be null. After all, the discussion about whether level-splitting can be rational is, or at least is meant to be, as far as we're concerned, a discussion about whether level-splitting can be a rational epistemic attitude for human beings, who are less-than-perfect rational thinkers.

Now, one could be more charitable to Wedgwood by taking it, *prima facie*, that perfectly rational thinkers are out there, or among us. Even so, issues remain. Say Wedgwood's view applies to perfectly rational thinkers. It appears that, in this case, his motivation to accept that level-splitting is the most rational response an agent could have in cases of competent inference followed by higher-order counter-evidence stems from his subscribing to an externalist theory of epistemic justification, of the sorts embraced, for instance, by realiabilism. The core insight of reliabilism is that the status

of a doxastic attitude, such as belief or credence, as justified or unjustified, as thus as rational or irrational, depends more upon whether or not the psychological process of which the attitude stems is a reliable one (where reliability is understood in terms of truth-conduciveness), than upon the sort of evidence the agent may have. According to Wedgwood, "whether or not a mental process (like drawing an inference) is rational depends on the *real* nature of that mental process—not on what the thinker *believes* (perhaps mistakenly) about that process"(Wedgwood, 2012, p. 291, italics original).

That is to say, within reliabilist-like accounts of rationality, an epistemic attitude such as a belief, or a credence, is justified or not (and thus rational or not) depending on whether it was formed by a reliable truth-indicator. Since a perfect rational thinker, as conceived by Wedgwood (someone who applies the rules of inference perfectly and whose assessment of the premises, prior to drawing an inference, are utterly unimpeachable), is one whose cognitive processes are reliable par excellence, such a creature is a reliable truth-indicator par excellence. This agent's inferences, therefore, will be consistently competent; and beliefs obtained through them will always be rational to have.

On reliabilist-like accounts, evidence of whatever kind (higher-order or not) plays a small role, if any, in the rationality of endorsements, once the status of a given belief as reasonable or unreasonable depends more upon whether the belief-forming process from which it stems is *in fact* reliable, than on information which may cast doubt either on its reliability or on the truth of its conclusion. In fact, the conduct recommended by Wedgwood for cases involving competent inference followed by higher-order counter-evidence – level-splitting – resembles the conduct often sanctioned by reliabilists for similar cases. One of such cases

is that of the Clairvoyant discussed by Bonjour (1985, Ch. 3). According to him, from a reliabilist standpoint, it would be rational for someone who was in fact a reliable clairvoyant to hold on to a belief formed on the basis of clairvoyance, even when faced with evidence that strongly suggests that he or she does not possess the gift of clairvoyance, or that the belief in question is actually false. This person may even be convinced that they're not a real clairvoyant and that their belief is actually false; even so, they would be justified in sticking to it, since it was formed through a process that is indeed reliable. The clairvoyant, in this situation, would be responding to the situation by holding a belief such as "***p***, but I shouldn't be confident that ***p***".

Likewise, following Wedgwood, a perfectly rational thinker may be convinced that he is not a perfectly rational thinker, and that a given belief of his is actually false; even so, it would be rational for him to hold on to it. But this does not seem to work out well, and here we come to the second pitfall within Wedgwood's position. Let's set aside for a minute the fact that one can be a reliable truth-indicator and yet perform poorly at one particular time (if even thermometers and smoke detectors, that are reliable truth-indicators par excellence, can be wrong sometimes, why would things be any different for any other reliable truth-indicator?). Even when a reliable truth-indicator is performing correctly, doesn't it seem that the very idea of perfect rational thinking on Wedgwood's terms is more consistent with the so-called steadfast response to cases of competent inference followed by higher-order counter-evidence, rather than the level-splitting response? It's not difficult to see why.

Think like this. If one is a perfectly rational thinker, one whose beliefs are *perfectly* rational in the way Wedgwood has said (the degrees of confidence this person attaches to his beliefs are utterly unimpeachable), then, if he believes that a

given conclusion he has drawn *is* rational, no amount of countering higher-order evidence will ever be enough to convince him that he acted irrationally in endorsing said conclusion. Thus, whenever he is confident that, say, ***p***, he would have no reason to accept that he shouldn't be confident that ***p***. But that equals precisely the steadfast verdict, according to which higher-order evidence should have no impact on first-order beliefs. A perfect rational thinker cannot, it seems, be persuaded that a given belief of his stemmed from what was actually an incompetent process, i.e., he cannot fall prey to misleading higher-order evidence. Even if misleading higher-order evidence is out there, and this person comes to access it, it won't make a difference in terms of what he believes. In this sense, he is just like a *perfectly* clairvoyant entity. Just like a perfect clairvoyant would immediately know that the evidence before her is misleading, and then exonerate it, a perfectly rational thinker would immediately know that the new evidence countering his former reasoning is misleading, precisely because it counters his former reasoning, and then he would just exonerate such evidence, instead of compromising between keeping his first-order conclusion and believing he reasoned incompetently. In fact, it would seem rather irrational of this person if he did such a compromising knowing that he has been right all along.

5. Final Remarks

In this paper, I have argued that, in some situations, reaching a certain conclusion calls for the revision of an inference, rather than for its own endorsement, in spite of the inference having been drawn competently. If I am right, then not all cases of competent inference are cases of rational inference, i.e., not every belief one comes to have as a result of a rational inference will be rational for him to hold. The

example I provided is one whereby a scientist reasons from a set of premises that includes a false premise, and then reaches a false conclusion, in spite of applying the rules of inference correctly. The false conclusion attained is boldly at odds with reasonable assumptions the scientific community agrees on. In this type of situation, contrary to the prognosis offered by Wedgwood, I have argued that the most adequate attitude the agent should have involves questioning the premises and rejecting the conclusion – i.e., be conciliatory – in spite of the inference having been drawn competently. Wedgwood recommends level-splitting as the only rational response an agent could have in cases of competent inference followed by high-order counter-evidence. But, if my example is plausible, and if it is true that not all cases of competent inference are cases of rational inference, then the range of cases to which Wedgwood's recommendation is meant to apply becomes significantly more restricted. That was my first objection to Wedgwood's view.

In addition, I have argued, in my second objection, that if we presuppose a perfectly rational thinker in the way Wedgwood conceives of it, then the very idea of misleading higher-order evidence loses its grip. A perfectly rational thinker, as conceived by Wedgwood, is an agent whose evaluation of the premises being adduced in inference is pristine. Such an agent would have no reason to take higher-order evidence of any sort (misleading or not) as a basis to change the level of confidence he has attached to the premises. Therefore, he would have no reason to change the level of confidence he attaches to his conclusion. What matters for the rationality of his conclusion is that the cognitive process by means of which he connects premises and conclusion – i.e., the inference – is reliable; which means evidence shouldn't impact it. If evidence shouldn't impact it, then no amount of higher-order evidence will ever provide him with reasons to

lower the level of confidence in the rationality of his own inference. That is the same as saying that this agent has no reason to accept high-order evidence, in the first place. But that equals exactly the steadfast response, whereby an agent maintains the same level of confidence in his first-order conclusion and rejects higher-order evidence as having a bearing. With this, the set of cases to which Wedgwood's pro-level-splitting recommendation is meant to apply becomes null.

References

BONJOUR, Lawrence. The Structure of Empirical Knowledge. Cambridge: Harvard University Press, 1985.

CHRISTENSEN, David. Epistemology of Disagreement: the Good News. The Philosophical Review, n. 116, v.2, 2007, p. 187–217. doi: 10.1215/00318108-2006-035

CHRISTENSEN, David. Rational Reflection. Philosophical Perspectives, v. 24, Epistemology, 2010, p. 121-140. doi: 10.1111/j.1520-8583.2010.00187.x

COATES, Allen. Rational Epistemic Akrasia. American Philosophical Quarterly, v. 49, n. 2, 2012, pp. 113-124. doi: 10.2307/23213349

DOUVEN, Igor. Abduction. The Stanford Encyclopedia of Philosophy (Summer 2017 Edition), Edward N. Zalta (ed.), URL = https://plato.stanford.edu/archives/sum2017/entries/abduction/>. [consulted: 20/12/2024].

ELGA, Adam. Reflection and Disagreement. Noûs, v. 41, n. 3, 2007, pp. 478-502. doi: 10.1111/j.1468-0068.2007.00656.x

ELGA, Adam. The puzzle of the unmarked clock and the new rational reflection principle. Philosophical Studies, v. 1, n.164, 2013, p. 127-139. doi: 10.1007/s11098-013-0091-0

FELDMAN, Richard. Respecting the Evidence. Philosophical Perspectives, v. 19, n. 1, 2005, pp. 95-119. doi: 10.1111/j.1520-8583.2005.00055.x

GOLDMAN, Alvin. “What is Justified Belief?” In George Pappas, Justification and Knowledge: New Studies in Epistemology. Boston: D. Reidel, 1979, pp. 1–25. doi: https://doi.org/10.1007/978-94-009-9493-5_1

HARMAN, Gilbert. Change in view: Principles of reasoning. Cambridge: MIT Press, 1986.

HAWTHORNE, John; ISAACS, Yoaav; LASONEN-AARNIO, Maria. The rationality of epistemic akrasia. Philosophical Perspective, v. 1, n. 35, 2015, pp. 206-208. doi: 10.1111/phpe.12144

HOROWITZ, Sophie. Epistemic Akrasia. Noûs, v. 4, n.48, 2014, p. 718–744. doi: 10.1111/nous.12026

JOSHI, Hrishikesh. What’s the matter with Huck Finn? Philosophical Explorations, v. 20. n.1, 2016, 1-18. doi: 10.1080/13869795.2016.1246669

KELLY, Thomas. The Epistemic Significance of Disagreement. In: John Hawthorne and Tamar Szabo Gendler (Eds.) Oxford Studies in Epistemology, Vol. 1. Oxford: Oxford University Press, 2005, pp. 167-196.

KELLY, Thomas. Peer Disagreement and Higher-Order Evidence. In: Richard Feldman and Ted A. Warfield (Eds.), Disagreement. New York: Oxford University Press, 2010, pp. 111–174.

LO GUERCIO, Nicolás. Moderate Epistemic Akrasia. Crítica. v. 5, n. 148, 2018, pp. 69-97. doi: 10.22201/iifs.18704905e.2018.04

ROUSH, Sherrilyn. Epistemic Self-doubt. In: The Stanford Encyclopedia of Philosophy (Winter 2017), Edward N. Zalta (Ed.), 2017. Available at: <https://plato.stanford.edu/archives/win2017/entries/epistemic-self-doubt/>. [consulted: 20/12/2024].

SCHOENFIELD, Miriam. An Accuracy Based Approach to Higher Order Evidence. Philosophy and Phenomenological Research, v. 96, n. 3, 2018, pp. 690-715. doi: 10.1111/phpr.12329

SMITHIES, Declan. Moore's Paradox and the Accessibility of Justification. Philosophy and Phenomenological Research, v. 2, n.85, 2012, pp. 273–300. doi: 10.1111/j.1933-1592.2011.00506.x

WEATHERSON, Brian. Do Judgments Screen Evidence? [Unpublished manuscript], 2010. Available at: <http://brian.weatherson.org/JSE.pdf > [consulted: 20/12/2024].

WEDGWOOD, Ralph. Justified Inference. Synthese, v. 189. n. 2, 2012, pp. 273-295. doi: 10.1007/s11229-011-0012-8

WILLIAMSON, Timothy. Knowledge and its Limits. Oxford: Oxford University Press, 2000.

WILLIAMSON, Timothy. Improbable knowing. In: Evidentialism and its Discontents, Trent Dougherty (Ed.), Oxford: Oxford University Press, 2011, pp. 147–164.

WILLIAMSON, Timothy. Very Improbable Knowing. Erkenn, v. 1, n.79, 2014, pp. 971–999. doi: 10.1007/s10670-013-9590-9

(Submissão: 20/12/24. Aceite: 24/07/25)