Level separation in epistemology

Abstract

This paper formulates and defends the approach of level separation in epistemology. After motivating and formulating level separation, I show how level separation is already used within epistemology to protect state norms from procedural defeat as well as to relocate wrong-kind reasons for doxastic attitudes. Then I show how level separation can be used to give attractive treatments of transitional attitudes and rational delay. I address an objection to level separation due to Susanna Rinard, then conclude with lessons for future research.

1 Introduction

Consider three questions raised by recent work in epistemology:

> **(Rational delay)** May agents rationally delay updating their beliefs in response to new evidence on the grounds that it takes time to think through the implications of new evidence? (Meacham 2015; Na'aman 2021a; Podgorski 2017)
>
> **(Transitional attitudes)** What attitudes may agents rationally take towards the targets of active inquiry? (Friedman 2017; Palmira 2020; Staffel 2025)
>
> **(The epistemic and the zetetic)** Could rational inquiry ever produce beliefs which violate traditional epistemic norms? (Flores and Woodard 2023; Friedman 2020; Thorstad 2021)

My aim in this paper is to suggest that all of these questions can be given a unified, principled and explanatorily powerful answer. That answer invokes a strict *level separation* between normative questions about rational attitudes and normative questions about rational processes of inquiry.

Roughly put, level separation holds that the rationality of inquiry does not determine the rationality of the attitudes that inquiry produces. Instead, the rationality of both

inquiry and the resulting attitudes must be separately determined by distinct applications of the norms governing each. We will see in Sections 2-3 how level separation generalizes to other objects of normative assessment.

Although level separation is commonly invoked in practical philosophy (Driver 2012; Kagan 2000; Parfit 1984), it is less commonly invoked within epistemology and is sometimes overlooked entirely in epistemological discussions. My aim in this paper is to give a full statement of level separation and make the case for more expansive use of level separation within epistemology.

Here is the plan. Sections 2-3 introduce, motivate and develop an account of level separation. Section 4 illustrates two ways in which level separation is already invoked within practical and theoretical philosophy: to protect state norms from procedural defeat (Section 4.1) and to relocate wrong-kind reasons for attitudes (Section 4.2). Sections 5-6 show how level separation can be used to give attractive treatments of transitional attitudes (Section 5) and rational delay (Section 6). Section 7 addresses an objection due to Susanna Rinard (2019). Section 8 concludes.

2 Level separation

In the 1970s, many philosophers held that consequentialism is self-refuting.[1] To see the problem, suppose you are shopping for coffee beans. There are many varieties of beans on the shelf. One, the Sumatra, is slightly better than the others. Let A be the action of buying the Sumatra. Consequentialists say that you ought to intend to do A.

How should you deliberate about which coffee beans to buy? Consequentialists typically deny that you should explicitly calculate the expected utilities of each option (Parfit 1984; Railton 1984). Instead, you might employ a simpler procedure, perhaps choosing the first in-budget single-origin beans that smell nice and are familiar to you. If you deliberate in this way, you will often choose a high-quality option. But sometimes, you

[1] See (Bales 1971) for a good discussion of early arguments, and (Parfit 1984; Railton 1984) for the orthodox response.

will pick something other than the Sumatra — perhaps the Yirgacheffe. Let B be the action of buying the Yirgacheffe.

Here we have a puzzle. Consequentialism says that you should intend to do A, but deliberate in a way that will sometimes lead you to intend B instead. But how could it be that rational inquiry leads you to intend B, yet you are rationally required to intend A and rationally forbidden from intending B? This has seemed self-undermining or self-refuting to many critics. Furthermore, we will see in Section 7 how this puzzle might be extended to an argument that consequentialism violates the principle that ought implies can. Checkmate, consequentialists?

Most consequentialists have thought that this criticism is too hasty. To see why, we will need a bit of machinery (Kagan 2000). There are various *evaluative focal points*, or objects of normative assessment, including intentions and deliberation procedures. At each focal point, we can apply *normative terms* such as 'rational', 'ought' or 'blameworthy'.

We might represent the facts about rationality in our coffee-buying case as in Figure 1. Here the agent's deliberation procedure is rational, because expected utility maximizing, but her resulting intention is irrational, because not expected utility maximizing. There is something troubling about this situation, but what is it?

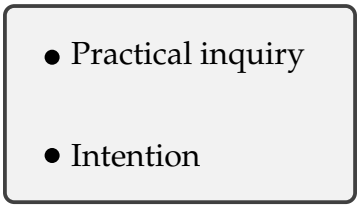

Figure 1: Intention and practical inquiry

We might expand Figure 1 to represent other focal points, such as belief, in Figure 2. It may turn out that the agent's coffee-related beliefs are rational, or that they are irrational. But neither verdict seems to be in tension with our judgments about the rationality of the agent's intentions or practical inquiries. We don't bat an eye at saying that someone had rational beliefs about coffee but intended or acted irrationally anyways. For this reason, Figure 2 represents belief as orthogonal to intention, but practical inquiry as bearing

some hierarchical relationship to intention. This hierarchical relationship is what makes us uncomfortable about a mismatch between rationality at the evaluative focal points 'intention' and 'practical inquiry', but not uncomfortable about mismatch between either of these focal points and 'belief'. But what does this hierarchical relationship represent?

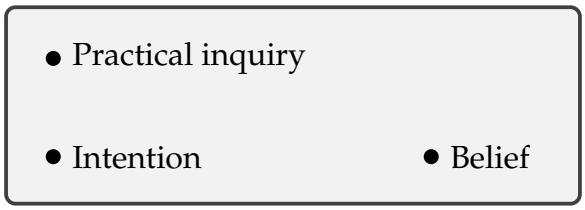

Figure 2: Belief, intention and practical inquiry

Here is a natural proposal.[2] Say that focal point X *governs* focal point Y if members of the governed class Y are typically produced and modified by members of the governing class X. Practical inquiry governs intention because practical inquiry is the process through which intentions are typically produced and modified. However, belief does not govern intention. Although beliefs enter into practical deliberation, it would not be right to say that typical intentions are directly produced and modified by beliefs. That role belongs to practical inquiry.

The tension which troubles us arises when the same normative term occurs with different valence across governing focal points. For example, our coffee buyer inquires rationally but intends irrationally, giving rise to tension across governing focal points. This seems troubling, because the very same process of practical inquiry through which intentions are typically produced and modified led through fully rational inquiry to an irrational intention. Call tension of this sort *level tension* to reflect how governing focal points are diagrammatically represented a level above the focal points they govern.

Traditionally, consequentialists have reacted to level tension in one of two ways. Indirect consequentialists have reacted by linking the normative status of governed focal points, such as intention, to the normative status of governing points, such as practical inquiry (Adams 1976; Brandt 1959; Harsanyi 1977). For example, we might say that an in-

[2]See Thorstad (2021) for discussion.

tention is rational just in case it results from a rational process of practical inquiry (Figure 3).

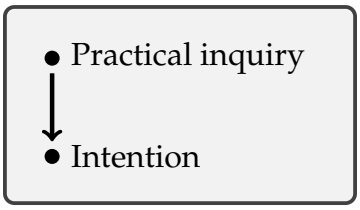

Figure 3: Indirect consequentialism

By contrast, direct consequentialists have reacted by retaining level tension, arguing that level tension is not a contradiction but rather an important fact about normative theory (Driver 2012; Kagan 2000; Parfit 1984). On this view, there is some tension between the statements that our coffee buyer inquired rationally and intended irrationally. However, this tension is not a contradiction and should be treated as an important feature of the normative landscape.

Direct consequentialism allows us to directly apply normative theories at each focal point. For example, we can say as I did above that a practical inquiry is rational just in case it is expected utility maximizing, and that an intention is rational just in case it is expected utility maximizing as well. Or, if we like, we can apply different normative theories directly at each focal point. We could, for example, couple a fittingness-based account of rational attitudes with a consequentialist account of rational inquiry.

Direct consequentialism is one of many ways to apply a strict *level separation* between governing focal points. Level separation involves postulating separate normative theories governing each focal point, then directly applying the relevant normative theory at each focal point, rather than evaluating the normative status of some focal points in terms of the normative status of other focal points. At a first pass, this means:

> **(LS-1)** For all normative terms N and governing pairs F, F' of focal points, the N-status of F is independent of the N-status of F'.

where focal points F, F' constitute a *governing pair* just in case F governs F' or F' governs F. We will see in Section 3 that this first pass is not quite right, but it will do for now.

Theoretical inquiry governs belief in exactly the same way that practical inquiry governs intention (Figure 4). Just as seemingly rational practical inquiry can give rise to seemingly irrational intentions, seemingly rational theoretical inquiry can give rise to seemingly irrational beliefs (Friedman 2020). It is natural to apply a strict level separation not only to practical inquiry and intention, but also to theoretical inquiry and belief. This allows us to retain traditional normative theories of rational belief, such as evidentialism or coherentism, while adopting any theory of rational inquiry on offer (Falbo 2023; Kelp 2021; Willard-Kyle 2023), including an instrumentalist (Cowie 2014; Friedman 2018; Steglich-Petersen forthcoming) or consequentialist theory (Stich 1990; Thorstad 2024; Schurz and Hertwig 2019).

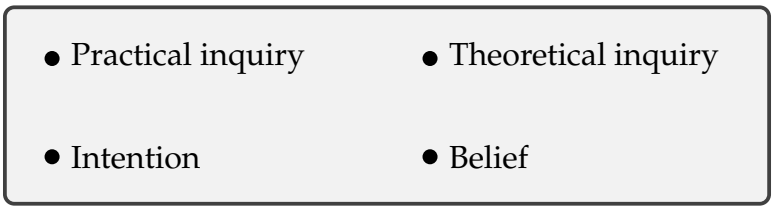

Figure 4: Belief, intention and their governing points

That is the basic idea. It turns out that the basic idea needs a good deal of elucidation.

3 Formulating level separation

Section 2 formulated level separation as the claim that:

> **(LS-1)** For all normative terms N and governing pairs F, F' of focal points, the N-status of F is independent of the N-status of F'.

LS-1 requires at least four modifications.

3.1 Grounding

Direct consequentialists wish to ground the value of acts in the value of their outcomes. However, they do not wish to ground the value of outcomes in the value of the acts

that produce them. This creates a problem for LS-1 because dependence can run in both directions. For example, many direct consequentialists think it is a necessary truth that the value of a constant act is the value of its unique outcome. Then it is not clear that the value of a constant act is independent of the value of its outcome, which would violate LS-1.

The natural solution is to replace independence with a more heavyweight, asymmetric metaphysical notion. I will reformulate level separation in terms of metaphysical grounding, but readers who prefer other notions such as supervenience or even a souped-up notion of dependence are welcome to read this discussion with those notions in place of grounding. I take grounding to be a relation between facts, treating facts as true propositions and using $[p]$, $< p >$ to denote facts and propositions, respectively. All mentions of grounding in this paper refer to weak grounding unless otherwise specified.[3]

In these terms, level separation becomes:

> **(LS-2)** For all normative terms N, governing pairs F, F' of focal points, and elements f of F, if $< Nf >$ is true then $[Nf]$ is not grounded in facts involving the N-status of F'.

LS-2 solves our initial problem. Although the value of constant acts covaries with the value of their outcomes, this need not imply that facts of the form [o has value r] ground facts of the form [a_o has value r], where a_o is the constant act producing outcome o. So far so good.

3.2 Immediate grounding

Direct consequentialists may not ground the value of outcomes in the value of acts. But as Julia Staffel (forthcoming) notes, direct consequentialists do seem to ground the value of acts in the value of their outcomes. On a traditional theory, direct consequentialists

[3] A popular definition of weak grounding follows Fine (2012): facts Γ fully weakly ground $[p]$ just in case, for any $[q]$, if $[p]$ fully strictly grounds $[q]$ then Γ fully strictly grounds $[q]$.

evaluate acts by combining the evidential probabilities Pr_E of worlds w with the values $V(a, w)$ of acts in worlds, as:

$$V(a) = \sum_w Pr_E(w)V(a, w). \tag{1}$$

It is natural to read (1) as grounding the value of an act a in the value of its possible outcomes. More precisely, let V_a be the outcome values that a has a nontrivial probability of bringing about. For each $v \in V_a$, let $V(a) = v$ be the event that a takes value v. Then (1) is naturally read as supporting:

> **(Outcome Values Ground Act Values)** For any act a, [a has value $V(a)$] is fully grounded in the facts $\{$[a has evidential probability $Pr_E(V(a) = v)$ of producing an outcome with value v]$\}_{v \in V_a}$.

Outcome Values Ground Act Values contradicts LS-2 by grounding the values of acts in the values of their outcomes.

One response would be for direct consequentialists to claim that (1) expresses a relationship of necessary coextensiveness rather than grounding.[4] After all, a mature axiology tells us which features of outcomes matter and how much they matter. Coding outcomes as tuples $< o_i >= (o_1^i, \ldots, o_n^i)$ which specify them along all dimensions that matter, a mature axiology will give a function f mapping outcomes to their values. (1) then implies a relationship between the features of outcomes and the value of acts, as:[5]

$$V(a) = \sum_i Pr_E(< o_i > |a) f(< o_i >). \tag{2}$$

Letting O_a be the possible outcomes of a, (2) could be used to argue that:[6]

> **(Outcome Features Ground Act Values)** For any act a, [a has value $V(a)$] is fully

[4] Kagan (2000) sketches some thoughts in this direction.

[5] As is well known, (2) may need to be modified if we allow violations of act-state independence (Joyce 1999).

[6] See Berker (2018a) for an argument in favor of excluding f from the statement of Outcome Features Ground Act Values.

grounded in the facts $\{[a$ has evidential probability $Pr_E(< o > |a)$ of producing outcome $< o >]\}_{<o> \in O_a}$.

Outcome Features Ground Act Values does not contradict LS-2 because it does not ground the value of acts in the value of their outcomes.

The problem is not that Outcome Features Ground Act Values is false, but rather that Outcome Features Ground Act Values is compatible with Outcome Values Ground Act Values. After all, many facts are multiply grounded.

A hard-line response would be to deny Outcome Values Ground Act Values. Perhaps there are grounds for this response.[7] At the same time, there is no need to bet the house on denying Outcome Values Ground Act Values.

Some grounding relationships are immediate. For example, [Tom Lehrer was a singer] immediately grounds [Tom Lehrer was an entertainer]. Other grounding relationships are mediate, consisting of chains of immediate grounding. For example, [Tom Lehrer was a singer] mediately grounds [someone was an entertainer] because [Tom Lehrer was a singer] immediately grounds [Tom Lehrer was an entertainer], which immediately grounds [someone was an entertainer].

When direct consequentialists say that the values of acts are grounded in the features of their outcomes, they mean to make a stronger claim: that the values of acts are immediately grounded in the features of their outcomes. (2) expresses an immediate grounding relation that does not need to be mediated through (1). If this is right, then level separation should claim:

> **(LS-3)** For all normative terms N, governing pairs F, F' of focal points, and elements f of F, if $< Nf >$ is true then $[Nf]$ is fully immediately grounded in a collection of facts not involving the N-status of F'.

[7] For example, we might be concerned to avoid having one thought too many (Williams 1981), fetishizing value by making it an object of pursuit in its own right (Smith 1994) or ignoring what Ross called the *resultant* nature of moral properties: they can only be possessed in virtue of possessing other properties (Berker 2018a; Ross 1930).

LS-3 does not rule out the possibility that the value of acts may also be fully grounded in facts about the value of their outcomes, nor even that this grounding should be immediate. LS-3 simply requires that the value of acts be immediately grounded in a collection of facts that are not about the value of their outcomes.

One way to see how LS-3 constitutes an improvement is to note that LS-3 is the first principle violated by traditional indirect theories such as motive consequentialism. On a first pass, motive consequentialists think that:[8]

> **(Motive Values Immediately Ground Act Values)** For any act a, [a has value $V(a)$] is fully immediately grounded in the facts [a results from motives M] and $\{[m \text{ has value } V(m)]\}_{m \in M}$.

Letting O_m be the outcomes that motive m has a nontrivial probability of bringing about, motive consequentialists could also hold:

> **(Outcome Features Immediately Ground Motive Values)** For any motive m, [m has value $V(m)$] is fully immediately grounded in the facts $\{[m \text{ has evidential probability } Pr_E(< o > | m) \text{ of producing outcome } < o >]\}_{<o> \in O_m}$.

But even if we allowed the motive consequentialist to drop the grounding fact [a results from motives M] from Outcome Values Immediately Ground Motive Values, this would only imply that:

> **(Outcome Features Mediately Ground Act Values)** For any act a, [a has value $V(a)$] is fully mediately grounded in the facts $\{[m \text{ has evidential probability } Pr_E(< o > | m) \text{ of producing outcome } < o >]\}_{m \in M, <o> \in O_m}$.

It would not imply that outcome features immediately ground act values. On a traditional motive consequentialist story, facts about the value of outcomes bear on the value of acts only through bearing on the value of the motives that produce them.

[8] Some motive consequentialists might prefer to evaluate motive sets rather than individual motives. The discussion below will be broadly similar on these reformulated views.

3.3 Naturalism and ungrounded facts

LS-3 implies that all normative facts are fully grounded. This rules out the possibility of ungrounded normative facts, a possibility which is important to many traditions including normative non-naturalism (Berker 2018a,b). This problem can be fixed by restricting LS-3 to normative facts that have full grounds:

> **(LS-4)** For all normative terms N, governing pairs F, F' of focal points, and elements f of F, if $< Nf >$ is true and fully immediately grounded, then $[Nf]$ is fully immediately grounded in a collection of facts not involving the N-status of F'.

LS-4 does not rule out the possibility of ungrounded normative facts, because the antecedent of LS-4 exempts them from its consequent grounding requirements.

3.4 Aboutness

The last order of business is to say what it means for facts to involve the N-status of F'. Again, we will need a heavyweight metaphysical notion for the task. In particular, involvement cannot be implied by necessary coextensiveness or we will be back in the soup pot that grounding was meant to rescue us from.

The most natural thing to say is that the grounds for facts of the form $[Nf]$ should not say anything about the N-status of F'. And in fact we can say just this.

Philosophers of language increasingly recognize that in addition to their truth conditions, propositions have subject matters that they make claims about (Goodman 1961; Lewis 1988; Yablo 2014). For example, if o is an outcome then the proposition $< o$ has value $v(o) >$ makes a claim about **the value of o**, where bold text is used to denote subject matters. If a_o is the constant act producing outcome o, then the proposition $< a_o$ has value $v(a_o) >$ makes a claim about **the value of $\mathbf{a_o}$**. Most leading theories of aboutness allow for **the value of *o*** and **the value of $\mathbf{a_o}$** to be distinct subject matters, even if o and a_o necessarily

take the same value.[9]

Reading involvement as aboutness, LS-4 becomes:

> **(LS-5)** For all normative terms N, governing pairs F, F' of focal points, and elements f of F, if $< Nf >$ is true and fully immediately grounded, then $[Nf]$ is fully immediately grounded in a collection of facts, none of which are about **the N-status of F′**.[10]

This completes our account of level separation.

I think that level separation, thus formulated, is a productive approach that should be judged in large part by what it can do for epistemology. My primary project in this paper is to illustrate two new areas of epistemology in which level separation gives plausible and explanatorily powerful results (Sections 5-6). However, it is worth recalling that level separation already plays an important role in epistemology paralleling its role in practical philosophy.

4 Traditional applications of level separation

In this section, I give two examples to illustrate what level separation currently does for epistemology. Many readers may already accept the applications in question, in which case they may already have reason to accept level separation. In each example, I show how the same application is realized in both practical and theoretical philosophy in order to draw support from both domains.

[9]Indeed, Peter Hawke (2018) includes a version of this requirement as a desideratum on theories of aboutness.

[10]Since Lewis (1988), many theories have allowed propositions to be *partly about* subject matters. On many theories, facts of the form [a has evidential probability $Pr_E(V(a) = v)$ of producing an outcome with value v] will be only partly about the value of outcomes. Given Outcome Values Ground Act Values, this means that LS-5 can be true because facts of the form [a has value $V(a)$] are fully grounded in the facts $\{$[a has evidential probability $Pr_E(V(a) = v)$ of producing an outcome with value v]$\}_{v \in V_a}$. This is not really what advocates of level separation had in mind. A natural way to block it is to replace 'about' with 'partially about' in LS-5, though the details will depend on which notion of partial aboutness is used.

4.1 Protecting state norms from procedural defeat

Many approaches to practical rationality allow that an inquiry can be rational even if it leads to an irrational intention. We saw one case of this phenomenon in Section 2, where mainline consequentialists allow that an expected-utility-maximizing practical inquiry need not produce an expected-utility-maximizing intention. Without level separation, this leads to a quick charge that consequentialism is self-defeating, and in fact the same charge has been leveled against many non-consequentialist views (Stocker 1976). However, we saw that level separation removes the danger of self-defeat by allowing consequentialists and others to say without contradiction that the rationality of inquiries may differ from the rationality of the intentions they produce.

In the theoretical domain, Jane Friedman (2020) has made the parallel observation that plausible norms of rational inquiry come apart from received theories of rational belief. For example, while you are busy counting windows it may be irrational to inquire about other matters because this would distract from your larger inquiry. This verdict follows from norms such as:

> **(Zetetic Instrumental Principle (ZIP))** If one wants to figure out $Q^?$, then one ought to take the necessary means to figuring out $Q^?$. (Friedman 2020, p. 503)

Inquiring about unrelated matters is irrational, according to ZIP, because focusing on the question at hand is a necessary means to answering it.

However, if you were to irrationally interrupt your window-counting to form evidentially supported beliefs about other matters, those beliefs would count as rational by the lights of evidentialism and many other theories of rational belief. This verdict follows from norms such as:

> **($\mathbf{EP_a}$)** If one has excellent evidence for p at t, then one is permitted to judge p at t. (Friedman 2020, p. 504)

So long as you have excellent evidence for unrelated propositions, EP_a permits you to judge them true.

Friedman (2020) suggests that these verdicts show a deep tension between plausible norms of rational inquiry and traditional norms of rational belief. Friedman suggests that this tension should probably be resolved by rejecting traditional norms of rational belief.

By contrast, level separation allows us to capture both verdicts in Friedman's cases without revising traditional norms of rational belief.[11] Given level separation, it is not a contradiction to say that a wasteful inquiry would be irrational, even if it would lead to a belief that is rational because evidentially supported. Norms of inquiry such as ZIP are fully compatible with traditional epistemic norms so long as they are understood in their traditional sense as norms governing belief.

For example, we might hold:

> **($\mathbf{EP_b}$)** If one has excellent evidence for p at t, then one is permitted to believe p at t.

By ZIP, an agent who interrupted their window-counting would inquire rationally, even if the belief that she formed was evidentially supported and thereby rational by EP_b. This is not a contradiction so long as we do not read evidentialism as a norm governing inquiry, along the lines of:

> **($\mathbf{EP_i}$)** If one has excellent evidence for p at t, then one is permitted to initiate some process of inquiry at t terminating in the belief that p.

EP_i would permit an agent to initiate the same irrelevant inquiry that ZIP forbids. However, EP_b does not permit this or any other inquiry.

Following evidentialist orthodoxy (Feldman 2000), a level-separated approach denies EP_i because it interprets evidentialism as a claim about rational belief rather than a claim about rational inquiry. The fact that a wastefully irrational and distracting inquiry would terminate in a rational belief does not change the fact that this inquiry is irrational.

[11] See Thorstad (2021) for discussion.

In this way, level separation protects traditional norms of rational belief from procedurally driven challenges in exactly the same way that it protects traditional norms of rational intention from the same challenges.

4.2 Relocating wrong-kind reasons

What reasons bear on whether I should intend to drink a bottle of toxins (Kavka 1983)? Two categories of reasons suggest themselves. On the one hand, there are unproblematic reasons such as the fact that toxins will harm me and the fact that they taste bad. On the other hand, there are strange apparent reasons for intending to drink toxins, such as the fact that you will pay me a million dollars for intending to drink them or that my so intending will make you happy. There is something intuitively more problematic about reasons in the second class, which has led philosophers to label them as reasons of the wrong kind, as opposed to the reasons in the first class, which appear to be of the right kind to rationalize intention (Gertken and Kiesewetter 2017; Hieronymi 2005; Parfit 2011).

Much ink has been spilled on the nature and importance of the distinction between right- and wrong-kinds of reasons (Gertken and Kiesewetter 2017; Hieronymi 2005; Schroeder 2012). But one prominent tradition has it that:

> **(WKR Skepticism)** Apparent wrong-kind reasons for attitudes are actually reasons for something else. (Gibbard 1990; Kolodny 2005; Way 2012)

WKR skeptics press this case by noting that wrong-kind reasons for intention lack some features we might expect reasons for intention to have. For example, wrong-kind reasons cannot easily motivate intention, do not figure directly in deliberation about what to intend, and seem intuitively less apt to be reasons for intention (Schroeder 2012).

Instead, skeptics hold, apparent wrong-kind reasons for intention are actually reasons to *get oneself to intend* to drink a toxin. This is accomplished through various processes, including inquiry. I might, for example, reflect on the good features of the toxin or the bad features of life in order to help myself acquire an intention to drink the toxin. Applying

level separation, we may admit that such an inquiry is rational while denying that the resulting intention to drink a toxin is rational.

Similar questions arise in epistemology. What reasons bear on whether I should believe that I am a good dancer? On the one hand, there are unproblematic reasons such as my knowledge of dancing or the compliments I have received on my dancing. On the other hand, there are more controversial reasons such as the fact that believing I am a good dancer will make me confident and happy. Traditionally, epistemologists have held that these apparent practical reasons for belief are not, in fact, reasons to believe I am a good dancer, but rather reasons to get myself to believe I am a good dancer. They have held this on much the same grounds as in the practical case: apparent practical reasons for belief may not play the right role in motivation (Kelly 2002, 2003) or deliberation (Shah 2003, 2006), and seem intuitively less apt to be reasons for belief (Firth 1956).

As before, this view applies level separation to allow that it may be rational to inquire in a way that leads me to believe that I am a good dancer, while denying that it is rational for me to believe I am a good dancer, since I have abundant evidence that I am a terrible dancer. This line of argument has been viewed as an integral part of the defense of evidentialism and other traditional epistemic norms, providing a principled reason for the traditional view that apparent practical reasons for belief are actually reasons to get oneself to believe and thereby avoiding pragmatic encroachment on rational belief (Kelly 2003; Singh 2021). And here the appeal to level separation is unavoidable, since inquiry governs belief, yet it may be rational to inquire in a way that leads me to believe something which is irrational for me to believe. If we want to retain this traditional argument, we need to appeal to level separation.

4.3 Taking stock

So far, we have seen how level separation can be applied within theoretical and practical philosophy to protect state norms from procedural defeat and to relocate apparent wrong-kind reasons for attitudes as right-kind reasons to get oneself to hold attitudes. But these

applications are familiar. What else can level separation do for us? In the next two sections, I show how level separation can be applied to give attractive treatments of transitional attitudes (Section 5) and rational delay (Section 6).

5 Transitional attitudes

5.1 Two types of rationality

Julia Staffel distinguishes between *transitional attitudes* held towards the objects of ongoing deliberation and *terminal attitudes* held as conclusions of deliberation (Staffel 2019, 2021a,b, 2025). On Staffel's view, transitional and terminal attitudes play distinct functional roles:

> **(Representational Function for Terminal Attitudes**) The function of terminal attitudes is to represent the world.
>
> **(Placeholder Function for Transitional Attitudes**) The function of transitional attitudes is to serve as placeholders for the agent's position in an ongoing process of reasoning.

This functional difference between transitional and terminal attitudes gives rise to descriptive and normative differences between them.

Staffel holds that transitional attitudes typically differ from terminal attitudes in three descriptive respects.

> **(Stability Datum**) Terminal attitudes are typically more stable than transitional attitudes are.
>
> **(Guidance Datum**) Terminal attitudes are typically more available to guide action than transitional attitudes are.
>
> **(Updating Datum**) Terminal attitudes are typically more available to update other attitudes than transitional attitudes are.

Staffel also holds that transitional attitudes differ from terminal attitudes in normative respects.

Consider:

> **(Framed)** Manny has committed a murder, and tries to frame Fred for it. Detective Fletcher, upon initially inspecting the evidence, responds as Manny has planned, and becomes 90% confident that Fred committed the murder. However, as she evaluates the evidence more carefully, she discovers incongruencies that ultimately lead her to conclude that Fred was framed, so she reduces her confidence that Fred is the murderer to 5%. (Staffel 2021a, p. 387)

Staffel takes it as a datum that:

> **(Credence Rational in Framed)** Manny's initial credence is rational in Framed.

Leading theories of epistemic rationality imply that Manny's initial credence is irrational, for example because it is not supported by Manny's total evidence. Capturing Credence Rational in Framed therefore requires a new normative theory.

Staffel holds that transitional attitudes are governed by a distinctive type of *pro tem* rationality, to be distinguished from the *pro toto* rationality governing terminal attitudes. Although traditional theories may give correct accounts of pro toto rationality, Staffel holds that they are inadequate accounts of pro tem rationality because they wrongly classify the beliefs of agents such as Manny as irrational.

Staffel offers the following account of pro tem propositional rationality:

> **(Pro Tem Propositional Rationality)** A doxastic transitional attitude d is pro tem rational for an agent to adopt towards an answer to some question q at some time t just in case
>
> (I) The agent adopts d while using a permissible cognitive process to settle the question q, where a permissible cognitive process (or combination of

processes) is one that gives the agent a sufficiently high probability of reaching a doxastically rational terminal attitude, given the input attitude to the process(es); and

(II) At t, d is supported by the combination of (a) the truth-relevant evidence the agent has considered up to t, (b) the agent's take on the logical and probabilistic relevance of this evidence for q at t; and

(III) Evidence regarding the quality of the agent's reasoning up to t is factored into d when level-merging [combined consideration of first- and higher-order evidence] is beneficial. (Staffel 2025)

Staffel shows how this account may be extended to a theory of pro tem doxastic rationality, although the details fall beyond the scope of this paper.

Below, I offer (Section 5.2) and evaluate (Section 5.3) an alternative level-separated approach.

5.2 A level-separated alternative

If we accept level separation, then we may readily grant the datum that:

(Something Rational in Framed) There is an important sense in which Manny cognizes rationally in Framed.

But we cannot treat as a datum the claim that:

(Credence Rational in Framed) Manny's initial credence is rational in Framed.

Credence Rational in Framed is not an uncontroversial datum, since it directly contradicts most leading theories of rational credence including evidentialism. Instead, a level-separated alternative vindicates Something Rational in Framed by holding that Manny's process of inquiry is rational. Just as a rational process of inquiry may produce irrational attitudes at the conclusion of inquiry, it may likewise produce irrational attitudes during

the course of inquiry. We may argue that these attitudes are blamelessly irrational, but they remain irrational nonetheless.

A level separated view accepts Representational Function for Terminal Attitudes but extends the same function to transitional attitudes:

> **(Representational Function for Transitional Attitudes)** The function of transitional attitudes is to represent the world.

On this view, doxastic attitudes are intentional attitudes with a mind-to-world direction of fit. Doxastic attitudes succeed in their function when they correctly represent the world. While attitudes may become more likely to serve their representational function successfully as inquiry proceeds, it remains true that the agent's earliest transitional attitudes are her best attempt to represent the world as it is. This view concedes the causal fact that an agent's doxastic attitudes are determined by her current place in reasoning. However, it does not equate this causal fact with the representational function of doxastic attitudes, which is to represent their intentional objects and not their own causal origins.

A level-separated approach draws on traditional explanations for the Stability, Guidance and Updating Data, rather than explaining them by positing a novel function for transitional attitudes. While the correct explanation of each datum remains an object of ongoing debate, there are good options for a level-separated approach to explain each datum.

Begin with the Stability Datum. A level-separated approach offers three reasons to expect terminal attitudes to be more stable than transitional attitudes. First, transitional attitudes are the targets of ongoing processes of inquiry whose typical effect is to modify agents' attitudes towards the question under consideration. Second, because terminal attitudes are formed at the conclusion of inquiry, they tend to be better supported and hence more difficult to overturn. Finally, on many views the attitudes that conclude inquiry come with a commitment to resist reopening inquiry without good reason (Bratman 1987; Fraser forthcoming; Holton 2014). This commitment contributes additional stability to terminal attitudes.

Turn next to the Guidance Datum. A level-separated approach can explain at least two ways in which transitional attitudes are less available to guide action than terminal attitudes are. First, agents who hold transitional attitudes may prefer to postpone decisionmaking until deliberation is concluded. This is traditionally explained by the fact that agents have the option to postpone decisionmaking rather than choosing immediately, and should take this option when the expected value of information yielded by additional deliberation exceeds expected costs.

A second way in which transitional attitudes are less available to guide action is that agents may be more willing to treat terminal attitudes as premises in deliberation about how to act. In the case of belief, this can be explained if we take on board two premises. First, it is appropriate for agents to take full beliefs, but not non-beliefs as premises in practical deliberation (Ross and Schroeder 2014; Williamson 2000). Second, inquiry is descriptively or normatively incompatible with full belief (Fraser forthcoming; Friedman 2019; Millson 2021). It follows that appropriate premises for practical deliberation must be full beliefs, and hence must be terminal attitudes. This explanation is, of course, controversial, but readers who deny one or more premises may wish to deny the datum.

Turn finally to the Updating Datum. It is not uncontroversial to say that transitional attitudes are less apt to be used to update other beliefs than terminal attitudes are.[12] However, those who wish to deny that attitudes formed during reasoning are always used to update other beliefs have many explanations at their disposal. They may point to limits on cognitive resources and updating abilities, the expectation that beliefs will soon change and require further updates, or the lack of pressing practical need to perform a full update. Staffel herself appeals to many of these same factors. But precisely for this reason, there are ample resources to explain the Updating Datum without positing a functional or normative difference between transitional and terminal attitudes. For example, we might hold that given scarce resources and abilities, agents must update rarely, and given

[12]Many leading approaches in Bayesian cognitive science such as hierarchical Bayesian approaches (Friston 2010; Howhy 2013) and even leading theories of resource-rationality (Icard ms; Lieder and Griffiths 2020) often treat updating as relatively costless, rapid, and continuously performed.

the relatively high likelihood that transitional attitudes will be revised soon, it would be irrationally wasteful to update upon transitional attitudes unless there were an immediate and pressing need to do so.

5.3 Evaluating the alternative

The level-separated explanation of the function, norms and descriptive characteristics of transitional attitudes has three important advantages. First, it is parsimonious. We do not need to posit a new type of attitude, a new type of rationality, a new normative theory, or a new functional role for doxastic attitudes in order to accept the level-separated view.

Second, the view is continuous with existing normative and descriptive theories. Normatively speaking, level separation retains compatibility with traditional epistemological theories such as evidentialism by separating the rationality of deliberative processes from the rationality of attitudes formed during deliberation. Level separation also preserves a traditional view on which doxastic attitudes have a world-to-mind direction of fit, function to represent the world, and succeed in their function when they correctly represent the world.

Descriptively speaking, level separation draws on a number of descriptive views that many philosophers already accept, including views about the role of belief in reasoning, the attitudes and commitments that settle inquiry, and the option to postpone decision-making in order to gather further evidence. If Staffel takes these explanations to be descriptively adequate, then there is no need for further explanation, and if they are taken to be descriptively inadequate, then there is a real risk of producing descriptive theories which force revision of existing descriptive views.

Finally, Staffel (2025) notes that the Stability, Guidance and Updating Data do not always reflect categorical differences between transitional and terminal attitudes. First, there is often a graded structure whereby transitional attitudes begin to function increasingly like terminal attitudes as the end of deliberation nears. Second, in edge cases some transitional attitudes may play many descriptive roles more strongly than some terminal

attitudes do. It would be good to have a principled explanation for these phenomena. Why do the descriptive differences between terminal attitudes seem to come in degrees? Why do they lessen as the end of inquiry nears? And why can they reverse in edge cases? The descriptive explanations offered by a level-separated view give plausible explanations of the graded differences between transitional and terminal attitudes which explain why these differences tend to lessen as inquiry progresses, and why they sometimes reverse.

Consider, for example, the Stability Datum. Let B be a transitional attitude formed near the end of a very rigorous, demanding, and lengthy inquiry. Plausibly, B becomes increasingly stable as inquiry nears its end, since ongoing inquiry is less likely to reveal new insights. Now let C be a terminal attitude formed at the conclusion of a less-demanding, everyday inquiry. C has some stability advantages over B, since C may come with a commitment to resist reopening inquiry, whereas B is the target of ongoing inquiry.

However, these advantages may well be offset by the fact that B results from more careful consideration of a greater quantity of dispositive evidence, such that it is unlikely that anything could come to light that would significantly revise B. This explanation follows naturally from the idea that the typical instability of transitional attitudes is driven by their being the targets of active inquiry, resulting from less thorough consideration of less-dispositive evidence, and lacking a commitment to reopen inquiry. At least one of these drivers, thoroughness of evidence and consideration, can come apart from the distinction between transitional and terminal attitudes. When it does, transitional attitudes may be more stable than terminal attitudes are.

If this is right, then level separation offers an attractive alternative to Staffel's theory of transitional and terminal attitudes. This alternative allows beliefs to retain their traditional functional role of representing the world and to be subject to traditional normative theories. A level-separated alternative explains why there is something importantly rational in cases such as Framed, while also offering a nuanced and traditional descriptive explanation of the differences between transitional and terminal attitudes. This explanation has the advantage of allowing us to explain the graded structure of the descriptive

data as well as edge cases in which the descriptive data reverse.

In the next section, I show how level separation can be applied to account for another phenomenon: rational delay.

6 Rational delay

In a series of papers, Abelard Podgorski defends a process-oriented view on which all fundamental epistemic norms govern processes (Podgorski 2016a,b,c,d, 2017). Podgorski applies this view to formulate novel versions of traditional views such as epistemic conservatism (Podgorski 2016a) and epistemic permissivism (Podgorski 2016b). Others have extended Podgorski's arguments to develop a procedural account of rational emotions (Na'aman 2021a,b).

In this section, I present Podgorski's view (Section 6.1) and a level-separated alternative (Section 6.2). I argue that the alternative has important advantages over Podgorski's process-oriented view, while retaining the advantages claimed for the process-oriented view.

6.1 Rational delay

Podgorski argues that theories of rational attitudes should accommodate:

> **(Rational Delay)** It takes time for agents like us to update our attitudes in response to changes in our mental state. (Podgorski 2017, p. 1)

However, Podgorski argues, norms governing attitudes cannot easily accommodate rational delay.

Consider, for example, the synchronic requirement to hold the attitudes we have most reason to hold:

> **(Synchronic State-Oriented Reasons)** If an agent's reasons favor attitude X at t, she is rationally required to have X at t. If her reasons disfavor attitude X at

t, she is rationally required not to have *X* at *t*. (Podgorski 2017, p. 3)

For example, evidentialists hold that reasons for belief are provided by an agent's total evidence. It follows from Synchronic State-Oriented Reasons that agents are rationally required to believe what their current evidence supports.

As Podgorski notes, Synchronic State-Oriented Reasons makes no room for rational delay. When an agent's reasons change, for example because she receives new evidence, the attitude that it is rational for her to hold immediately changes. No grace period is allowed for attitudinal adjustment. Podgorski argues that other leading synchronic norms likewise fail to make room for rational delay.

Perhaps diachronic norms fare better? We might, for example, build an interval of rational delay into a diachronic version of Synchronic State-Oriented Reasons. We might account for the interval of delay in terms of the agent's cognitive abilities. And we might include a placeholder for events which change or delay the requirements of rationality in the interim. We would then arrive at something like the following:

> **(Diachronic State-Oriented Reasons* (DSR*))** An agent is rationally required, if her reasons favor attitude *X* at t_1 and $t_2 - t_1$ is a duration appropriate to the agent's cognitive abilities as applied to the problem under consideration, and no cancelling or delaying event happens between t_1 and t_2 to preclude or forestall the need to form *X*, to have attitude *X* at t_2. (Podgorski 2017, p. 12)

For example, agents might be required to conditionalize on new evidence in an interval of time equal to the duration required for them to feasibly conditionalize.

However, Podgorski argues, DSR* does not look like a good candidate for a fundamental epistemic norm. When we ask questions about the parameters in DSR*, these questions are increasingly answered by citing procedural norms. For example, I might equally well have begun by stating the procedural norm that agents are rationally required to apply a cognitive process of conditionalization in order to update their beliefs, whose duration is as short as feasible. Agents who follow this procedural norm would also conform to

DSR*, with the interval $t_2 - t_1$ read directly from the interval of the rational process. Other parameters in DSR*, such as the nature and importance of cancelling and delaying events, would likewise be read off from the procedural norm.

But if compliance with diachronic state norms follows automatically from compliance with norms governing processes, and if fundamental parameters in state norms are read directly off of procedural norms, it becomes increasingly tempting to view norms governing processes as the fundamental epistemic norms. Therefore, Podgorski proposes a process-oriented picture of rationality:

> **(Process-oriented picture)** Rationality does not fundamentally govern states of mind like belief or intention at all. It governs processes such as reasoning and deliberation. (Podgorski 2017, p. 12)

On the process-oriented picture, there are no fundamental epistemic norms governing belief. We can perhaps say in some derivative sense that beliefs are rational when they result from rational processes. But this should be understood as a weak and non-fundamental shadow of the more fundamental norms governing attitudes.

Podgorski argues that the process-oriented picture satisfies two important theoretical desiderata.

> **(Delay Desideratum)** Theories of rationality should accommodate rational delay.

> **(Belief Norm Desideratum)** Theories of rationality should explain and derive norms governing belief.

The process-oriented picture satisfies the Delay Desideratum because rational processes take time to execute. The process-oriented picture satisfies the Belief Norm Desideratum in the sense that it allows us to explain and derive parameters in belief norms, such as the interval of rational delay.

6.2 Process norms without process-orientation

Advocates of level separation can agree with Podgorski in emphasizing the existence of norms governing processes of belief formation and revision, as exemplified by recent work in the epistemology of inquiry (Friedman 2020; Kelp 2021). They might even argue for one sense in which epistemology should be process-focused: in cases where procedural and attitudinal norms come apart, procedural norms often provide a more complete, charitable and grounded picture of the epistemic lives of agents like us (Morton 2017; Simon 1976; Thorstad 2022, 2024, forthcoming). However, level separation allows us to say these things without rejecting the existence or fundamentality of norms governing belief.

Although Podgorski states rational delay as a descriptive fact, presumably it is meant to be glossed as a normative datum. As with Staffel's motivating examples, we cannot take the datum to be that:

> **(Beliefs Rational Under Delay)** Beliefs retained during periods of rational delay are themselves rational.

Beliefs Rational Under Delay is not an uncontroversial datum, since it contradicts most leading epistemological views. But we can readily accommodate the datum that:

> **(Something Rational Under Delay)** Agents who exhibit rational delay are in an important sense cognizing rationally.

Indeed, on many views all rational processes of inquiry exhibit delay. Since ought implies can, rational processes must be processes that agents can execute. Since processes that agents can execute take time, all rational processes exhibit delay.

In this way, level separation accommodates rational delay. It also has three further advantages. First, the account is faithful to linguistic and philosophical practice. In our philosophical lives as well as our daily interactions, we regularly speak in normative terms

about belief, often saying directly of a belief that it is rational or justified. On a process-oriented picture, this talk cannot be taken at face value as expressing fundamental norms of rationality. Instead, it must be re-glossed, perhaps as a useful way of getting traction on processes of belief-formation which are difficult to directly introspect or discuss. By contrast, level separation allows us to take ordinary talk and philosophical theorizing at face value. When we say that a belief is rational, this is because there is a genuine standard of rationality governing belief that the belief has met.

Second, just as level separation poses no obstacle to the existence of norms governing belief, it also poses no obstacle to their fundamentality. Because level separation does not incorporate rational delay into attitudinal norms, we are not forced to posit unsightly and seemingly non-fundamental norms such as Diachronic State-Oriented Reasons*. Instead, we can remove the notion of an appropriate interval of rational delay and read cancelling events as changes in reasons, giving:

> **(Diachronic State-Oriented Reasons******)** An agent is rationally required, if her reasons favor attitude X at t_1 and her relevant reasons do not change between t_1 and t_2, to have attitude X at t_2.

There is nothing in the form of Diachronic State-Oriented Reasons** to suggest it must be derivative on some further procedural norm.

Finally, level separation is non-revisionary. It allows us to write down and defend synchronic attitudinal norms such as evidentialism or diachronic attitudinal norms such as conditionalization and Diachronic State-Oriented Reasons in their standard forms. By contrast, a process-oriented view not only denies the fundamentality of these norms, but also revises them to forms such as Diachronic State-Oriented Reasons*.

In addition to these advantages, level separation satisfies both of Podgorski's desiderata. Level separation satisfies the Delay Desideratum by treating rational delay as a constraint on rational inquiry. Level separation satisfies the Belief Norm Desideratum not by modifying belief norms and deriving them from process norms, but instead by retaining traditional belief norms and rendering them compatible with rational delay. We saw

in this section that level separation goes beyond the process-oriented picture's account of the Belief Norm Desideratum by not only deriving epistemic norms, but also making those norms true and fundamental.

6.3 Taking stock

In this section, we have seen how level separation provides an alternative to Podgorski's process-focused approach. We saw that a level-separated alternative has the advantages of capturing discourse about rational belief, permitting the existence of fundamental attitudinal norms, and avoiding revision to traditional epistemic norms. We also saw that level separation retains the advantages of Podgorski's process-focused approach.

Together with the arguments of Sections 5, this provides some positive evidence in favor of level separation. However, level separation is not without challenges. In the next section, I address a challenge to level separation due to Susanna Rinard.

7 Agglomeration

Susanna Rinard (2019) objects to level separation on the grounds that it violates a principle of agglomeration.

> **(Unrestricted Deontic Agglomeration)** For all agents S and all options A, B for S, if S ought to A and S ought to B then S ought to $(A \wedge B)$.

For example, consider:

> **(Pascal's Wager)** S's total evidence supports the claim that God does not exist. However, if S comes to believe that God exists, this will increase her chance of going to heaven.

The classic level separationist response to Pascal's Wager holds that S should A, get herself to believe that God exists, but should also B, not believe that God exists. It follows from

unrestricted deontic agglomeration that S ought to take the combined option $A \wedge B$ of getting herself to believe that God exists and not believing that God exists.

This leads to contradiction under two assumptions. The first assumption is that ought implies can:

> **(Ought Implies Can)** For all agents S and all options A for S, if S ought to A then S can A.

From this, it follows that S can perform $A \wedge B$.

The second assumption is that S cannot perform $A \wedge B$. This is because the combined option $A \wedge B$ involves both coming to believe that God exists and believing that God does not exist, which we might take to be impossible. After all, coming to believe that God exists entails believing that God exists, which is incompatible with the belief that God does not exist.

An immediate problem for this objection is that Unrestricted Deontic Agglomeration is ill-formed unless options also agglomerate. That is, Unrestricted Deontic Agglomeration requires:

> **(Option Agglomeration)** For all agents S and all options A, B for S, $A \wedge B$ is an option for S.

If Option Agglomeration fails, then it does not always make sense to say that S ought to $A \wedge B$, since $A \wedge B$ is not always a genuine option.

However, Option Agglomeration is not a plausible principle. Consider:

> **(Crossroads)** S is being chased by a tiger. She comes to a crossroads with two paths leading left and right. The tiger will follow her down whichever path she chooses and block her retreat.

In Crossroads, S has the option to take the left path and the option to take the right path. But S does not have the option to (take the left path and take the right path). Even if we

creatively shuffle time indices to make sense of this compound option, it is ruled out by the presence of a tiger who will block the way back to the second path. Crossroads is therefore a counterexample to Option Agglomeration.

We can remove the dependence on Option Agglomeration by restricting Unrestricted Deontic Agglomeration to cases in which Option Agglomeration holds. That is:

> **(Restricted Deontic Agglomeration)** For all agents S and all options A, B for S, if S ought to A and S ought to B *and* $A \wedge B$ *is an option for* S, then S ought to $(A \wedge B)$.

At this point, defenders of level separation have two options. First, they could lean on existing challenges to agglomeration principles, most of which would, if successful, falsify Restricted Deontic Agglomeration (Williams 1965; van Fraassen 1973; Brown 2005). Second, they could accept Restricted Deontic Agglomeration and deny that the combined option $A \wedge B$ of, say, getting oneself to believe that God exists while also believing that God does not exist is in fact an option. Readers sympathetic to existing challenges to agglomeration principles are welcome to rest content with these. Without pronouncing on these challenges, I think we can make do with the second strategy of arguing that the combined action $A \wedge B$ is not, in fact, a genuine option.

I suspect that even Rinard herself may agree with this claim, since it is hard to make sense of what the act $A \wedge B$ would be. In the case of Pascal's wager, the combined action involves both coming to believe that God exists and believing that God does not exist. But as we saw above, coming to believe that God exists entails believing that God exists, which is incompatible with the belief that God does not exist. So it is unclear that there is a genuine option described by the compound act $A \wedge B$. Indeed, if Rinard thought that there were such an option, she would need to say something to explain what taking this option involves. Rinard has not attempted this task, in my view for good reason.

And in fact, when we examine leading accounts of the metaphysics of options, most accounts hold directly that options are things that agents can do. For example, Brian

Hedden holds that options are decisions, which are necessarily things that the agent can bring about.

> **(Options-as-Decisions)** A set of propositions is a set of options for agent S at time t iff it is a maximal set of mutually exclusive propositions of the form S decides at t to ϕ, each of which S is able to bring about. (Hedden 2012, p. 352)

Similarly, Richard Jeffrey defines an option as "a proposition which is within the agent's power to make true if he pleases" (Jeffrey 1965, p. 84), and this definition is explicitly followed by many theorists (Joyce 1999; Solomon 2019). Indeed, the leading alternative picture on which options are arbitrary functions from states to outcomes (Savage 1954) has been criticized precisely on the grounds that it admits options which agents cannot perform (Broome 1991; Joyce 1999). However, if options are things that agents can do, then it follows immediately from Rinard's contention that the agent cannot $A \wedge B$ that $A \wedge B$ is not a genuine option.

Even if we remove from these accounts the stipulation that options are things that agents can do, we may nevertheless be able to make do with their contention that options are propositions.[13] On this view, the option A becomes the proposition that S gets herself to believe that God exists and the option B becomes the proposition that S believes that God does not exist. The joint option becomes the conjunctive proposition $A \wedge B$. But since coming to believe that God exists involves believing that God exists, $A \wedge B$ entails $\neg B$. As a matter of first-order logic, $A \wedge B$ entails B, which means that $A \wedge B$ expresses a contradictory proposition.

On this view, there is no option described by $A \wedge B$ any more than there is an option described by the proposition that you both sit and stand at the same time. The problem is not about what agents can do, as though some other agent might be able to both sit and stand at the same time, or both believe and disbelieve something at the same time. The

[13]Indeed, Hedden follows many theorists in requiring that *all* options be mutually exclusive, so that agglomeration over an agent's options at a given time will never produce a well-defined option. While this is a common stipulation, we do not need to accept it in order to see that the propositions involved in this case are inconsistent.

problem is rather that there is no such thing for any agent to do.

So far, we have seen that agglomeration principles are controversial, and that leading theories of options tend to imply twice over that the compound options purported to cause trouble for level separation are not genuine options. I hope that many readers will be satisfied with this line of response to agglomeration-based challenges. To those who are not, I would urge that the shoe always pinches somewhere. If this is the place where level separation pinches, then I think that the pinch is relatively mild and worth bearing when compared to the explanatory benefits of level separation.

8 Conclusion

This paper introduced and motivated the strategy of level separation (Section 2). After formulating level separation (Section 3), we saw how level separation is already applied within several areas of practical and theoretical philosophy (Section 4). We also saw how level separation can be used to give a unified, principled, and explanatorily powerful account of transitional attitudes (Section 5) and rational delay (Section 6), as well as how the view meets an agglomeration-based objection due to Susanna Rinard (Section 7).

These results join earlier applications (Section 4) of level separation in lending support to the adoption of level separation as a strategy throughout practical and theoretical philosophy. Further work might productively address applications of level separation to evaluative focal points beyond inquiry and the attitudes that inquiry produces. It might also be productive to explore the extent of analogies between the motivations for, and applications of level separation across practical and theoretical philosophy. Finally, there is scope to explore in further detail the implications of level separation for a broad range of recently proposed norms governing inquiry.

References

Adams, Robert Merrihew. 1976. "Motive utilitarianism." *Journal of Philosophy* 73:467–81.

Bales, R. Eugene. 1971. "Act-utilitarianism: Account of right-making characteristics or decision-making procedure?" *American Philosophical Quarterly* 8:257–65.

Berker, Selim. 2018a. "The explanatory ambitions of moral principles." *Noûs* 53:904–36.

—. 2018b. "The unity of grounding." *Mind* 127:729–777.

Brandt, Richard. 1959. *Ethical theory*. Prentice-Hall.

Bratman, Michael. 1987. *Intentions, plans, and practical reasoning*. Harvard University Press.

Broome, John. 1991. *Weighing goods: Equality, uncertainty and time*. Blackwell.

Brown, Campbell. 2005. "Blameless wrongdoing: A response to Streumer." *Utilitas* 17:222–5.

Cowie, Christopher. 2014. "In defense of instrumentalism about epistemic normativity." *Synthese* 191:4003–4017.

Driver, Julia. 2012. *Consequentialism*. Routledge.

Falbo, Arianna. 2023. "Inquiring minds want to improve." *Australasian Journal of Philosophy* 101:298–312.

Feldman, Richard. 2000. "The ethics of belief." *Philosophy and Phenomenological Research* 60:667–95.

Fine, Kit. 2012. "The pure logic of ground." *Review of Symbolic Logic* 5:1–25.

Firth, Roderick. 1956. "Ultimate evidence." *Journal of Philosophy* 53:732–9.

Flores, Carolina and Woodard, Elise. 2023. "Epistemic norms on evidence-gathering." *Philosophical Studies* 180:2547–71.

Fraser, Rachel. forthcoming. "The will in belief." *Oxford Studies in Epistemology* forthcoming.

Friedman, Jane. 2017. "Inquiry and belief." *Noûs* 53:296–315.

—. 2018. "Junk beliefs and interest-driven epistemology." *Philosophy and Phenomenological Research* 97:568–83.

—. 2019. "Checking again." *Philosophical Issues* 29:84–96.

—. 2020. "The epistemic and the zetetic." *Philosophical Review* 129:501–36.

Friston, Karl. 2010. "The free-energy principle: A unified brain theory?" *Nature Reviews Neuroscience* 11:127–38.

Gertken, Jan and Kiesewetter, Benjamin. 2017. "The right and wrong kind of reasons." *Philosophy Compass* 12:e12412.

Gibbard, Allan. 1990. *Wise choices, apt feelings: A theory of normative judgment*. Harvard University Press.

Goodman, Nelson. 1961. "About." *Mind* 70:1–24.

Harsanyi, John. 1977. "Rule utilitarianism and decision theory." *Erkenntnis* 11:25–53.

Hawke, Peter. 2018. "Theories of aboutness." *Australasian Journal of Philosophy* 96:697–723.

Hedden, Brian. 2012. "Options and the subjective ought." *Philosophical Studies* 158:343–360.

Hieronymi, Pamela. 2005. "The wrong kind of reason." *Journal of Philosophy* 102:437–57.

Holton, Richard. 2014. "Intention as a model for belief." In Manuel Vargas and Gideon Yaffe (eds.), *Rational and social agency: Essays on the philosophy of Michael Bratman*, 12–37. Oxford University Press.

Howhy, Jacob. 2013. *The predictive mind*. Oxford University Press.

Icard, Thomas. ms. *Resource rationality*.

Jeffrey, Richard. 1965. *The logic of decision*. University of Chicago Press.

Joyce, James. 1999. *The foundations of causal decision theory*. Cambridge University Press.

Kagan, Shelly. 2000. "Evaluative focal points." In Brad Hooker, Elinor Mason, and Dale E. Miller (eds.), *Morality, rules and consequences*, 134–55. Edinburgh University Press.

Kavka, Gregory. 1983. "The toxin puzzle." *Analysis* 43:33–36.

Kelly, Thomas. 2002. "The rationality of belief and some other propositional attitudes." *Philosophical Studies* 110:163–96.

—. 2003. "Epistemic rationality as instrumental rationality: A critique." *Philosophy and Phenomenological Research* 66:612–40.

Kelp, Christoph. 2021. *Inquiry, knowledge, and understanding*. Oxford University Press.

Kolodny, Niko. 2005. "Why be rational?" *Mind* 114:509–63.

Lewis, David. 1988. "Statements partly about observation." *Philosophical Papers* 42:1–31.

Lieder, Falk and Griffiths, Thomas. 2020. "Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources." *Behavioral and Brain Sciences* 43:E1.

Meacham, Christopher. 2015. "Understanding conditionalization." *Canadian Journal of Philosophy* 45:767–797.

Millson, Jared. 2021. "Seeking confirmation: A puzzle for norms of inquiry." *Analysis* 80:683–93.

Morton, Jennifer. 2017. "Reasoning under scarcity." *Australasian Journal of Philosophy* 95:543–59.

Na'aman, Oded. 2021a. "Emotions and process rationality." *Australasian Journal of Philosophy* 99:531–46.

—. 2021b. "The rationality of emotional change: Toward a process view." *Noûs* 55:245–69.

Palmira, Michele. 2020. "Inquiry and the doxastic attitudes." *Synthese* 197:4947–73.

Parfit, Derek. 1984. *Reasons and persons*. Oxford University Press.

—. 2011. *On what matters*, volume 1. Oxford University Press.

Podgorski, Abelard. 2016a. "Dynamic conservatism." *Ergo* 3:349–76.

—. 2016b. "Dynamic permissivism." *Philosophical Studies* 173:1923–39.

—. 2016c. *Process-oriented rationality*. Ph.D. thesis, University of Southern California.

—. 2016d. "A reply to the synchronist." *Mind* 125:859–71.

—. 2017. "Rational delay." *Philosophers' Imprint* 17:1–19.

Railton, Peter. 1984. "Alienation, consequentialism, and the demands of morality." *Philosophy and Public Affairs* 13:134–71.

Rinard, Susanna. 2019. "Equal treatment for belief." *Philosophical Studies* 176:1923–50.

Ross, Jacob and Schroeder, Mark. 2014. "Belief, credence, and pragmatic encroachment." *Philosophy and Phenomenological Research* 88:259–288.

Ross, William David. 1930. *The right and the good*. Clarendon Press.

Savage, Leonard. 1954. *The foundations of statistics*. Wiley.

Schroeder, Mark. 2012. "The ubiquity of state-given reasons." *Ethics* 122:457–88.

Schurz, Gerhard and Hertwig, Ralph. 2019. "Cognitive success: A consequentialist account of rationality in cognition." *Topics in Cognitive Science* 11:7–36.

Shah, Nishi. 2003. "How truth governs belief." *Philosophical Review* 112:447–82.

—. 2006. "A new argument for evidentialism." *Philosophical Quarterly* 56:481–98.

Simon, Herbert. 1976. "From substantive to procedural rationality." In T.J. Kastelein, S.K. Kuipers, W.A. Nijenhuls, and R.G. Wagenaar (eds.), *25 years of economic theory*, 65–86. Springer.

Singh, Keshav. 2021. "Evidentialism doesn't make an exception for belief." *Synthese* 198:5477–94.

Smith, Michael. 1994. *The moral problem*. Wiley-Blackwell.

Solomon, Toby. 2019. "Causal decision theory's predetermination problem." *Synthese* 198:5623–54.

Staffel, Julia. 2019. "Credences and suspended judgments as transitional attitudes." *Philosophical Issues* 29:281–94.

—. 2021a. "Pro tem rationality." *Philosophical Perspectives* 35:383–403.

—. 2021b. "Transitional attitudes and the unmooring view of higher-order evidence." *Noûs* 57:238–60.

—. 2025. *Unfinished business*. Oxford University Press.

—. forthcoming. "On the relationship between attitude and process norm: Comments on David Thorstad's *Inquiry under bounds*." *Philosophical Studies* forthcoming.

Steglich-Petersen, Asbjørn. forthcoming. "An instrumentalist unification of zetetic and epistemic reasons." *Inquiry* forthcoming.

Stich, Stephen. 1990. *The fragmentation of reason: Preface to a pragmatic theory of cognitive evaluation.* MIT Press.

Stocker, Michael. 1976. "The schizophrenia of modern ethical theories." *Journal of Philosophy* 73:453–66.

Thorstad, David. 2021. "Inquiry and the epistemic." *Philosophical Studies* 178:2913–28.

—. 2022. "Two paradoxes of bounded rationality." *Philosophers' Imprint* 22:1–16.

—. 2024. *Inquiry under bounds*. Oxford University Press.

—. forthcoming. "The zetetic turn and the procedural turn." *Journal of Philosophy* forthcoming.

van Fraassen, Bas. 1973. "Values and the heart's command." *Journal of Philosophy* 70:5–19.

Way, Jonathan. 2012. "Transmission and the wrong kind of reason." *Ethics* 122:489–515.

Willard-Kyle, Christopher. 2023. "The knowledge norm for inquiry." *Journal of Philosophy* 120:615–40.

Williams, Bernard. 1965. "Ethical consistency." *Proceedings of the Aristotelian Society, Supplement* 39:103–124.

—. 1981. "Persons, character and morality." In *Moral luck: Philosophical Papers 1973-80*, 1–19. Cambridge University Press.

Williamson, Timothy. 2000. *Knowledge and its limits*. Oxford University Press.

Yablo, Stephen. 2014. *Aboutness*. Princeton University Press.