



Why reasons and reasoning don't come apart

Christian Kietzmann¹

Received: 25 February 2023 / Accepted: 16 September 2023 / Published online: 2 November 2023
© The Author(s) 2023

Abstract

In recent years several philosophers have proposed what has come to be called the Reasoning View of normative reasons, according to which normative reasons are premises of sound reasoning. The reasoning view has come under some criticism, which chiefly consists in counterexamples that purport to show that something can be a premise of sound reasoning without intuitively being a normative reason, or can be a normative reason without being a premise of sound reasoning. I here consider and reject three examples that were recently put forward. Discussing them will allow me to clarify the Reasoning View in important respects.

In recent years, several philosophers have proposed what has come to be called the “Reasoning View” of normative reasons.¹ On such an account, normative reasons are premises of sound reasoning. Reasoning is sound when its premises are correct and it follows a correct pattern. This is intended to be an explanation of what normative reasons *are*.

The reasoning view has come under some criticism, which chiefly consists in counterexamples that purport to show that the account over- or undergenerates normative reasons. Some purport to show that something can be a premise of sound reasoning without, intuitively, being a normative reason, while others aim to convince us that some genuine normative reasons cannot be premises of sound reasoning.

¹ Some recent proponents are Müller (2004), Setiya (2014), Silverstein (2016), Way (2017), Asarnow (2016 and, 2017), Paakkunainen (2017) and Hieronymi (2021). There is also an older tradition of explaining reasons as premises of reasoning; see e.g. Raz (1975 and, 1978), Williams (1981), Harman (1986) and Grice (2001). For a very similar view, see Korsgaard (2009, p. 28).

✉ Christian Kietzmann
christian.kietzmann@fau.de

¹ Friedrich-Alexander-Universität Erlangen-Nürnberg, Institut für Philosophie, Bismarckstr. 1, 91054 Erlangen, Germany

In this paper, I will address three examples put forward by Eva Schmidt and Artūrs Logins in recent papers.² Responding to them will allow me to bring out some important clarifications of the Reasoning View.

1 Some general discussion of the Reasoning View

I will begin by situating the Reasoning View within the contemporary philosophical landscape. Many philosophers treat reasons as an irreducible and as the most basic normative concept. There are two distinct claims. Reasons Fundamentalism is the claim that the concept of “being a reason for” is basic. It cannot be analysed in simpler terms (see e.g. Scanlon, 1998 and, 2014; Parfit, 2011). Reasons First is the claim that other normative concepts, such as the concepts of value, norm, fittingness, ought or reasoning, can be analysed in terms of reasons (see e.g. Scanlon, 1998).

The Reasoning View rejects both Reasons Fundamentalism and Reasons First. It rejects the former by accounting for the concept of a reason in terms of correct reasoning. For this reason, the Reasoning View is a competitor of other accounts of reasons that have recently been defended: accounts of reasons as explanations for ought facts (Broome, 2013), reasons as explanations of the goodness of an action (Raz, 1999), or reasons as evidence for ought propositions (Kearns & Star, 2008, 2009).

Notice, however, that the Reasoning View need not be understood as a reductive project. It can equally well be put forward in the spirit of a connective analysis of a web of interrelated and mutually irreducible concepts.³ The Reasoning View would still be incompatible with Reasons Fundamentalism, because reasoning and reasons would be understood as necessarily interrelated and mutually illuminating concepts. The Reasoning View would also reject Reasons First to the extent that the concept of reasoning would not be explainable in terms of reasons.

Something weaker than the Reasoning View is the *Reasoning Constraint* on reasons. According to this constraint, whenever an agent *A* has a normative reason *p* to *F* in circumstances *C*, it is in principle possible for him to reason from *p* to the intention or intentional action to *F*.⁴ The reasoning constraint is weaker because it does not explain what a reason is in terms of correct reasoning. Nevertheless, it holds that the possibility of reasoning with a reason is a necessary condition that every candidate for a normative reason must meet.

Now, what is at stake with the alleged counterexamples that Schmidt and Logins put forward? Let’s look first at the examples in which the Reasoning View seems to overgenerate reasons. These examples suggest that flawless reasoning doesn’t reliably track normative reasons, because not all reasons that appear in flawless reasoning are intuitively normative reasons. You can do everything correctly in your reasoning and still not have genuine normative reasons in view. This is bad news not only for the Reasoning View but for reasoning itself. For how, then, can correct reasoning ever hit upon genuine reasons? Why isn’t it, conceptually speaking, a mere

² See Schmidt (2021) and Logins (2020).

³ For the idea of connective analysis, see Strawson (1992).

⁴ Compare e.g. Schroeder (2007); Paakkunainen (2017); Kietzmann (2022a).

accident if reasoning manages to track reasons? Moreover, if we cannot rely on reasoning to track reasons, one wonders: what good is reasoning anyway? Why engage in reasoning if even our very best efforts do not lead to genuine reasons?⁵

The examples that purport to show that the Reasoning View undergenerates reasons pose a different problem. If successful, they undermine not only the Reasoning View, but also the weaker Reasoning Constraint. In other words, the counterexamples threaten to cut normative reasons loose from reasoning and deliberation. They suggest that playing a role in reasoning is, at best, an additional fact about reasons, i.e. a fact that is not entailed by the very concept of a normative reason. They suggest that it is, conceptually speaking, an accident if reasons play this additional role.

A *prima facie* argument against divorcing reasons from reasoning in this manner is that it threatens to undermine the distinction between genuine normativity and merely evaluative normativity. For if there are reasons that in principle cannot be taken up in reasoning, it is not clear in what sense such reasons still genuinely count in favour of the actions they are reasons for. How is such favouring different, for instance, from the sunshine favouring a sun-flower's turning its head towards it, or a gazelle's proximity favouring the lion's pursuit of it? Of course, the sunshine and the gazelle's proximity are factors we bystanders must take into account when we want to evaluate the appropriateness of the sunflower's or the lion's behaviour. But is the sunshine a genuine normative reason for the sunflower to turn its head, and is the proximity of the gazelle a real normative reason for the lion to chase it? Intuitively this seems strained, presumably since sunflowers and lions do not understand reasons as reasons, and so aren't capable of treating them as reasons in reasoning. They quite simply lack the necessary cognitive and reasoning capacities. But how, then, are things different for us humans with respect to those reasons which we cannot, in principle, accommodate in our thinking and reasoning? I will return to such worries in Sect. 4 of this paper.

2 Schmidt's wedding anniversary example

Let's now turn to the counterexamples. A particularly challenging one was recently put forward by Eva Schmidt:

⁵ One might object, for instance, that we are justified to adopt reasoning patterns that don't fully align with our reasons on pragmatic grounds, e.g. because reasoning patterns that fully match our reasons would be excessively cumbersome and not workable in practice. (Compare e.g. Kahneman, 2011, on heuristics we use in reasoning, e.g. in order to be able to react quickly to challenging situations.) Of course, such reasoning patterns would still have to secure *some* measure of alignment with the reasons we actually have. But the fit need not be perfect. In response, I would like to point out that we only get a meaningful contrast between reasoning heuristics that don't perfectly track reasons and these reasons themselves when we are at least in principle able to get a grip on the reasons we actually have through reasoning, however cumbersome it may be. Otherwise, we cannot determine whether and to what degree reasoning heuristics actually *deviate* from the reasons we have. But if this is so, then the objection merely shows that we need to distinguish between imperfect but handy reasoning heuristics and those reasoning patterns that fully match our reasons.

It's Peter and his husband Bob's wedding anniversary tomorrow. Every year, Peter gives Bob ten red roses as an anniversary gift. Otherwise, neither of them has any interest in flowers, but this is their wedding anniversary tradition. In the evening, Peter goes out and gets extremely drunk. When he wakes up hungover the next morning, he has forgotten about the anniversary, but suddenly it hits him: 'I'm going to go to the florist's today!' I.e., he finds himself intending to go to the florist's and so comes to believe that he intends to go to the florist's. Via inference to the best explanation, he then correctly infers from the fact that he intends to go to the florist's to the true proposition that today is his and Bob's wedding anniversary. His belief that today is their wedding anniversary, together with his standing intention to give Bob ten red roses for their wedding anniversary every year, moves him to form the intention to give Bob ten red roses today. He executes this intention by going to the florist's, buying and then giving ten red roses to his husband – his reasoning terminates in this action. (Schmidt, 2021, p. 766)

Peter reasons abductively from the fact that he intends to see the florist to the fact that today is his and Bob's wedding anniversary, and he reasons instrumentally from that fact and the intention to present Bob with some roses on anniversaries to the intention to give him ten roses today. His reasoning seems sound. Sound reasoning gets Peter from the fact that he intends to go to the florist's today to the intention to give Bob ten roses. The Reasoning View, Schmidt suggests, is therefore committed to claiming that the fact that he intends to go to the florist's is a normative practical reason for Peter to give Bob ten red roses. But intuitively this seems wrong. The fact that Peter intends to go to the florist's provides him at best with a theoretical reason for thinking that today is his wedding anniversary, and the fact that today is his wedding anniversary is in turn a practical reason for him to intend to give Bob the roses. But the theoretical reason is not in itself a practical reason for having the intention.

In response to this challenge, I propose two moves on behalf of the Reasoning View. First, we should distinguish between trains of reasoning and reasoning patterns, and second, we should distinguish between theoretical and practical reasoning. I will sketch two ways these distinctions can be made more precise. Note that, even if it turns out that the particular way in which I spell out these distinctions needs to be improved, my two-step strategy may still work against the counterexample.

The first move is a clarification of what the best version of the Reasoning View ought to claim. Its proponents ought to distinguish between trains of reasoning and reasoning patterns. A train of reasoning is an exercise of a person's reasoning capacities that occurs in space and time.⁶ This is a psychological notion. A pattern of reasoning, by contrast, is a logical notion. It denotes a relation among propositions or attitudes, and thus something that does not occur in space and time. It is *patterns of reasoning* that are correct or incorrect. By contrast, it is *trains of reasoning* by which

⁶ In the current literature, it is common to encounter talk about "processes of reasoning". However, I am sceptical that "process" is the right ontological category for episodes of reasoning. See Kietzmann (2022b). I therefore prefer the phrase "train of reasoning".

particular persons adopt conclusion attitudes on account of premise attitudes.⁷ We can now specify that the Reasoning View is about patterns of reasoning, not about trains of reasoning. It claims that normative reasons are premises in sound patterns of reasoning.

The crucial question concerning Schmidt's counterexample is thus not whether Peter engages in a single train of reasoning. Perhaps he does, but that is irrelevant as far as the Reasoning View is concerned. Rather, the crucial question is whether his reasoning exemplifies a single pattern. In what follows, I will argue that it doesn't. In my view, the example consists of two independent parts which instantiate distinct patterns: a pattern of abductive theoretical reasoning and a pattern of instrumental practical reasoning. The theoretical part of Peter's reasoning concludes in his belief that today is his and Bob's wedding anniversary, whereas the practical part terminates in his intention to give Bob ten red roses today.

How do we determine whether one single or two different patterns are involved here? What we need is a robust distinction between theoretical and practical reasoning. A first proposal for drawing the distinction is in terms of their respective conclusions: theoretical reasoning terminates in belief, disbelief, or suspension of judgement, whereas practical reasoning terminates in intention or intentional action.⁸ However, in the counterexample we are considering, reasoning terminates in an intention. So if we went by the first proposal we are considering, as a whole it would count as a piece of practical reasoning, even though it contains a stretch of reasoning that terminates in belief and is thus theoretical.

I would like to make a different proposal for drawing the distinction between theoretical and practical reasoning, which focuses on their different *constitutive aims*.⁹ The idea of a constitutive aim is familiar from the theory of action. According to J. David Velleman, a constitutive aim of action is "an aim with respect to which behavior must be somehow regulated in order to qualify as action" (2004, p. 234). Similarly, I want to suggest that theoretical and practical reasoning qualify as such because they are regulated by different aims. Let's suppose that theoretical reasoning patterns constitutively aim at true belief, whereas practical reasoning patterns constitutively aim at good action.¹⁰ Then we can say that theoretical reasoning patterns are *theoretical* because they are regulated by the former aim, whereas practical reasoning patterns are *practical* because they are regulated by the latter aim. The concept of a constitutive aim is often illustrated with the game of chess. For example, Paul

⁷ What I call "patterns of reasoning" are sometimes called "arguments" and what I call "trains of reasoning" "inferences"; see e.g. Harman (1986, pp. 3–4). Thanks to an anonymous reviewer for pointing this out.

⁸ See Schmidt (2021, pp. 766–767 and 773), where she also considers several other suggestions for drawing the distinction, which I will disregard in what follows.

⁹ Schmidt (2021, p. 773) mentions this possibility and cites Silverstein (2016, pp. 6–8) as having made this proposal.

¹⁰ The precise nature of the constitutive aims of theoretical and practical reasoning is of course contentious. For example, some argue that the constitutive aim of theoretical reasoning is knowledge, and others surmise that the constitutive aim of practical reasoning is self-knowledge, or autonomy, or efficient action. I will ignore these issues because, for the purposes of my argument, nothing hangs on whether I get the precise nature of these constitutive aims exactly right.

Katsafanas points out that “part of what it is to play chess is to aim at checkmating your opponent. This aim simply must be present in order for a series of movements to count as an episode of chess-playing” (2013, p. 1). Aiming at checkmating is the aim that makes chess the kind of game it is. By contrast, spending an entertaining evening, or humiliating your opponent, or pocketing the prize money, are further or external aims you may pursue in playing a particular game of chess. They are external aims in the sense that whether or not you pursue them makes no difference to the question whether or not you’re playing chess.

Does one invariably pursue the constitutive aim of a reasoning pattern whenever one instantiates the pattern? Consider that there is a sense in which one need not intend to checkmate one’s opponent when playing chess. Perhaps you are playing with your boss, who can’t bear to lose at chess, and so you intend to lose in order to please him. Or you are teaching your daughter how to play and so you intend to let her win the game. This is possible of course. But it doesn’t undermine the point that, even in those cases when you don’t intend to win the game, there is still a sense in which the point of your moves remains that of checkmating your opponent. You’re just intentionally playing badly so as to lose. First, you still have to make moves according to the rules, and it is one of the rules of chess that winning consists in checkmating. And second, your opponent may protest if they realise that you are intentionally letting them win: they may, for instance, say “but you are not *really* playing!” So even if you intend to lose, you still at least have to keep up the appearance of trying to checkmate.

Reasoning patterns are abstract things. Can *abstracta* have aims, let alone constitutive aims? I think they can. Consider, for example, a recipe in a cookbook. A recipe is something abstract, a pattern that can be instantiated by individual actions of cooking. Yet it seems intuitively intelligible to say that recipes have an aim or point: they aim at a well-cooked meal of the type the recipe is a recipe for. This arguably entails that every act of cooking that uses the recipe is towards that aim. Similarly, I want to claim, reasoning patterns have aims, where this entails that the patterns can be instantiated by individual actions of reasoning, and that these actions individually aim at true belief or good action.

My point about reasoning patterns having constitutive aims allows us to respond to Schmidt’s example. The part of Peter’s train of reasoning that gets him from his intention to go to the florist’s to the thought that today is his wedding anniversary is *theoretical* because it instantiates a pattern the point of which is to establish a true belief. The relevant question that reasoning according to this pattern helps answer is whether today is Peter and Bob’s wedding anniversary, which is a factual question. The reasoning thus constitutively aims at finding out how things are. Whether or not engaging in that piece of reasoning *also* serves a practical purpose, such as figuring out how to celebrate the wedding anniversary, is neither here nor there where the reasoning’s constitutive aim is concerned. If the reasoning *also* has such a practical aim, this is an *additional* aim that is *external* to the reasoning itself. On the other hand, the bit of reasoning that gets Peter from his insight that today is his wedding anniversary to the intention to give Bob ten red roses today is *practical* because it instantiates a pattern whose point is to secure the goodness of Peter’s action. The relevant question it helps answer is whether he is acting well in giving Bob the roses today.

Schmidt is right to insist that any view that denies that her example instantiates a single correct pattern of reasoning should not cut such patterns too finely. In particular, it should allow for cases where a single pattern consists of several steps, as it often does in instrumental reasoning. If you reason from a good end A to a sufficient means B, and from B in turn to a sub-means C, it is generally correct to think of A as a reason not only for B but also for C. So the Reasoning View should think of the whole as a single train of reasoning which instantiates a single correct pattern.

My proposal to distinguish between theoretical and practical reasoning in terms of constitutive aims allows us to preserve this important thought. For we can allow that the premises of sound reasoning are reasons for that which figures in its conclusion, whether the reasoning consists of one or several steps, as long as *all the steps are of one kind*, i.e. all instantiate either theoretical or practical sub-patterns. Thus we can say that, when you reason from end A to means B to sub-means C, this is a single train of practical reasoning with two practical parts, whereby A is a reason for B, but also for C. We can also say that, when you reason from fact p to end D, and from end D to means E, this is a single train of practical reasoning with two practical parts, such that p is a normative reason for both D and E. And we can say that, when you reason from evidence p to proposition q, and from evidence q to proposition r, this is a single train of theoretical reasoning, whereby p is evidence for both q and r. But what we *cannot* say is that, when you reason from evidence p to proposition q, from proposition q to end F and from end F to means G, p is both a theoretical reason for q *and* a normative practical reason for doing F and doing G. Why? Because here the train of reasoning consists of parts that are not of a single kind, and therefore the whole train does not instantiate one single reasoning pattern.

Schmidt considers a version of this reply and objects that theoretical and practical reasoning are often intertwined, as when you simultaneously figure out appropriate means and possible unintended side-effects:

To borrow an example from Bratman (2009, p. 42), imagine that a combat pilot keen on helping to win the war correctly reasons towards the conclusion that she is going to bomb the munitions factory and that she will thereby promote the war effort, but also destroy the nearby school. Taking seriously the distinction between the intended goal of an action and its unintended side-effects, the combat pilot's conclusion involves both an intention (to bomb the munitions factory and thereby to promote the war effort) and a belief about the unintended consequences of her action (that she will thereby destroy the nearby school). Correspondingly, this reasoning has both a practical and theoretical aim, both hitting upon the best course of action and figuring out the truth about its side-effects. But then the reasoning is both practical and theoretical. (Schmidt, 2021, p. 774)

Schmidt means to say that in the case she considers, numerically one and the same reasoning concludes in two things, an intention to bring about intended effects and a belief about unintended side-effects, and has both a practical and a theoretical aim. She concludes that numerically one and the same reasoning is both practical and theoretical.

However, this conclusion follows only if one overlooks that the proposal Schmidt objects to distinguishes between theoretical and practical reasoning not simply in terms of aims of reasoning, but in terms of *constitutive* aims. In the bomber pilot's means-end reasoning, the pilot's *constitutive aim* is practical: she aims to figure out what to do in the situation at hand. Given her aim to support the war effort, she settles upon the appropriate means of bombing the munitions factory. However, acting well requires not only adopting appropriate means to one's ends, but also considering possible side-effects, and negotiating between envisaged means and their unintended effects. Reasoning that uncovers effects – “if I do this, such and such is likely to happen...” – is theoretical reasoning, because its constitutive aim is true belief. This fact is in no way changed or undermined if the theoretical reasoning is embedded in practical reasoning, and thus serves the practical end of securing the agent's acting well. For this practical aim is not a *constitutive*, but only an *external* aim of that reasoning. In other words, although theoretical reasoning is here in the service of practice, as we might say, this doesn't make it practical reasoning, because the reasoning retains its constitutive aim of true belief.¹¹ For this reason I disagree with Schmidt's conclusion that one and the same reasoning is both theoretical *and* practical. Some stretch of the bomber pilot's reasoning is theoretical because it instantiates a pattern whose constitutive aim is true belief. It concludes in beliefs about effects. Other stretches are practical because they instantiate patterns with the constitutive aim of acting well. They conclude in intentions to bring about certain (intended) effects.

I thus submit that the distinction between intended effects and unintended yet foreseen effects is bound up with and explained by the different kinds of reasoning that lead up to the beliefs and intentions in which these different effects figure. Schmidt agrees that the distinction itself is important and should be respected for other reasons. This distinction is important not only in order to get an adequate grip on our own psychologies, but also for moral and legal reasons: we treat consequences differently depending on whether they were intended or merely foreseen. Moreover, the distinction should not be ad hoc or arbitrary. There should be some principled ground for drawing it. My proposal gives us a principle for making the distinction because it holds that intentions to bring about some effect are the conclusions of reasoning that instantiates patterns whose constitutive aim is practical, whereas beliefs about effects are conclusions of reasoning that instantiates patterns whose constitutive aim is theoretical. By contrast, if one asserts, as Schmidt does, that in some cases theoretical and practical reasoning are inextricably intertwined, and that one and the same reasoning issues in beliefs about side-effects and in intentions, it remains somewhat unclear on what grounds it is still possible to distinguish between intended effects and unintended effects that are merely foreseen. Indeed, if one assumes that there is a single reasoning pattern at work, it is inviting to conclude that there is no deep difference here and that all effects of one's actions that one is aware of are intended effects.¹²

¹¹ For this important point, compare Müller (1979 and, 1992).

¹² Some such reasoning led authors in the consequentialist tradition to completely abandon the distinction. They thought, roughly, that reasoning here consists in determining which effects our actions will have, and how these effects will promote some designated overall end. See e.g. Sidgwick (1907, III.1, pp. 201–202), and Hare (1952, pp. 56–57).

3 Schmidt's white Christmas example

Schmidt offers a second counterexample in which the Reasoning View seems to over-generate reasons:

Reliable weather forecast website 1 predicts that there will be a white Christmas, whereas reliable weather forecast website 2 predicts that there will be no white Christmas. In response to these conflicting forecasts, Haya believes, first, that website 1 predicts that there will be a white Christmas and she believes, second, that website 2 predicts that there will be no white Christmas. Reflecting on her beliefs leaves her unsettled on the issue – i.e. it terminates in her suspending on whether there will be a white Christmas. (2021, p. 769)

Haya reasons from two conflicting pieces of reliable evidence – the weather forecast of website 1 and the weather forecast of website 2 – to adopting an attitude of suspension of judgement.¹³ The reasoning pattern she thereby instantiates seems correct: if the available evidence on an issue conflicts, the right thing to do is to suspend judgement. On the Reasoning View, a reason is a correct premise in a correct pattern of reasoning. So, Schmidt argues, on the Reasoning View, since each of the conflicting pieces of evidence is a premise in the reasoning pattern underlying Haya's inference, each of them, *taken on their own*, counts as a reason for suspending judgement. But intuitively this seems wrong. A reliable piece of evidence – like, say, the weather forecast of website 1 – is not a reason to suspend judgement.¹⁴

I'd like to begin with a methodological remark. Schmidt suggests that the best kind of response to the one counterexample should also work for the other (p. 774). I disagree. In my view, the problems with this example are different from the difficulty with the example I considered in the previous section. Different problems call for different solutions.

I'll now turn to the solution. In response to the counterexample, I recommend refining the Reasoning View.¹⁵ The example shows that premises of correct reasoning aren't normative reasons when taken in isolation. When we talk of reasons, we're referring to the whole set of premises of a correct piece of reasoning. Even if we only name a single reason, what we're actually referring to when we call something a reason is the whole argument in favour of the conclusion. In such cases, normally the context makes clear what the rest of the argument looks like, and we name only a single reason because we are not interested in those parts of the argument that we do not explicitly state.

How does this refined version of the Reasoning View apply to Schmidt's example? The weather forecast of website 1 is a reason to suspend judgement on whether there will be a white Christmas only *in the context of* conflicting and equally strong

¹³ Following Friedman (2013; 2017), Schmidt takes suspension of judgement to be a question-directed attitude (of being unsettled with respect to a question). See Schmidt (2021, p. 769).

¹⁴ See Schmidt (2021, pp. 770–771).

¹⁵ A similar proposal has been made independently by Asarnow (2016 and, 2022). Thanks to an anonymous reviewer for drawing my attention to Asarnow's work.

counter-evidence. This much seems uncontroversial on any reasonable view of evidential support. The Reasoning View, if read in the way I suggest, captures this point nicely. For it says that weather forecast 1 is a reason to suspend judgement only if it is a premise of a reasoning pattern that also contains another premise containing equally strong counter-evidence – and no further premises that tilt the scales one way or another.

4 Logins on Moore paradoxical reasons

Artūrs Logins puts forward the following challenge to the Reasoning View (RV), in a version proposed by Conor McHugh and Jonathan Way that gives prominence to the fittingness of attitudes:¹⁶

The fact that (*r*) “the building is on fire, but John doesn’t believe that the building is on fire” is, intuitively, a reason for John to check/consider/reconsider/investigate the hypothesis (*h*) “the building is on fire.” Yet, there is no possible good pattern of reasoning for John from a fitting belief in (*r*) to reconsidering/investigation/etc. of (*h*). This is so simply because it is not possible for John to have a fitting belief that the building is on fire and that he doesn’t believe that the building is on fire. John cannot have a fitting belief that (*r*) because it is impossible for such a belief to be true. One cannot truly believe that the building is on fire and that one doesn’t believe that the building is on fire. The belief in the first conjunct contradicts the belief in the second. On the (RV) account, given that John cannot (in any sense) have a true belief in (*r*), (*r*) cannot possibly be a reason for any response from him. However, it appears pretheoretically plausible to think that (*r*) speaks in favor of some response for John. (2020, pp. 709–710)

A crucial assumption in this argument is that John cannot believe that (*r*). In what follows, I will argue that this assumption is wrong.

Notice that Logins, in presenting his argument, switches back and forth between two different versions of (*r*). On the face of it, (*r*) looks like an instance of Moore’s paradox. It is well known, however, that there is an important difference between first-person and third-person versions of Moore paradoxical statements. Consider the assertion that (*r**) “the building is on fire, but John doesn’t believe that the building is on fire.” This isn’t paradoxical at all: John might just be ignorant about the fire. In fact, it is not uncommon for something to be the case and for John not to believe that it is. So it is quite possible for people to truly believe that (*r**), and to reason from this belief to various conclusions. The situation doesn’t change fundamentally when the belief that gets expressed in the assertion (*r**) is *John’s own*. For John might be ignorant that *he himself* is the same John whom the second part of his assertion is about. Where this is the case, it is still possible for John to believe that (*r**) and to use that

¹⁶ See McHugh and Way (2016).

belief – or its content – as a premise in his reasoning.¹⁷ The assertion becomes paradoxical only when we switch, within (*r*), from the name “John” to the first-person pronoun: (*r****) “the building is on fire, but *I* don’t believe that the building is on fire”. Quite plausibly, it is not possible for John to truly believe that (*r****) . And thus (*r***) cannot serve as a premise in John’s reasoning.

However, if there is any version of (*r*) that favours checking/reconsidering (*h*), it must be (*r**), the third-person version, and not (*r****) , the first-person version. And as I have argued, it is possible for John to believe that (*r**) and reason from it. What is not possible, by contrast, is for him to believe and reason from (*r****) , i.e. the first-person version of (*r*). But intuitively, (*r****) doesn’t favour checking/reconsidering etc. (*h*), because (*r****) is Moore paradoxical, a contradiction of sorts, and thus doesn’t favour anything.

One might object that what counts in favour of actions in the sense of reasons are facts. Thus, reasons should be individuated as finely as facts. In this case, (*r**) and (*r***) are *the same* fact, for in general a fact doesn’t change if its subject is represented differently. For example, the fact that John enters the room and the fact that last year’s winner of the Ryle Prize enters the room are one and the same fact, given that John was last year’s winner of the Ryle Prize. So, the objection concludes, facts favour certain responses quite independently of how they are represented. If the fact that John enters is a good reason to, say, leave the room, *ipso facto* the fact that last year’s winner of the Ryle Prize enters is a good reason to leave the room.

However, I doubt that this is true. In my opinion, there are good reasons to individuate reasons more finely than facts as they are understood in the objection.¹⁸ For the favouring relation does seem to be sensitive to how facts are represented. An example brings this out. Assume that the fact that Peter runs into his wife is a reason for him to greet her with a kiss. By contrast, the fact that Peter runs into the Russell Professor of Natural Philosophy isn’t a reason for him to greet her with a kiss. Intuitively, this is still the case even if Peter’s wife happens to be the Russell Professor of Natural Philosophy, because the reason for kissing her is tied up with her being his wife, not with her being the Russell Professor of Philosophy.¹⁹

Even if I’m wrong on this point and there’s only one single fact that (*r*), of which (*r**) and (*r***) are merely different descriptions, and that (*r*) is John’s reason however it is described, what I have argued above still undermines Logins’s challenge. For he assumes that “[o]ne cannot truly believe that the building is on fire and that one

¹⁷ But can John believe this *truly*? If the first conjunct expresses one of John’s beliefs, this seems to guarantee that the second conjunct is false, and so can’t be truly believed. Notice, however, that in certain (admittedly rare) scenarios true beliefs of this sort are in fact possible. Perhaps John is in two minds about the question whether the house is on fire, part of him believing it and part of him repressing the belief and in this sense not believing it. If the first conjunct of (*r**) expresses the non-repressed belief, and the second conjunct describes the repressed state of non-belief, John can in fact truly believe (*r**) about himself.

¹⁸ Let me emphasise, however, that the Reasoning View is not as such committed to this view of the individuation of reasons. For example, one might claim instead that reasons are facts, and that fact F is a reason for A to phi iff F has at least one guise that satisfies the Reasoning View conditions. Thanks to an anonymous reviewer for pointing this out.

¹⁹ Consider also that the relation of explaining something seems in many ways to be similar to the relation of being a reason for something. But it is rather uncontroversial that explanation creates intensional contexts. For explanation, it matters how *explanans* and *explanandum* are represented.

doesn't believe that the building is on fire. The belief in the first conjunct contradicts the belief in the second" (Logins, 2020, p. 709). But as we have seen, this is not true. My argument shows that there is a way for John to truly believe that (r) – namely in the guise of (r^*). True, he can't believe that (r) in the guise of (r^{**}). But if a belief that (r^*) and a belief that (r^{**}) are merely different ways of representing one and the same fact – namely, that (r) – then the fact that at least *one* of these beliefs is available to John is enough to refute Logins's assumption.

At this point, someone could object that Logins's example is just one among a much wider class of counterexamples. Elusive reasons, as they have come to be called, apply to an agent even though the agent cannot truly believe them.²⁰ Perhaps the most extensively discussed elusive reason is Mark Schroeder's surprise birthday party:

[C]onsider the case of Nate, who hates all parties except for successful surprise parties thrown in his honor. Given Nate's situation, the fact that there is a surprise party waiting for him now at home is a reason for him to go home. But it isn't a reason that Nate could know about or act on. Still, someone Nate trusts might tell him that there is a reason for him to go home now. (Schroeder, 2007, p. 33)²¹

Several philosophers have offered ingenious responses to this and other examples.²² I won't discuss them here. Instead, I would like to offer a principled objection to the possibility of elusive reasons.

The objection, in a nutshell, is that elusive reasons go against a combination of two seemingly platitudinous assumptions: that normative reasons are normative, and that normativity is constrained by what is possible for its addressees – in other words, that ought implies can. Here are more precise formulations of these assumptions:

- (1) If p is a reason for A to F for the reason that p , then A *pro tanto* ought to F for the reason that p .
- (2) If A *pro tanto* ought to F for the reason that p , then A can F for the reason that p .

Assumption (1) expresses the normativity of normative reasons. One may suggest replacing (1) with a weaker version that is formulated not in terms of "ought", but in terms of a weaker normative modality – such as "should". However, "ought" is the default expression of normativity. And so in order to retain a connection between having a reason and normativity, there should be a connection between what one has reason to do and what one ought to do. Or one may suggest weakening (1) from say-

²⁰ Compare Rossi (2021, p. 83): "a blindspot reason p is a normative reason for A to ϕ that is such that there is no possible world in which (a) A believes that p and (b) p is true".

²¹ For other examples, compare Markovits (2014) and Rossi (2021). Such examples are usually put forward in order to undermine motivational and deliberative constraints on normative reasons. But they seem equally well to speak against the Reasoning View.

²² See e.g. Paakkunainen (2017), and McKeever and Ridge (2012).

ing that having a reason to do something entails that one *pro tanto* ought to F for the reason that p, to saying simply that one *pro tanto* ought to F – for whatever reason. Notice, however, that doing F for whatever reason would not be a genuine response to the reason that one has. One would merely by chance do what one has reason to do. And intuitively this is not what having a reason demands of one. A third suggestion is to weaken (1) and (2) in order to accommodate *prima facie* reasons or enticing reasons. The idea is that, while these reasons are genuine reasons, you don't necessarily flout normative demands – *pro tanto* oughts – if you don't go along with them. Conversely, this would mean that having such a reason doesn't imply that you ought to comply with it. This is certainly true. But, intuitively, even a *prima facie* reason or an enticing reason confronts you with *some* kind of normative demand. Perhaps it's not the case that you *pro tanto* ought to do what it recommends to you. But you nevertheless cannot completely ignore it, normatively speaking. And this in turn must have something to do with what you *pro tanto* ought to do. So even if it's true that no *pro tanto* ought follows directly from a *prima facie* or an enticing reason, such reasons must nevertheless be indirectly tied to *pro tanto* oughts.

Turn now to (2). This is a widely accepted truth about normativity.²³ “Ought implies can” says, intuitively, that the normative demands on us cannot outstrip what we in principle can do. Conversely, *ultra posse nemo obligatur* – i.e. beyond what is possible, nobody is obliged. If it's in principle not the case that you can do something, then it's not the case that you ought to do it. The rationale behind this is obvious enough: demands that you cannot live up to are both unfair and pointless. They are unfair because complying with them is not up to you. You thus incur criticism or blame for not complying with the demand no matter what you do. They are pointless because a demand on you that you cannot in principle meet cannot make any difference to you. So what would the demand be good for?

If we accept (1) and (2) as they stand, it follows that

(3) if p is a reason for A to F for the reason that p, then A can F for the reason that p.

But (3) is in tension with the existence of elusive reasons. They are reasons for A to F, for which A cannot – and in principle cannot – act. And if it's not the case that A can F for the reason that p, by *modus tollens* and (3), it is not the case that p is a reason for A to F – contrary to what the proponents of elusive reasons assume.

5 Conclusion

I have argued that several counterexamples that have recently been put forward against the Reasoning View of reasons fail in interesting ways. Two lessons have emerged. The first is that such counterexamples threaten to undermine more than the Reasoning View. If valid, these examples also put pressure on the Reasoning

²³ Streumer (2007) derives “ought implies can” from the idea that, if one has a reason to F, then it must in principle be possible for the person to F.

Constraint on reasons, and thereby put into question core intuitions about reasons and reasoning, namely, that reasons register in reasoning and that reasoning tracks reasons. The second lesson that has emerged from the discussion is that the Reasoning View should be refined in several ways: it should distinguish between practical and theoretical reasoning – perhaps in the way sketched here; it should conceive of reasons as premises within a whole pattern of reasoning, and not in isolation; and it should be sensitive to modes of presentation used in the premises.²⁴

Funding Open Access funding enabled and organized by Projekt DEAL.

Declarations

Conflicts of interest This work was supported by Deutsche Forschungsgemeinschaft (Grant numbers 391136454 and 439616221). The author has no competing interests to declare that are relevant to the content of this article.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Asarnow, S. (2016). *Rational Internalism Ethics* 127:1, 147–178.
- Asarnow, S. (2017). The reasoning view and defeasible practical reasoning. *Philosophy and Phenomenological Research*, 95:3, 614–636.
- Asarnow, S. (2022). Action and rationalization. *Australasian Journal of Philosophy*, 100:4, 758–773.
- Bratman, M. (2009). Intention, belief, practical, theoretical. In S. Robertson (Ed.), *Spheres of reason: New Essays in the philosophy of Normativity* (pp. 29–61). Oxford University Press.
- Broome, J. (2013). *Rationality through reasoning*. Oxford University Press.
- Friedman, J. (2013). Question-directed attitudes. *Philosophical Perspectives*, 27, 145–174.
- Friedman, J. (2017). Why suspend judging? *Noûs* 51, 302–326.
- Grice, H. P. (2001). *Aspects of Reason*, edited by Richard Warner, Oxford: Oxford University Press.
- Hare, R. M. (1952). *The Language of Morals*. Clarendon Press.
- Harman, G. (1986). *Change in View: Principles of reasoning*. MIT Press.
- Hieronymi, P. (2021). Reasoning First. In R. Chang, & K. Sylvan (Eds.), *Routledge Handbook of practical reasoning* (pp. 349–365). Routledge.
- Kahneman, D. (2011). *Thinking, fast and slow*. Penguin.
- Katsafanas, P. (2013). *Agency and the foundations of ethics: Nietzschean constitutivism*. Oxford University Press.
- Kearns, S., & Star, D. (2008). Reasons: Explanations or evidence? *Ethics*, 119, 31–56.
- Kearns, S., & Star, D. (2009). Reasons as evidence. *Oxford Studies in Metaethics* (4 vol., pp. 215–242). Oxford University Press. Russ Shafer-Landau.

²⁴ I would like to thank Erasmus Mayr, Anselm W. Müller, Eva Schmidt and two anonymous reviewers for their comments on earlier drafts of this paper, and Aaron Shoichet for improving my English.

- Kietzmann, C. (2022a). Can realists reason with reasons? *Philosophical Explorations*, 25:2, 159–169.
- Kietzmann, C. (2022b). Inference is not a process. *Australasian Journal of Philosophy*, 100:3, 539–549.
- Korsgaard, C. M. (2009). The Activity of Reason. *Proceedings and Addresses of the American Philosophical Association* 83:2, 23–43.
- Logins, A. (2020). Normative reasons without (good) reasoning. *Ethics*, 130:2, 208–210.
- Markovits, J. (2014). *Moral reason*. Oxford University Press.
- McHugh, C., & Way, J. (2016). *Fittingness First Ethics* 126, 575–606.
- McKeever, S., & Ridge, M. (2012). Elusive reasons. *Oxford Studies in Metaethics* (7 vol., pp. 110–137). Oxford University Press. Russ Shafer-Landau.
- Müller, A. W. (1979). How theoretical is practical reason? In C. Diamond, & J. Teichman (Eds.), *Intention and intentionality: Essays in Honour of G. E. M. Anscombe* (pp. 91–108). Harvester.
- Müller, A. W. (1992). Mental Teleology. *Proceedings of the Aristotelian Society*, 92, 161–183.
- Müller, A. W. (2004). Acting well. In O. H. Anthony (Ed.), *Modern Moral Philosophy: Royal Institute of Philosophy supplement 54* (pp. 15–46). Cambridge University Press.
- Paakkunainen, H. (2017). Can there be Government House reasons for action? *Journal of Ethics and Social Philosophy*, 12:1, 56–93.
- Parfit, D. (2011). *On what matters* (2 vols). Oxford.
- Raz, J. (1975). *Practical reason and norms*. Oxford University Press.
- Raz, J. (1978). Introduction. In J. Raz (Ed.), *Practical reasoning* (pp. 1–17). Oxford University Press.
- Raz, J. (1999). *Engaging reason*. Oxford University Press.
- Rossi, B. (2021). Elusive reasons and the motivational constraint. *Journal of Ethics and Social Philosophy*, 20:1, 82–110.
- Scanlon, T. M. (1998). *What we owe to each other*. Harvard University Press.
- Scanlon, T. M. (2014). *Being realistic about reasons*. Oxford University Press.
- Schmidt, E. (2021). Where reasons and reasoning come apart. *Noûs*, 55(4), 762–781.
- Schroeder, M. (2007). *Slaves of the Passions*. Oxford University Press.
- Setiya, K. (2014). What is a reason to Act? *Philosophical Studies*, 167, 221–235.
- Sidgwick, H. (1907). *The methods of Ethics* (7th ed.). Macmillan.
- Silverstein, M. (2016). Reducing reasons. *Journal of Ethics and Social Philosophy*, 10:1, 1–22.
- Strawson, P. F. (1992). *Analysis and Metaphysics*. Oxford University Press.
- Streumer, B. (2007). Reasons and impossibility. *Philosophical Studies*, 136, 351–384.
- Velleman, J. D. (2004). Précis of *the possibility of practical reason*. *Philosophical Studies*, 121:3, 225–238.
- Way, J. (2017). Reasons as premises of good reasoning. *Pacific Philosophical Quarterly*, 98, 251–270.
- Williams, B. (1981). Internal and external reasons. *Moral Luck: Philosophical Papers 1973–1980* (pp. 101–113). Cambridge University Press.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.