Forthcoming in *Philosophical Review*

Coherence and Incoherence

Daniel Fogal (NYU) and Olle Risberg (Uppsala, IFFS, SCAS)

Penultimate version - please cite published version where possible

1. Introduction

Consider the following combinations of attitudinal mental states (henceforth 'attitudes'):

Combo 1: {belief that it is raining; belief that if it is raining, then the streets are wet; belief that the streets are wet}

Combo 2: {preference for strawberries over apples; preference for apples over cherries; preference for strawberries over cherries}

Combo 3: {intention to go for a run; belief that in order to go for a run, you have to get out of bed; intention to get out of bed}

Combo 4: {belief that it is raining; belief that if it is raining, then the streets are wet; belief that the streets are not wet}

Combo 5: {preference for strawberries over apples; preference for apples over cherries; preference for cherries over strawberries}

Combo 6: {intention to go for a run; belief that in order to go for a run you have to get out of bed; intention to not get out of bed}

Combo 7: {belief that it is raining; belief that if you drop an egg, it will break; belief that Washington D.C. is the U.S. capital}

Combo 8: {preference for strawberries over apples; preference for Star Wars over Star Trek; preference for talking with people over talking with parrots}

Combo 9: {intention to go for a run; belief that in order to make dinner you need to buy pasta; intention to call your mom}

Intuitively, in Combos 1-3 the attitudes *cohere* or *fit well together*, in Combos 4-6 the attitudes *clash* or *fit poorly together*, and in Combos 7-9 the attitudes are *random* and *unrelated*—they neither clash nor cohere. These judgments are equally clear and stable—there is no obvious reason to prioritize one set of intuitions while ignoring the

others. They also form part of a larger pattern, as similar tripartite distinctions between (something like) a "positively valenced" property, a "negatively valenced" property, and a "neutral" property arise in many other areas:

	Positively valenced	Negatively valenced	Neutral
Psychological	pleasurable	painful	neither
	happy	sad	neither
	love	hate	neither
Evaluative	good	bad	neither
	benefit	harm	neither
	beautiful	ugly	neither
Miscellaneous	rich	poor	neither
	hot	cold	neither
	tall	short	neither

In the terminology of Berker (2022, 2024), pairs of properties that conform to this pattern are *polar* opposites. In such pairs, intuitively, each property is the "inversely charged flipside" or "metaphysical mirror image" of the other (2024: 1058). These contrast with *privative* opposites, where at least one property is the mere lack or absence of the other. An object is invisible, for example, just in case it's not visible; *P* is impossible just in case it's not possible; and a chess move is illegal just in case it's not legal.

Privative and polar opposites differ with regard to the status of the "neither" verdict. In the case of privative opposites, the neither verdict is only plausible in cases that involve a presupposition failure or category mistake. The fact that two plus two equals four, for instance, might be neither visible nor invisible, but that's just because mathematical facts fall outside the realm of the physical (and hence of the visible) altogether. Similarly, prior to the introduction of legal systems, dancing was neither legal nor illegal, but that's just because the (il)legality of an action presupposes the existence of a legal system. By contrast, in the case of polar opposites, the neither verdict can also be a genuine, substantive form of assessment — what we're calling a "neutral" verdict. A particular foot massage, for example, can be pleasurable, painful, or neither; a particular painting can be beautiful, ugly, or neither; a particular person can be rich, poor, or neither; and so on.

If you think these observations border on the banal, good—we agree. But not everyone does. Here's Schopenhauer:

...pain, suffering [is] *what is positive, what is immediately felt*. The nature of satisfaction, pleasure, or happiness, by contrast, consists only in a privation's being removed, a pain's being stilled... Thus the *positive*, that which makes itself known of itself, is pain: satisfaction and pleasures are the *negative*, the mere removal of the former. (2009 [1839]: 202; italics in original)

In this passage, Schopenhauer effectively treats pleasure and pain as privative rather than polar opposites.¹ An obvious worry is that this obscures the difference between positively pleasurable states, such as the pleasure generated by an engaging conversation, and neutral states that consist in the mere absence or removal of suffering, such as the cessation of a migraine.

Perhaps surprisingly, the vast majority of recent work on coherence and (what's commonly called) structural rationality is Schopenhauerian. It focuses almost exclusively on the negative status – namely, incoherence – that Combos 4-6 have, and assumes (if only implicitly) that the lack of incoherence suffices for coherence. And this assumption is questionable: it ignores the fact that Combos 1-3 appear to enjoy a positive status – namely, coherence – that Combos 7-9 lack. As a result, nearly all recent accounts of coherence and structural rationality face a Schopenhauer-style problem: they fail to account for the difference between neutral sets of attitudes, like Combos 7-9, and positively coherent ones, like Combos 1-3.

In this paper, we have two aims. The first is to show that this problem arises for the two most prominent approaches to coherence and structural rationality, centered (respectively) around *wide-scope requirements* (§3.1) and *guaranteed failures* of some normatively significant kind (§3.2). The second is to develop a non-Schopenhauerian, *support-centric* alternative which avoids the problem (§4). The central idea is that in coherent combinations of attitudes, some attitudes support others; in incoherent combinations of attitudes, no attitudes oppose others; and in neutral combinations of attitudes, no attitude either supports or opposes any other.

We begin, however, with some preliminary remarks (§2).

2. Preliminaries

The focus of this paper is on coherence and incoherence, together with the distinctive sort of (ir)rationality that (in)coherent attitudes display—what is standardly called structural (ir)rationality.² A few remarks are thus in order about each and their relation.

¹ Nothing of substance (as opposed to fun) depends on whether we're correct in taking Schopenhauer's claims at face value. But for scholarly defense, see Fox (2023).

 $^{^{2}}$ We use parenthetical prefixes for ease of expression – '(in)coherence', for example, stands for both coherence and incoherence.

Let's start with coherence and incoherence. Although there are many different uses of 'coherence' (and its cognates), we use it to pick out the phenomenon illustrated by Combos 1-3-the salient property, whatever it is, that they have in common. Similarly, we use 'incoherence' to pick out the phenomenon illustrated by Combos 4-6-the salient property, whatever it is, that *they* have in common. Incoherence is thus the polar opposite of coherence, both of which apply to many different combinations of attitudes. Other examples of attitudes that cohere, or fit together, include believing in accordance with our assessment of the evidence, having intentions (or desires, preferences, fears, etc.) that we think we should have, hoping for an outcome that we take to be good, admiring someone that we take to be admirable, and being upset by something that we take to be bad. Other examples of attitudes that clash or are otherwise in tension include believing against our assessment of the evidence, having intentions (or desires, preferences, fears, etc.) that we think we shouldn't have, hoping for an outcome that we take to be bad, admiring someone that we take to be despicable, and being upset by something that we take to be good.³ What's more, although for simplicity we will focus on coarse-grained, "categorical" attitudes like outright beliefs and intentions, similar observations apply to various finer-grained, "graded" attitudes, such as partial belief (or confidence) and partial intention, as well as to various "inquiry-related" attitudes, such as suspension of judgment and curiosity. Being curious about whether it is true that *p*, for example, clashes with being certain that it is false that *p*, and hence is (to some degree) incoherent, as is being extremely upset by what one considers merely a minor inconvenience.

The relevant notions of coherence and incoherence are thus broad ones, encompassing a wide range of combinations of different types of attitudes that fit together (in the case of coherence) or clash (in the case of incoherence), where this extends beyond logical, probabilistic, evidential, and explanatory relations between the attitudes themselves and between their contents. What exactly it is for attitudes to fit together or clash in this sense is a difficult question, and one we'll consider at length below. Our hope, though, is that the examples above provide a sufficient grip on the relevant phenomenon to enable theorizing about it, and that they help reinforce the central observation that coherence requires more than merely the absence of incoherence: random assortments of attitudes that fail to clash do not thereby cohere.⁴

³ Some philosophers working on rationality and (in)coherence focus only on a subset of these attitudes. Staffel's (2019) sophisticated Bayesian approach, for example, focuses primarily on credences: roughly, on her view, the degree to which a set of credences is incoherent depends on its "distance" to the closest ideally rational set of credences (where the relevant notion of distance can be defined in various ways; see ch. 3). While much of what we say is compatible with much of what Staffel says, our approach has the virtue of being more general (see §4.4).

⁴ Related observations are made by Sylvan (2024) and Daoust (2023: 180-181), though neither makes the same use of the point as we do.

Our use of 'coherent' and 'incoherent' to pick out polar opposites differs from the Schopenhauerians we're concerned with in this paper, who use them to pick out privative opposites.⁵ This is because in the contemporary debates about (in)coherence and structural (ir)rationality, "[i]t is typically assumed that [what] it is for attitudes to be coherent is merely for them not to be incoherent" (Worsnip 2025: 2).⁶ Our use of 'coherent' and 'incoherent' to pick out polar opposites may thus seem to be contentious.

However, the assumption highlighted by Worsnip is just that: an assumption. No one, to our knowledge, has *argued* that what it is for attitudes to fit together in any intuitive, not-purely-technical-and-stipulative sense is merely for them not to clash, and the intuitive reflections above cast doubt on such an assumption. So although we're happy to follow others in using 'incoherence' to pick out the negative property of clashing, the use of 'coherence' to pick out its absence is clearly stipulative. If one treats coherence and incoherence as privative opposites, it's more natural to characterize incoherence in terms of the lack of coherence (as suggested by the prefix 'in-'), thereby prioritizing the positive, rather than to characterize coherence in terms of the lack of incoherence (a double negative). This accords with all the lexical entries for 'incoherence' we have seen-Merriam-Webster, for example, defines 'incoherent' simply as "lacking coherence."⁷ The alternative is akin to defining harmfulness as the absence of harmlessness rather than defining harmlessness as the absence of harmfulness. Ultimately, however, what matters most is not the terminology one uses but the phenomenon itself - in this case, the intuitive and intuitively important tripartite distinction illustrated by Combos 1-3, Combos 4-6, and Combos 7-9.8 The accounts we criticize are inadequate because they fail to capture this difference, whatever one decides to call it.

The attitudes capable of standing in relations of (in)coherence—beliefs, intentions, hopes, fears, and the like—are evaluable not only in terms of *structural* rationality, but also in terms of *substantive* rationality.⁹ Whereas structural rationality is

⁵ Berker thinks that 'coherent' and 'incoherent' express privative opposites "in everyday parlance" (2024: 1068). Whether or not he's right (we ourselves think this is at best true of some but not all everyday uses of these terms), that's compatible—as Berker himself notes—with 'coherent' and 'incoherent' being used as (semi-)technical terms to pick out polar opposites.

⁶ Cf. Kiesewetter and Worsnip (2023), who note at the outset of their *Stanford Encyclopedia of Philosophy* entry on structural rationality that they (as is standard) operate with a "fairly broad," and fundamentally negative, sense of 'coherence' that refers to "the absence of conflicts" between attitudes.

⁷ <u>https://www.merriam-webster.com/dictionary/incoherent</u> (accessed May 26, 2024).

⁸ We suspect one of the main reasons the tripartite distinction has been overlooked has been the tendency to view coherence as a generalization of the broadly logical notion of consistency, which is the privative opposite of inconsistency. But, as we're arguing, that's a mistake.

⁹ This paper assumes rather than argues for this (increasingly common) sort of "dualism" about rationality. For defense, see Fogal (2020) and Worsnip (2021).

roughly a matter of being coherent, substantive rationality is roughly a matter of responding to one's normative reasons (henceforth 'reasons').¹⁰ Slightly more fully, but still roughly: attitudes are substantively rational to the extent that they're supported by one's reasons, whereas attitudes are structurally rational to the extent that they cohere or fit together, independently of whether and to what extent they're supported by one's reasons.¹¹ (We'll say more about reasons in §4.1.) Whereas attitudes like beliefs, intentions, and so on are rationally evaluable in both these ways, not all mental states are. Feelings of pain or hunger, for example, are not rationally evaluable in either way, as they are not the sort of states that can stand in coherence relations or that we can have reasons for or against.

What's the relationship between (in)coherence and structural (ir)rationality? One option is to follow Worsnip (2021) and take facts about structural (ir)rationality to be grounded in facts about (in)coherence. On this view, attitudes are structurally (ir)rational *in virtue of* their (in)coherence. Another possibility is that the relationship between (in)coherence and structural (ir)rationality has a "common cause" structure: facts about both are grounded in facts about some third thing. Our own view, developed in §4, takes this latter form: facts about (in)coherence and facts about structural (ir)rationality are both grounded in facts about structural support (and structural opposition—cf. §4.2). On either approach, however, (in)coherence and structural (ir)rationality typically co-occur.

A final preliminary note: we set aside questions concerning the so-called "normativity" of structural rationality.¹² As will become apparent, on our view, the fact that a set of attitudes is structurally (ir)rational doesn't entail anything in particular about normative reasons; for instance, it doesn't guarantee that there are reasons (not) to have those attitudes. But we don't take this to undermine the normative significance of structural rationality – the assumption that the significance of everything genuinely significant can be understood in terms of reasons is non-trivial, and in our view questionable.¹³ However, we lack the space to properly address this issue.

¹⁰ We will set aside questions concerning whether the attitudes in question are "properly based," along with diachronic processes more generally (see instead Fogal and Risberg ms for discussion of such issues). We're thus concerned with (synchronic) "ex ante" rationality as opposed to "ex post" rationality, though we realize that talk of "responding to reasons" suggests the latter.

¹¹ More accurately: attitudes are substantively rational to the extent that they are supported by one's *evidence-relative* reasons. For the need to distinguish an intermediate, "evidence-relative" notion of reasons from the more objective, "fact-relative" notion as well as the more subjective, "attitude-relative" notion, see Fogal and Worsnip (2021).

¹² For helpful overviews, see Kiesewetter (2017) and Worsnip (2021).

¹³ For critical discussion, see Worsnip (2021: 29–30). For critical discussion of the idea that reasons are normatively fundamental, see Fogal and Risberg (2023a; 2023b).

3. Requirement- and Guarantee-Based Views

The current debate over the nature of coherence and structural rationality is complicated and, at times, convoluted. We won't attempt to provide a systematic overview – for that, see Kiesewetter and Worsnip (2023). Instead, we'll argue that the two main types of views proposed in the recent literature are Schopenhauerian and therefore problematic.¹⁴ Although the critical discussion is not intended to be exhaustive or conclusive, it should suffice to motivate the alternative, non-Schopenhauerian view developed in the following section.

According to (what we'll call) the *requirement-centric view*, structural rationality corresponds to a distinctive set of rational requirements. On this view, a set of attitudes is structurally irrational if it violates those requirements and structurally rational otherwise.¹⁵ Although the exact content, nature, and form of the relevant requirements is subject to controversy, most of the details won't matter—our main objection will apply regardless of how they are spelled out. One exception, however, concerns the "scope" of the requirements. On the *wide scope* (or, rather, *wide-scope-in-spirit*¹⁶) view, structural rationality prohibits incoherent combinations of attitudes, and that's it. Structural rationality is thus silent (at least typically¹⁷) about which individual attitudes you may or ought to have—it only requires that you *not* have incoherent combinations of attitudes. This stands in contrast to the *narrow scope* view, according to which structural requirements don't just ban incoherent combinations of attitudes, but also (often if not always) single out particular attitudes as required. Our discussion will focus on the currently-dominant wide scope version of the requirement-centric view, as the narrow scope view has largely fallen out of favor.¹⁸

The second type of view we'll consider is *guarantee-based*. Guarantee-based views posit a close connection between incoherence and the guaranteed lack of some other, broadly normative property. The most well-known version appeals to *reasons*. In its simplest form, it says that in any set of incoherent attitudes, at least one attitude is

¹⁶ See Worsnip (2021: §6.3) for discussion.

¹⁴ For an interesting alternative view that we lack space to discuss, see Lasonen-Aarnio (2021).

¹⁵ For book-length defenses, see Broome (2013) and Worsnip (2021). More specific requirementcentric views are commonly assumed in formal work on rationality, such as in decision theory and Bayesian epistemology. Their popularity is perhaps best illustrated by the fact that there exists a *PhilPapers* category called 'Rational Requirements'.

¹⁷ We say 'typically' because some think it's possible (albeit rare) for individual attitudes to be structurally irrational (cf. Broome 2013: 153; Worsnip 2021: 17).

¹⁸ For a survey of objections to the narrow scope view, see Worsnip (2021: ch. 6). It's an underappreciated virtue of the narrow scope view, however, that it can capture the tripartite distinction between coherence, incoherence, and neutral sets of attitudes. We take our own view to accommodate the insights motivating the narrow scope view while avoiding its problems – see, e.g., footnote 72.

guaranteed to be an incorrect response to the subject's reasons (cf. Kolodny 2007, Kiesewetter 2017, Lord 2018).¹⁹ We'll argue against these views in §3.2.

3.1 Wide-Scope Requirements

The wide scope view consists of two main claims: (a) structural rationality corresponds to a distinctive set of requirements and (b) these requirements prohibit incoherent combinations of attitudes, and that's it. The core idea is thus that for each kind of incoherent combination of attitudes there exists a corresponding requirement of structural rationality which prohibits it, and combinations of attitudes that violate these requirements are structurally irrational.²⁰ Consider, for example:

Modus ponens consistency	Structural rationality prohibits one from {believing
	that <i>p</i> , believing that if <i>p</i> then <i>q</i> , believing that not- <i>q</i> }.
Acyclicality	Structural rationality prohibits one from {preferring
	X to Y, preferring Y to Z, preferring Z to X $\}$.
Means-end consistency	Structural rationality prohibits one from {intending
	to <i>E</i> , believing that <i>E</i> -ing requires <i>M</i> -ing, intending to
	not- <i>M</i> }. ²¹

Requirements like these are then appealed to in explaining our judgments about particular (kinds of) cases. Recall, for example:

Combo 4: {belief that it is raining; belief that if it is raining, then the streets are wet; belief that the streets are not wet}

Combo 5: {preference for strawberries over apples; preference for apples over cherries; preference for cherries over strawberries}

Combo 6: {intention to go for a run; belief that in order to go for a run you have to get out of bed; intention to not get out of bed}

¹⁹ Some philosophers (e.g. Kiesewetter 2017) defend related but slightly less simple versions of the guarantee-based view. The differences won't matter for our purposes.

²⁰ For some additional nuances and qualifications, see especially Worsnip (2021: ch. 6).

²¹ Although it is more common to formulate the modus ponens-like requirement and the meansend requirement as "closure" principles (prohibiting one from believing that p and believing that if p then q while not believing that q, and prohibiting one from intending to E and believing that E-ing requires M-ing while not intending to M), we have chosen to frame them as "consistency" requirements instead. The reason is that there are various worries about closure requirements that do not generalize to their corresponding consistency requirements, and our arguments against the wide scope view go through no matter which option one chooses here. See Worsnip (2021: ch. 9.4) and Fullhart and Martinez (2024: §5) for discussion of the modus ponens requirement and Broome (2013: ch. 9.4) for discussion of the means-end requirement.

Each of these combinations of attitudes is structurally irrational, and wide-scopers have a story about why: Combo 4 violates *Modus ponens consistency*, Combo 5 violates *Acyclicality*, and Combo 6 violates *Means-end consistency*.

Our main objection to the wide scope view is simple: it is Schopenhauerian.²² Regardless of the merits of its verdicts about Combos 4-6, no wide scope requirements have been proposed to explain the intuitive difference between Combos 1-3 and Combos 7-9. And given the nature of wide scope requirements as mere prohibitions of incoherent combinations of attitudes, it's doubtful that such an explanation is possible – neither positively coherent combinations of attitudes nor random, unrelated ones are incoherent, and so neither will be prohibited by such requirements. Absent supplementation, then, the wide scope view is committed to viewing Combos 1-3 and Combos 7-9 as structurally on a par – both count as structurally rational in the same sense (they are not structurally irrational) and for the same reason (they don't violate any wide scope requirements). And insofar as structural rationality and coherence are closely related, as is standardly assumed, Combos 1-3 and Combos 7-9 are also on a par with respect to their coherence. That's wrong.

The argument just made can be summarized as follows:

- P1. If the wide scope view is true, then Combos 7–9 are structurally rational and/or coherent in the same way as Combos 1-3.
- P2. But Combos 7–9 are not structurally rational and/or coherent in the same way as Combos 1–3.
- C. So the wide scope view is false.

In response, wide-scopers might acknowledge that their view fails to capture the positive notion of coherence but nonetheless insist there is a distinct, weaker notion of coherence – one which only requires the absence of incoherence – that it does capture.²³ Call the former *strong* coherence and the latter *weak* coherence. While the argument above is sound if 'coherent' is understood in the strong sense, it's unsound if 'coherent' is understood in the strong sense, it's unsound if 'coherent' is false.

Even if one grants the notion of weak coherence, however, we also need an account of strong coherence. As noted at the outset, there seems to be no good reason to prioritize one set of judgments concerning strong coherence, incoherence, and the absence of each over the others. And it is reasonable to expect the correct account of strong coherence to be fundamentally similar to that of incoherence, since polar opposites typically share important similarities (despite, of course, being opposites). Both tallness and shortness, for instance, are naturally understood in terms of a shared

²² For a survey of other objections to the wide scope view, see Kiesewetter (2017); for a survey of responses, see Worsnip (2021).

²³ This reply was suggested by Alex Worsnip (personal communication).

dimension of comparison: height. Similarly, while harm and benefit differ in valence, they are both commonly understood in terms of (oppositely-valenced) impact on one's well-being. As far as we can tell, similar relationships obtain between all, or nearly all, polar opposites. So given that strong coherence and incoherence are polar opposites, we should expect them to conform to this pattern. Thus, if the wide scope view is incorrect as a theory of strong coherence, it is likely incorrect as a theory of incoherence as well.

A further problem for the wide scope view arises given dualism about rationality, which we assume (§2). Worsnip (2021), for example, is a dualist who explicitly accepts the wide scope view of structural rationality and pairs it with a version of the standard reasons-based view of substantive rationality. On this combination of views, a variety of significant differences between substantive and structural rationality arise, including the following:

- (1) Structural rationality is *formal*—"we can point to patterns of structural irrationality, using variables and schematic symbols, and without even fully specifying the content of the attitudes in question" (Worsnip 2021: 8)—while substantive rationality is not.
- (2) Structural rationality is *fundamentally negative* (it merely prohibits), while substantive rationality is not.
- (3) Structural rationality primarily concerns *combinations* of attitudes, while substantive rationality does not.
- (4) Substantive rationality is *graded*, in the sense that "[t]here are more and less severe failures to respond to one's... reasons" (Worsnip 2021: 47), while structural rationality is not.²⁴

On Worsnip's view, then, structural and substantive rationality are not just distinct and "equally genuine" kinds of rationality (2021: ix), with neither being reducible to the other – they are also *radically different*. Call this *Cartesian* dualism about rationality.

Cartesian dualism faces an immediate worry: given the radical differences between substantive and structural rationality, what makes them both kinds of *rationality*?²⁵ It can't just be a terminological accident or a simple case of homonymy, as it is with 'bank' being used to pick out both riverbanks and financial institutions. Worsnip's dualism about rationality is metaphysical, not merely linguistic. One might therefore expect there to be a story about why structural and substantive rationality count as forms of the same general kind of thing—rationality. It's one thing for distinct domains of evaluation (or anything else) to share a name or label; it's another thing for

²⁴ The assumption that structural rationality is non-graded is common but not strictly essential to the wide scope view; for discussion, see Fogal (forthcoming).

²⁵ For versions of this worry, see Way (2022), Kiesewetter (2024), and Singh (2025).

them to deserve it. Such a story is standardly available in cases of polysemy (i.e., when a word or concept has multiple related meanings) and when multiple species belong to the same genus, so we should expect one in the present case too.

Our support-centric view, developed in §4, is a non-Schopenhaurian, non-Cartesian form of dualism: substantive and structural rationality are distinct but not radically different. They share a variety of important features and they obtain in virtue of the same kinds of facts (namely, facts about support and opposition). It therefore ends up being no mystery that substantive and structural rationality are both kinds of rationality.²⁶

3.2 Guarantee-Based Views

Guarantee-based views offer a different kind of explanation of the irrationality of incoherent attitudes. We'll begin by considering the most prominent view, which appeals to (normative) reasons, before turning our attention to guarantee-based approaches more generally.

To illustrate the central idea behind the reasons-based version of the view, suppose that a subject S believes both that p and that not-p. According to adherents of the reasons-based view, it is impossible that S's reasons sufficiently support both these beliefs: either their belief that p is insufficiently supported by their reasons or their belief that not-p is. Thus, either way, S is irrational in virtue of having a belief that their reasons do not sufficiently support.

The reasons-based view seeks to generalize this style of explanation to all (or nearly all) cases of incoherence (cf. Kolodny 2007, Kiesewetter 2017, Lord 2018). This generalization relies on the following hypothesis ('RB' for 'reasons-based'):

The Guarantee Hypothesis_{**RB**}: If a set of attitudes $\{A_1...A_n\}$ is incoherent, then it is guaranteed that one of $A_1...A_n$ is an incorrect response to the subject's reasons.

This hypothesis (or some variation thereof) can then be used to provide an explanation of the irrationality of incoherent attitudes that is friendly to monism about rationality (cf. §2).²⁷ For if the Guarantee Hypothesis_{RB} is true, one might think, then there's no

²⁶ Cf. Singh (2025), who aims to explain "how rationality might ultimately be unified even in absence of a reduction (or elimination) of one kind of rationality to the other" (2). Even if Singh's view succeeds in being non-Cartesian, however, his account of structural rationality is Schopenhaurian (see §3.2 below).

²⁷ Note that neither Kiesewetter nor Lord put forward their guarantee hypotheses as accounts of what incoherence as such *is*. To address that question, which is the one that we're mainly interested in, they would have to endorse their hypotheses as at least giving both necessary and

need to posit a further, fundamentally distinct type of irrationality over and above that of failing to respond correctly to one's reasons. Instead, the intuitive irrationality of incoherent attitudes can be fully explained using the same basic material as the irrationality of attitudes that one's reasons don't sufficiently support.²⁸

An immediate problem for the Guarantee Hypothesis_{RB}, which we mention only to set aside, is that it faces counterexamples. One familiar kind concerns attitudes which are permitted but not required. If Selim's reasons to intend to drink coffee are equally as strong as his reasons to intend not to drink coffee, for example, then neither intention (considered on its own) is an incorrect response to his reasons. But he still seems incoherent, and irrational, if he has both intentions.²⁹

For present purposes, the more pressing problem is that the Guarantee Hypothesis_{RB} leads to a Schopenhauerian view of (in)coherence and (ir)rationality. As it stands, it only concerns incoherence, and hence is silent about the difference between sets of attitudes which are (strongly) coherent, like Combos 1-3, and those which are neutral, like Combos 7-9. Since neither are incoherent, the Guarantee Hypothesis_{RB} says nothing about the difference between them. It's possible for all the attitudes in those combinations to be sufficiently supported by the subject's reasons, and it's also possible for them not to be.³⁰

How might proponents of reasons-based views respond? One option is to grant that the Guarantee Hypothesis_{RB} only explains part of the phenomenon, but insist that its explanation of that part is nonetheless correct. For example, they might emphasize the distinction between strong and weak coherence and only claim to capture the weak notion of coherence (cf. §3.1). As noted above, however, when *F* and *G* are polar opposites, we should expect the correct account of *F* to mirror the correct account of *G*. Accordingly, if reasons-based views of strong coherence are false, we should expect reasons-based views of incoherence to be false as well.

Proponents of reasons-based views might instead try to generalize the Guarantee Hypothesis_{RB} so that it provides an explanation of the difference between strongly coherent attitudes and neutral ones. The problem, however, is that no such generalization seems plausible. The most straightforward suggestion would be that in any strongly coherent set of attitudes, at least one attitude is guaranteed to be a *correct*

sufficient conditions. They would also have to deny that they admit of exceptions (Lord, for example, thinks there are cases of rational incoherence (2018: ch. 2.5)).

²⁸ We say 'fully explained' because other views, including ours, can grant that (something like) the Guarantee Hypothesis_{RB} might be part of what underlies the intuitive irrationality in cases of incoherence. They just deny that it's the full explanation.

²⁹ See Lord (2018: ch. 2.4.2) for an attempted solution, and Risberg (2019) for criticism.

³⁰ Both Kiesewetter (2017) and Lord (2018) focus on incoherence and don't say much about how they understand coherence. Lord's arguments make clear, though, that he treats it as the privative opposite of incoherence.

response to the subject's reasons. But this suggestion clearly fails. Suppose, for instance, that *S* intends to cause themselves pain, believes that in order to cause themselves pain they must cause someone else pain, and intends to cause someone else pain. These attitudes are strongly means-end coherent—and not merely not means-end incoherent—and yet *S* may well lack sufficient reasons for any of them. (In particular, they might lack sufficient epistemic reasons for the belief and lack sufficient prudential and moral reasons for the intentions.) Other straightforward ways of generalizing the Guarantee Hypothesis_{RB} to explain the difference between strongly coherent attitudes and neutral ones run into similar problems.

In sum: while Worsnip's view is too Cartesian (as it makes substantive and structural rationality too different), the reasons-based view is not Cartesian enough, as it ties (in)coherence and structural (ir)rationality too closely to responding to one's reasons. What's more, both are Schopenhaurian – and problematically so.

There has recently been a raft of other guarantee-based proposals concerning the notion of (in)coherence at the heart of structural (ir)rationality. Though differing in details, they share a common focus and form. In terms of focus: the accounts provided are accounts of *in*coherence, rather than (strong) coherence, and insofar as coherence is mentioned at all it's treated merely as the absence of incoherence. This Schopenhauerian shortcoming is hardly surprising, given the history of the debate and its attendant negativity bias. But it's still regrettable, since there is no obvious way to modify or extend these accounts of incoherence to capture the positive (or "strong") notion of coherence, and hence the other (positive) half of structural rationality. That's partly because of their form: they all claim that what characterizes incoherent sets of attitudes is that in any such set, at least one attitude is guaranteed to lack some broadly normative property, such as being "satisfied" or "successful."³¹ Consider the following sample of recent (somewhat simplified) accounts of incoherence:

A set of attitudes $\{A_1...A_n\}$ is incoherent if and only if it is guaranteed that ...

- (a) not all of $A_1...A_n$ are functionally successful. (cf. Langlois 2014)³²
- (b) not all of $A_1...A_n$ have their constitutive aims satisfied. (cf. Brunero 2020)³³

³¹ See for example Fullhart and Martinez (2024), who hold that "coherence is fundamentally a matter of whether it's logically possible for one's attitudes to satisfy some condition" (2024: 316). ³² According to Langlois, the function of a belief is "to represent the way things are" while the function of an intention is "to change the way things are." So a belief succeeds, functionally speaking, just in case it's true, while an intention succeeds just in case it makes its object true (2014: 118–119).

³³ Brunero proposes this as an explanation of means-end incoherence only, as he thinks that there are other types of incoherence (including intransitive preferences and akrasia) which this view does not capture (2020: 205–206).

- (c) not all of $A_1...A_n$ amount to knowledge. (cf. Rosa 2022)³⁴
- (d) not all of $A_1...A_n$ are successful. (cf. Fink 2023)³⁵
- (e) not all of $A_1...A_n$ are satisfied. (cf. Fullhart and Martinez 2024)³⁶
- (f) someone who has all of $A_1...A_n$ fails to react to a question in the way she
- is committed to in virtue of having those attitudes. (cf. Lee 2024)³⁷
- (g) at least one of $A_1...A_n$ is an "attitudinal failure." (cf. Singh 2025)³⁸

All these views can be formulated as instances of the following:

Guarantee-based views of incoherence: A set of attitudes $\{A_1...A_n\}$ is incoherent if and only if it is guaranteed that not all of $A_1...A_n$ are *F*.

How could such views distinguish between strongly coherent sets of attitudes and neutral ones? To respect the appearance that coherence and incoherence are polar opposites, the most natural possibility is to endorse one of the following views of strong coherence:

Guarantee-based views of coherence #1: A set of attitudes $\{A_1...A_n\}$ is (strongly) coherent if and only if it is guaranteed that *at least one* of $A_1...A_n$ is *F*.

Guarantee-based views of coherence #2: A set of attitudes $\{A_1...A_n\}$ is (strongly) coherent if and only if it is guaranteed that *all* of $A_1...A_n$ are *F*.

One could then claim that a set of attitudes is neutral just in case it's neither coherent nor incoherent.

The problem, however, is that for any candidate property F that one might plausibly appeal to in a guarantee-based view of incoherence, a corresponding guarantee-based view of strong coherence is implausible. Consider again the set {intention to cause oneself pain; belief that in order to cause oneself pain one must cause

³⁴ Rosa proposes this as an account of when doxastic states are incoherent. He also sketches a related more general guarantee-based account of practical incoherence, on which what characterizes such sets of attitudes is that they cannot all be at their "best" (2022: §6).

³⁵ Fink understands the success of an attitude in terms of its constitutive aims (2023: §3).

³⁶ Fullhart and Martinez understand the satisfaction of an attitude in terms of whether there is a "fit" between the attitude and the world (2024: §4).

³⁷ Lee takes the relevant attitudes to be constituted by commitments to *act* (perhaps mentally), where a commitment is constitutive of an attitude just in case "it is partly *in virtue of* incorporating the commitment that the attitude is what it is" (2024: 625).

³⁸ Roughly, according to Singh, an attitude "fails" in the relevant sense just in case its constitutive correctness condition is not satisfied or it's substantively irrational.

someone else pain; intention to cause someone else pain}. Despite being strongly coherent (and not just not incoherent), it's false that at least one attitude in this set is guaranteed to be successful, or have a satisfied constitutive aim, or the like. A fortiori, it's false that *all* attitudes in the set are guaranteed to have the relevant desirable property. And if guarantee-based views fail to account for strong coherence, that raises significant doubts about guarantee-based views of incoherence as well.³⁹

4. The Support-Centric Approach

The support-centric view of rationality we favor is a non-Schopenhauerian and non-Cartesian form of dualism. It is non-Schopenhauerian since it doesn't treat coherence as simply the absence of incoherence: on our view, Combos 1-3 are (strongly) coherent, Combos 4-6 are incoherent, and Combos 7-9 are neither coherent nor incoherent. It is dualistic since it treats substantive rationality and structural rationality as distinct, genuine forms of rationality. Yet it is non-Cartesian since it doesn't entail that substantive and structural rationality are radically different – they each correspond to a distinct kind of *rational support* or (more generally) *rational pressure*.

In this section we'll begin by sketching the support-centric view of normative reasons and substantive rationality that we've defended elsewhere (Fogal and Risberg 2023a; 2023b), and that we'll be taking for granted here (§4.1). We'll then show how facts about support can be helpfully modeled using directed graphs (§4.2), which provides the basis for our support-centric view of structural rationality (§4.3) and

³⁹ An alternative proposal (suggested in personal communication by Samuel Fullhart and Camilo Martinez) is that, roughly put, a set of attitudes is strongly coherent iff it is guaranteed that, if some attitudes in the set are *F*, then the other attitudes in the set are also *F*. We have doubts about this approach-besides worries about extensional adequacy, it would give incongruent accounts of polar opposite properties if paired with the guarantee-based approach to incoherence. In response to this, an anonymous reviewer suggests that the guarantee-based view of incoherence could be reformulated so as to say that a set of attitudes is incoherent to some extent iff it contains two subsets of attitudes such that it is guaranteed that if each attitude in the first subset is F, then each attitude in the second subset is not F. This could then be combined with a parallel guarantee-based view of strong coherence, saying that a set of attitudes is strongly coherent to some extent iff it contains two subsets of attitudes such that it is guaranteed that if each attitude in the first subset is *F*, then each attitude in the second subset is F. While we have doubts about this approach too, we think it provides the most promising strategy for proponents of guarantee-style views, even though it abandons the guarantee-based view as usually stated. In particular, to our minds, the best version of this view is likely one that replaces 'F' with 'substantively supported' (and which is thus closely related to the nonguarantee-based view that we defend in §4). However, such a view would still face problems, for example concerning necessarily supported attitudes (see further footnote 62).

(in)coherence (§4.4). We'll conclude by showing how this view of (in)coherence makes correct predictions about the cases presented in the introduction (§4.5).

4.1 Substantive Pressure and Substantive Rationality

The slogan that substantive rationality is a matter of responding to one's reasons is just that: a slogan. A slightly more accurate slogan is that substantive rationality is a matter of one's attitudes being supported by *the overall balance* of reasons. The overall balance of reasons, in turn, is standardly thought to be determined by the *weight* of one's reasons. So whether one's attitudes are substantively rational depends not only on one's reasons, but also on their weight.

Our account of substantive rationality respects these slogans but doesn't stop with them. Instead of treating facts about reasons or their weight as primitive, we analyze both in terms of a third notion: *normative support*. On the view we favor, reasons are "sources" or "providers" of normative support, and the weight of a reason is simply a matter of how much support it provides: strong reasons provide lots of support, weak reasons provide only a little support, and so on. This view, which we call the *Sources of Support* (S.O.S.) view, can be formulated as follows (where *r* is a fact, *S* is a subject and *A* is an attitude⁴⁰):

- **S.O.S.**_{*R*} For *r* to be a reason for *S* to *A* is for *r* to provide (a particular amount of⁴¹) normative support for *S*'s *A*-ing.
- **S.O.S.**_{*W*} The weight of a reason, r, for S to A corresponds to the amount of normative support that r provides for S's A-ing.

To illustrate, suppose that the fact that the zoo is fun is a (strong) reason for Jens to intend to go there. On our view, this is because the fact that the zoo is fun provides (a lot of) support for Jens's intending to go there.⁴² Assuming that there are no other normatively relevant considerations, then, it is substantively rational for Jens to intend to go to the zoo. The weight of a reason is thus treated as a scalar quantity, like mass and speed, which are completely described by their magnitudes (in this case, *the amount*

⁴⁰ Elsewhere, we defend the S.O.S. view as an account of reasons for both actions and attitudes. Here we focus on reasons for attitudes, since sets of attitudes are standardly taken to be coherent and incoherent.

⁴¹ As Broome (2025) notes, 'amount' is ambiguous: "Suppose you have two bottles of wine, each containing one pint. In one sense of 'amount', they each contain the same amount of wine. In another sense, they contain two different amounts of wine" (6). When precision on this point is needed, we use 'particular amount' to express the second sense (Broome uses 'portion' instead). ⁴² Although S.O.S._R is stated as an analysis of what it is to be a reason, we take it to entail a corresponding grounding claim. (Compare: if to be a vixen is to be a female fox, then plausibly, one is a vixen in virtue of being a female fox.)

of support provided). Support, by contrast, is akin to vector quantities like force and velocity, as it has both magnitude (*amount* of support) and direction (support *for*).⁴³

A variety of different distinctions concerning reasons have been proposed, and similar distinctions can be drawn with respect to support. The S.O.S. view is neutral, for example, concerning how many different normative domains there are (e.g., whether there are distinct forms of moral support, prudential support, epistemic support, etc., grounding facts about moral reasons, prudential reasons, epistemic reasons, etc., respectively) and how many perspectival kinds of support there are within each domain (e.g., whether we should distinguish between fact-relative support, evidence-relative support, etc., which in turn ground fact-relative reasons, evidencerelative reasons, etc., across normative domains). In what follows, we'll use the terms 'normative reasons' and 'normative support' to include all these types of reasons and support, while using 'substantive reasons' and 'substantive support' to denote the types of reasons and support that are relevant to substantive rationality.

We understand the notion of provision in S.O.S._{*R*} and S.O.S._{*W*} in terms of the noncausal, non-pragmatic kind of explanation often called 'metaphysical' or 'grounding' explanation. Roughly, for *r* to provide a particular amount of normative support for *A* is for *r* to ground—and hence metaphysically explain—the fact that there is a corresponding particular amount of normative support for *A*. The S.O.S. view is thus a form of 'explanationism' about reasons, as it analyzes facts of the form [*r* is a reason for *S* to *A*] partly in terms of explanation-facts.

We assume that reasons are facts, and we'll usually talk as though they are nonmental facts – for instance, the fact that the streets are wet might be a reason for *S* to intend to take an umbrella. However, these assumptions are optional, and the S.O.S. view could be adapted to other views about reasons. For example, one might take substantive reasons to be (facts about) experiences (e.g., that it appears to *S* that the streets are wet) or knowledge (e.g., that *S* knows that the streets are wet). What matters for our purposes is that substantive reasons are not – or at least not generally – identical to (or constituted by or provided by or...) "mere" attitudes, such as beliefs and intentions, or by bare facts about such. For example, the mere fact that *S* intends to φ is,

⁴³ While the common-sensical notion of weight is a scalar notion, there are also vector notions of weight. In science, for example, weight is commonly defined as the force exerted on a body by gravity (though a scalar notion of weight as the magnitude of the gravitational force is also used). Note also that like forces and other abstract phenomena, normative support is not naturally thought of as a kind of "stuff," such as water or sand. At most, normative support can be said to have some "stuff-like features" – e.g., that there can be more or less of it – but the same is true of many other phenomena denoted by mass nouns, such as knowledge and wisdom, which are not intuitively stuff. This observation helps address some of Broome's (2025) criticisms of the S.O.S. view.

we assume, not in general a substantive reason for *S* to φ . This will matter for the way we distinguish between structural and substantive rationality (see §4.3 below).

It is important to distinguish between the verb 'support(s)' and the mass noun 'support' (each in their normative senses). The S.O.S. view is stated in terms of the mass noun. But there is obviously a close connection between the verb and mass noun, and as we argue elsewhere (Fogal and Risberg 2023a), the former is plausibly understood in terms of the latter: for *r* to normatively support *A* just is for *r* to provide normative support for A. The distinction between the verb and mass noun is subtle but significant – it is an instance of the more general distinction between the provision of something (e.g., an activity, process, or event that takes place) and that which is provided (e.g., an object or quantity of something). Compare: lights are sources of light, and the light (mass noun) a particular light (count noun) provides is distinct from the particular light (count noun) itself. The same goes for the relationship between pleasures (count noun) and pleasure (mass noun), sorrows (count noun) and sorrow (mass noun), and so on. What's more, the verb '(to) support' (in its normative sense) is closely related to the count noun 'reason(s)' (in *its* normative sense): to say that rsupports A-ing is essentially equivalent to saying that r is a reason to A, and both are essentially equivalent to the claim that r "favors" or "counts in favor of" A-ing. It should come as no surprise, then, that the count noun 'reason(s)' can similarly be analyzed in terms of the mass noun 'support': for *r* to be a reason to *A* is for *r* to provide, or be a source of, support for A-ing.

Though we use 'normative support' as a technical term, what it denotes is familiar. In ordinary thought and talk, the notion of normative support is often expressed by the mass nouns 'reason' and 'justification'. The support-centric view of reasons and their weight could thus be summarily put as follows: to be a reason is to provide reason, and the weight of a reason is the amount of reason it provides. (If these claims seem near-platitudinous, we take that to speak in their favor.) The popular view that one ought to do what there is most reason to do can be understood as expressing the idea that one ought to do what there is most support for doing. And the slogan that substantive rationality depends on the balance of (substantive) reasons can likewise be understood in support-centric terms: substantive rationality is a matter of having attitudes that there is on-balance substantive support for it.⁴⁴

What should one say about attitudes that are neither substantively supported nor substantively opposed, such as (let's suppose) intending to hum a random tune on

⁴⁴ We grant that the slogan may require modification in order to accommodate subtleties concerning the weight of reasons. For example, if there can be cases in which there is on-balance substantive support for intending to *A* and even more on-balance substantive support for intending to *B* (and one cannot both *A* and *B*), it might be considered irrational to intend to *A*. (Thanks to an anonymous reviewer for raising this possibility.)

one's way to work? One option is to say that such attitudes are substantively rational (because they are not substantively opposed); another option is to say that they are substantively irrational (because they are not substantively supported); a third option is to say that they are neither substantively rational nor substantively irrational. We think the best option is to say that such attitudes are substantively rational in a "weak" sense, as they are not substantively opposed (and hence not substantively irrational), but not in a "strong" sense, as they are not substantively supported – there is nothing of substance to be said in their favor, but also nothing of substance to be said against them. Similar distinctions between stronger and weaker senses arise with respect to many other broadly deontic terms. In deontic logic, for example, it is common to distinguish between weak (or "negative") permission, which is just a matter of not being prohibited, and strong (or "positive") permission, which is roughly a matter of being positively licensed (cf. von Wright 1963).⁴⁵ We can accordingly say that attitudes which are neither substantively supported nor substantively opposed are substantively rational in the sense of being weakly rationally permissible and deny that they are substantively rational in any stronger sense.

4.2 Modeling Substantive Support

As noted above, normative support is force-like. In particular, it is *directed*, has *magnitude*, and is *defeasible*, *transmissible*, and *separable* from its source:

- **Direction**: normative support is always support *for* some attitude.
- **Magnitude**: the strength and/or amount of normative support can vary (i.e., there can be more or less support for *A*), and so it is meaningful to ask *how strong* the normative support is (e.g., the support can be strong or weak) or *how much* support there is (e.g., there can be a lot or very little support for *A*).
- **Defeasibility**: normative support is capable of being *opposed* and *undermined*, and is in this sense "pro tanto." There might be some support for *A*, for example, but more support for some other attitude that is an alternative to *A*.
- **Transmissibility**: normative support can *transmit* or *transfer* from one (group of) attitude(s) to another (i.e., there can be normative support for the latter *in virtue of* there being normative support for the former).
- **Separability**: the normative support that a fact provides is distinct from the fact itself.

Direction, magnitude, defeasibility, transmissibility, and separability are all features of forces. It is thus unsurprising that facts about normative support—and hence facts about substantive support—can be naturally modeled using directed graphs, which are

⁴⁵ For a similar weak/strong distinction with respect to epistemic possibility (e.g. 'can') claims, see Przyjemski (2017). Compare also the weak/strong distinction regarding 'coherence' in §3.1.

often deployed to model force-like phenomena such as causal forces (cf. Pearl 2000). As an illustration, suppose we want to model a case in which (i) a subject *S* believes that *p*, (ii) *S* believes that *p* or *q*, and (iii) there is substantive support for their belief that *p*, and as a result there is substantive support for their belief that *p* or *q* as well. The following (pictorial representation of a^{46}) directed graph helps us represent some of these facts ('B' for belief):



Here and in what follows, the nodes in a graph represent attitudes held by a subject at a time – in this case, beliefs and their contents – and the arrows represent facts about (pro tanto) substantive support.⁴⁷ Thus, Figure 1 tells us that the subject believes that p and that p or q, and that substantive support for the former transmits to the latter, where substantive support for an attitude A "transmits" to an attitude B just in case the fact that there is substantive support for B obtains in virtue of the fact that there is substantive support for A. (Note that while there may be various psychological facts about the subject corresponding to these normative ones – for example, that her belief that p or q is based on her belief that p—what the arrow represents is not itself psychological.) Importantly, the source of support is not the subject's belief itself, since (we're assuming) merely believing that p does not, on its own, provide substantive support for believing that p or q. Instead, her belief that p or q is substantively supported in virtue of the fact that the subject's belief that p or q is not the subject for p or q. Instead, her belief that p or q is substantively supported in virtue of the fact that the subject's belief that p or q, where channeling is *not* provision but is instead understood in terms of

⁴⁶ As we'll mostly use the term 'graph', a graph is a certain kind of pictorial figure. Sometimes, a graph is instead defined as a certain abstract object; specifically, as an ordered pair consisting of a set of objects (intuitively, the nodes in the graph) and a set of ordered pairs of those objects (intuitively, the arrows in the graph). So understood, the figures we present below are not graphs but representations of them. Some figures we introduce below also require minor departures from the technical definition of a directed graph. One option is to view such figures as embellished representations of the relevant abstract objects (in that some information they provide us with does not correspond to anything in the abstract object that they represent); another is to take them to represent some related but more complex abstract object.

⁴⁷ In this regard our use of graphs differs from that of Berker (2015: 325), who uses them to represent facts about overall rather than pro tanto justification.

transmission: an attitude *A* channels support to an attitude *B* just in case support for *A* transmits to *B*.

To represent the fact that there's support for the initial attitude(s) – in this case, the belief that p – we can introduce a new line, embellishing Figure 1 as follows:⁴⁸



In Figure 1.1, the line ending with a circle represents the fact that there is support for the subject's belief that p (though it doesn't tell us what provides that support) and the arrow between the beliefs represents the fact that the belief that p channels support to the belief that q (in other words, that support for the former belief transmits to the latter).⁴⁹ We'll say more about what graphs like this one don't represent later, but for now, just note that it doesn't specify *how much* support there is for either belief.

While there is a fairly extensive literature in epistemology making use of directed graphs in modeling relations of epistemic support (or justification), as we have done above – see, e.g., Pollock (1995), Berker (2015) – this is just a special case of the more general project of modeling support relations of various kinds between mental states of various kinds. Consider, for instance, the following case, involving a belief and an intention ('I' for intention):



Figure 2

⁴⁸ For reasons familiar from the literature on "transmission failure" (e.g., Pryor 2012), it's important that the support for the initial attitude(s) be "independent" of the support for the "concluding" attitude. We take this to be part of what is represented by the new line.

⁴⁹ Given the technical definition of a directed graph (indicated in footnote 46) as a certain kind of abstract object, this embellished figure does not correspond to any such graph, since it represents both a relation (that of channeling support) and a property (that of being supported) while a directed graph technically only represents a relation.

Figure 2 represents both the fact that there's support for the belief that one ought to φ and the fact that this belief channels support to the intention to φ . This is compatible with there being other (and more direct) sources of support for that intention. If you have promised to φ , for example, that fact might directly support the intention to φ (rather than indirectly by way of supporting the belief that you ought to φ).⁵⁰ Figure 2 merely represents one source or route of support. (Analogous points apply to all our figures.)

The S.O.S. view thus provides an answer to a question raised by Hawthorne and Magidor (2018) concerning whether explanationist views of reasons can be developed in ways that go beyond simple, oft-criticized models involving "weights" and "weighing" reasons. While they don't themselves provide an answer to this question, they mention work that might serve as inspiration, including work in the philosophy of science on modeling causal forces (and their interactions) as well as work in epistemology that uses inference graphs to model "directed patterns and strength of support between 'prima facie' reasons and various propositions as well as phenomena of defeat" (2018: 134). The support-centric approach naturally lends itself to this kind of project given the force-like nature of support. As support is always support *for* some attitude, it is naturally represented by arrows. By contrast, the common-sense notion of weight lacks direction, which makes directed graphs less natural in representing facts about normative weight.

Figure 1.1 and Figure 2 both illustrate what we call substantive *I-support*: exactly one attitude channels support to exactly one other attitude. Following Berker (2015), I-support should be distinguished from both *V-support* and *Y-support*. In cases of V-support, some attitudes each independently channel support to an attitude, while in cases of Y-support, some attitudes *together* channel support to an attitude. Here is an example of each ('P' for preference):⁵¹

⁵⁰ Thanks to an anonymous reviewer for prompting this clarification, and for the example. We lack the space to discuss their worry that if the fact that you have promised to φ supports the intention to φ both directly and indirectly (via the belief that you ought to φ), then that would be a counter-intuitive form of overdetermination.

⁵¹ Figure 4, as well as some figures that will appear later, are strictly speaking (embellished representations of) so-called directed hypergraphs, which are generalizations of directed graphs.



Figure 4: Substantive Y-support

Figure 3 illustrates that the beliefs that p and that q are substantially supported (again, the graph does not tell us why), and that each of these beliefs individually channel support to the belief that p or q. Figure 4 illustrates that the preference for X over Y and the preference for Y over Z are substantially supported and that they together channel support for the preference for X over Z. Thus, in the first situation, it's overdetermined that there is support for the belief that p or q, whereas in the second situation no overdetermination occurs. Parallel distinctions are common in discussions of metaphysical grounding; it is standardly held, for example, that each of the facts [p] and [q] fully grounds the fact [p or q] on its own while they only together fully ground the fact [p and q], so that the disjunctive fact is overdetermined but the conjunctive fact is not.

While we have so far focused on reasons *for*, the S.O.S. view can be naturally extended to reasons *against*. The basic idea is that while reasons for are sources of support, reasons against are sources of opposition. A natural way of understanding

opposition is in terms of support for *not* having a certain attitude. An alternative approach is to posit a *sui generis* notion of opposition that is the negatively-valenced polar counterpart of support (rather than just a species of support distinguished by what is supported). We'll remain neutral on which of these approaches is preferable.⁵² What matters is that support and opposition are both force-like, and hence kinds of pressure, differing mainly in direction or valence: support is pressure "for" while opposition is pressure "against".

We represent opposition using solid red lines terminating in boxes. Here's an example of Y-opposition (V-opposition and I-opposition are also possible):



Figure 5: Substantive Y-opposition

At a later stage, it will be helpful to represent *merely possible attitudes*; i.e., attitudes the relevant subject does not in fact have, but which may be supported or opposed. We do so using dashed circles. For example, the following diagram represents a case in which the subject does not believe that p or q, even though her belief that p channels support for believing that p or q:

⁵² For discussion of reasons against, see, e.g., Snedegar (2018) and Tucker (2022). Note that if we want to represent situations in which a particular attitude is both pro tanto supported and pro tanto opposed, we need to assume that a particular figure can represent both support-relations and opposition-relations. While this is straightforward in practice (one just draws two kinds of arrows), it again requires a departure from the technical definition of a directed graph (cf. footnotes 46 and 49).



Figure 6: Merely possible attitude(s)

To summarize, here are the types of pressure-related facts that the graphs introduced so far represent:

- The *transmission* (and hence channeling) of substantive pressure.
- The *objects* of substantive pressure i.e., the attitudes being supported or opposed.
- The *valence* of the relevant substantive pressures—i.e., whether support or opposition is being provided—and hence the normative effects on the relevant attitudes—i.e., whether they are being supported or opposed.
- The *modal status* of the relevant attitudes—i.e., whether they are actually had or merely possible.

Moreover, the graphs can be used in more or less general ways. Although the graphs we've used so far are naturally understood as representing kinds of cases (e.g., a subject believing that *p* at some time), they can also be used to represent particular cases (e.g., Barack Obama's believing that grass is green on January 23, 2011 at 4:56 p.m. ET).

There are also several types of support-related facts that our graphs do not represent, including:

- The *magnitude* or *strength* of pressure; e.g., whether an attitude channels a lot or merely a little support.
- Whether the attitudes involved are *on balance* supported or opposed.
- Whether the support or opposition provided is *modified* in various ways. Possible kinds of modification include undermining, intensifying, attenuating, and that provided by higher-order evidence and higher-order reasons more generally.⁵³

⁵³ Of particular relevance here is the familiar distinction between component forces and resultant force. This distinction allows us to make sense not just of how normative support works but also how to distinguish opposition and undermining, which are the two main ways in which support is defeasible. In cases involving opposition, for instance, the initial component (pro tanto) support for having an attitude *A* can remain constant even though the overall resultant support is altered by competing component (pro tanto) opposition to *A*. In cases involving undermining, by contrast, the initial component support is itself weakened (and at

• The role, if any, that general normative *principles* play in the explanation of facts about support (for discussion, see, e.g., Fogal and Risberg 2020).

These details are important to account for in a complete theory of reasons and rationality, and many could in principle be included in more complex graph-theoretic representations.⁵⁴ We omit them both for simplicity and because they are not needed for the purposes of this paper. Thus, as with all models, our graphs should be seen as idealized representations of certain features of the relevant phenomenon, rather than as comprehensively describing all its aspects.⁵⁵

4.3 Structural Pressure and Structural Rationality

On our view, structural rationality and substantive rationality are distinct but importantly similar. In particular, we take each kind of rational evaluation to be associated with a distinct kind of pro tanto rational pressure or force: substantive rationality with substantive pressure and structural rationality with (what we'll call) structural pressure.⁵⁶ The main innovation is thus the introduction of the notion of structural pressure, which is importantly similar to – despite being distinct from – the more familiar notion of substantive pressure.⁵⁷ To illustrate the basic idea: whereas having *strong substantive reasons* to believe that *p* and *strong substantive reasons* to believe that *q* generates significant (pro tanto) *substantive* pressure to believe that *q*. Similarly, just as having strong substantive reasons to *E* and strong substantive reasons to believe that *E*-ing requires *M*-ing generates significant

the limit eliminated), and the resultant support changes as a result. So even if the resultant support – i.e., the all-things-considered balance of pro tanto support – ends up the same in both cases, the explanation of why it ends up that way differs. Similar things can be said about other ways in which the strength or significance of a reason may be altered, e.g. being "intensified."

⁵⁴ For example, so-called weighted directed graphs (or directed networks), which are graphs with weights assigned to their arrows, may be used to model (particular) amounts of support; arrows of different colors might be used to model kinds of support from different normative domains (moral, prudential, etc.); and different kinds of nodes may be used to model non-attitudinal mental states such as sensory experiences and feelings. See also Berker (2015) for discussion of how graph-theoretical tools can be used to inform traditional debates in epistemology over the structure of epistemic justification, such as that between coherentists, foundationalists, and infinitists.

⁵⁵ See further, e.g., Weisberg (2007) on idealization in science, Greco (2023) and Roussos (forthcoming) on modeling in epistemology, Roussos (2022) on modeling in normative ethics, and Williamson (2017) on modeling in all areas of philosophy.

⁵⁶ Fogal (2020) uses the terms 'justificatory' and 'attitudinal' pressure instead of 'substantive' and 'structural'. We opt for the latter in the interest of clarity.

⁵⁷ For an important precursor to our view within epistemology, see Pryor (2004; 2018).

substantive pressure to intend to *M*, so merely intending to *E* and believing that *E*-ing requires *M*-ing generates significant structural pressure to intend to *M*.

As this illustrates, the main difference between substantive and structural pressure lies in what generates them – i.e., their *grounds* or *sources*. Whereas substantive pressure is generated by one's reasons, structural pressure is generated by one's mere attitudes, regardless of their normative status. They can thus come apart in both directions: just as there can be substantive pressure without structural pressure, so there can be structural pressure without substantive pressure. The two kinds of rational pressure are otherwise very similar in nature. In particular, both kinds of rational pressure:

- (a) are *force-like*, having both magnitude and direction as well as being defeasible, transmissible, and separable from their sources (cf. §4.1);
- (b) include a variety of complex relations, and in particular relations of support, opposition, and undermining, all of which come in degrees; and
- (c) determine whether one's attitudes are (substantively/structurally) rational, and hence are more fundamental than non-graded facts about (ir)rationality.

Thus, structural pressure is so-called primarily because of its association with structural rationality, not because it behaves in a particularly structure-like fashion. It is force-like, just as substantive pressure is, rather than structure-like.⁵⁸

To elaborate on (c): the relationship between structural pressure and structural (ir)rationality is essentially the same as the relationship between substantive pressure and substantive (ir)rationality. Just as substantive rationality is (roughly) a matter of having attitudes that are on balance substantively supported, so structural rationality is (roughly) a matter of having attitudes that are on balance structurally supported. In neither case is an attitude's being (ir)rational a matter of it being supported or opposed by some particular amount or source of pressure considered on its own. Instead, in both cases, what determines the rational status of an attitude are global on-balance facts about rational pressure – facts which are in turn determined by various (possibly quite complex) local, graded facts about pro tanto rational pressure. Thus, if the structural support for attitude A generated by the subject's other attitudes is on balance sufficiently strong, then A is structurally rational (and possibly rationally required). Likewise, if the structural opposition to A that the subject's other attitudes generate is on balance sufficiently strong, then A is structurally irrational. The same goes for substantive pressure and substantive (ir)rationality. Thus, in both cases, facts about rational pressure – whether substantive or structural – determine the rational status of one's attitudes, and hence are more fundamental than non-graded facts about (ir)rationality.

⁵⁸ Thanks here to an anonymous reviewer.

Although philosophers have mostly focused on non-graded facts about (ir)rationality, our view also naturally allows for various "graded" notions of (ir)rationality, such that some (ir)rational attitudes count as more (ir)rational than others.⁵⁹ For example, in both the substantive and the structural domain, we can introduce graded notions of rationality and irrationality that are sensitive to degrees of support and opposition, in order to recognize differences between (in)adequately yet unequally supported attitudes. A mathematician's belief that 7+5=12, for example, might on balance be more substantively supported and hence more substantively rational than a kindergartener's belief that 7+5=12, even though both beliefs count as rational (as opposed to not rational). Similarly, a diehard conspiracy theorist's vaccine skepticism might be better structurally supported and hence more substantively rational (while also being more substantively opposed and hence more substantively *ir*rational) than that of a teenager whose main source of misinformation is TikTok.

What should we say about attitudes that are neither structurally supported nor structurally opposed by one's other attitudes, as in Combos 7-9? As before (cf. §4.1), one option is to say that such attitudes are structurally rational in virtue of not being opposed; another is to say that they are structurally irrational in virtue of not being supported; and a third is to say they are neither structurally rational nor irrational. Our view here mirrors our view of attitudes which are neither substantively supported nor substantively opposed: we think the appeal of all three options is best captured by saying that attitudes which are neither structurally supported nor opposed are structurally rational in a "weak" sense but not in a "strong" sense. In other words, such attitudes are structurally rational in the sense that they are weakly (but not strongly) rationally permissible.⁶⁰

Like substantive pressure, structural pressure can be fruitfully modeled using directed graphs (and embellishments thereof). As before, the nodes represent the subject's attitudes. However, to distinguish relations of pro tanto structural support from those related to substantive support, we introduce a new dashed arrow. A dashed arrow from (a node representing) attitude *A* to (a node representing) attitude *B* represents that *A* structurally supports *B*. Simple cases of structural I- and Y-support can thus be represented as follows (Figures 7 and 8), paralleling those of substantive I- and Y-support presented earlier (Figures 1.1 and 4):

⁵⁹ Cf. Siscoe (2021), Staffel (2019), and others.

⁶⁰ Thus, like our use of '(in)coherent', our use of 'structurally (ir)rational' to pick out polar opposites is (semi-)stipulative, and hence compatible with other uses, including ones that pick out privative opposites. For example, one might use 'structurally rational' in a broader, weaker way to mean that a (set of) attitude(s) is not structurally irrational, and so include both coherent (sets of) attitudes and neutral ones.



Figure 7: Structural I-support



Figure 8: Structural Y-support

We can introduce a similar convention for structural opposition, using a red dashed line terminating in a box. Such a line from (a node representing) attitude *A* to (a node representing) attitude *B* represents that *A* structurally opposes *B*. For example:



Figure 9: Structural Y-opposition

Structural V-support and V-opposition are possible as well.

In view of the similarities between structural pressure and substantive pressure, one may wonder whether there is anything more to say about how they are related. We think there is. For as the similarities (and differences) between the directed graphs representing substantive pressure and those representing structural pressure suggest, structural pressure strongly resembles "hypothetical" substantive pressure, and so hypothetical facts about substantive pressure can serve as a guide to (non-hypothetical) facts about structural pressure.⁶¹ In particular, though the following thesis may admit of exceptions, it is a good heuristic:

Proportionality Thesis: The strength and direction of structural pressure generated by attitudes $A_1...A_n$ is proportional to the strength and direction of substantive pressure that $A_1...A_n$ would channel if they were decisively substantively supported (holding everything else fixed).⁶²

To illustrate, return to Combos 1-3 and focus solely on the first two attitudes in each:

Combo 1.5: {belief that it is raining; belief that if it is raining, then the streets are wet}

Combo 2.5: {preference for strawberries over apples; preference for apples over cherries}

Combo 3.5: {intention to go for a run; belief that in order to go for a run, you have to get out of bed}

Suppose we want to know: what attitudes are structurally supported by the attitudes in Combos 1.5-3.5, and to what extent? To answer this question we can turn to the Proportionality Thesis. The answer it provides involves making a supposition and then

⁶¹ Cf. Pryor (2004, 2018) and Fogal (2020).

⁶² One potential problem for the Proportionality Thesis arises if there are attitudes which necessarily do not enjoy substantive support, such as (say) the preference for suffering over happiness. Whether this potential problem is a real problem depends not only on whether there are such attitudes but also on the best way to handle counterfactuals with impossible antecedents; we therefore set it aside for present purposes. Note that the Proportionality Thesis focuses on the substantive pressure that $A_1...A_n$ would *channel* if they were substantively supported, not just on the substantive pressure that there *would be* if $A_1...A_n$ were substantively supported. These may come apart, for example because some attitudes may necessarily enjoy substantive support: if a preference for happiness over suffering is necessarily supported, then it always holds that if (say) *S*'s belief that grass is hot pink were decisively substantively supported, then *S*'s preference for happiness over suffering would be substantively supported. We shouldn't therefore conclude that *S*'s belief that grass is hot pink structurally supports their preference for happiness over suffering would be substantively supports their we of coherence and incoherence that appeals to substantive support (cf. footnote 39).

asking a series of questions. To start: suppose (perhaps contrary to fact) that the attitudes in Combos 1.5-3.5 were to be decisively substantively supported, and hence substantively rational. Now ask: which other attitudes would these attitudes channel support to? The obvious answer includes at least the following: the belief that the streets are wet, the preference for strawberries over cherries, and the intention to get out of bed. We then ask: how much substantive support would they enjoy? Though perfect precision shouldn't be expected given the coarse-grainedness of the supposition, the obvious answer is: a lot! We thus arrive at an answer to our original question: the attitudes in Combo 1.5 provide a lot of – indeed, normally decisive – structural support for the belief that the streets are wet, the attitudes in Combo 3.5 do the same for the intention to get out of bed.⁶³

We can use directed graphs to illustrate the Proportionality Thesis as well. The figures below depict the four main stages. Figure 10.1 represents the *initial question* stage: which attitudes, if any, do the intention to go for a run and the belief that doing so requires getting out of bed together structurally support? Figure 10.2 represents the *supposition* stage: we suppose that there is decisive substantive support for those attitudes, in order to see what happens. Figure 10.3 then represents the *channeling* stage: those attitudes channel the (supposed) support to the intention to get out of bed. Finally, Figure 10.4 represents the *answer* stage: given the foregoing, the intention to go for a run and the belief that doing so requires getting out of bed structurally support (normally decisively) the intention to get out of bed.



Figure 10.1: Initial Question Stage

⁶³ What should we say about someone who has the attitudes in Combos 1.5, 2.5, or 3.5 but lacks the attitudes they support? Are they automatically thereby structurally irrational? Not obviously. For discussion, see Fogal (forthcoming).



Figure 10.4: Answer Stage

We can now see why the pressure-based view of structural and substantive rationality is not a form of Cartesian dualism. In contrast to the striking differences between the two dimensions of rational evaluation that arise on Worsnip's view (cf. §3.1), our view vindicates a corresponding set of similarities between them:

- (1*) Structural and substantive rationality are formal to a similar extent.
- (2*) Neither structural nor substantive rationality are fundamentally negative.
- (3*) Neither structural nor substantive rationality primarily concern combinations of attitudes.
- (4*) Structural and substantive rationality are both graded in the same way.

Concerning (1*), with both structural and substantive rationality we can (as we've seen) point to patterns "using variables and schematic symbols, and without even fully specifying the content of the attitudes in question" (Worsnip 2021: 8). Concerning (2*), facts about both structural and substantive rationality obtain in virtue of facts about support and opposition, so neither is fundamentally negative. Concerning (3*), individual attitudes can be assessed as both structurally and substantively rational. (It's true that structural rationality is concerned with combinations of attitudes in the sense that structural pressure is generated by one or more attitudes and directed at another attitude — but this does nothing to support Cartesianism.) Finally, concerning (4*), in both the substantive and the structural domain the degree to which an attitude is rational is a matter of how much support and opposition there is for and against it.

4.4 Structural Pressure and Coherence

We'll now use facts about structural pressure to provide an account of coherence and incoherence that, unlike its competitors, straightforwardly accounts for all nine cases presented in the introduction. (In what follows we use 'support' and 'opposition' to mean 'structural support' and 'structural opposition', unless otherwise indicated.)

To begin with, note that both requirement- and guarantee-based accounts treat coherence and incoherence as non-gradable, monadic properties of sets of attitudes. On our view, by contrast, the explanatory work is done by gradable relations among attitudes (rather than sets thereof). In particular, facts about structural (ir)rationality and facts about (in)coherence share a "common ground," as both obtain in virtue of facts about support and opposition. Thus, neither type of fact obtains in virtue of the other, though they reliably correlate.⁶⁴

⁶⁴ As noted earlier, this differs from Worsnip's view, on which "states are structurally (ir)rational *in virtue of* their incoherence; their incoherence grounds their structural irrationality" (2021: 4). It's not clear, however, how (if at all) such a picture can be reconciled with the standard Broomean idea that a set of attitudes is (structurally) irrational in virtue of violating a rational requirement (cf. Worsnip 2021: 169).

The central idea is that whether and to what extent a set of attitudes $\{A_1...A_n\}$ is coherent, incoherent, or neutral depends on whether and to what extent support and opposition relations obtain among $A_1...A_n$. When some members of $\{A_1...A_n\}$ support others, that contributes to making $\{A_1...A_n\}$ coherent; when some members of $\{A_1...A_n\}$ oppose others, that contributes to making $\{A_1...A_n\}$ incoherent; and when no member of $\{A_1...A_n\}$ either supports or opposes any other member, then $\{A_1...A_n\}$ is neutral, in that it is neither (strongly) coherent nor incoherent.⁶⁵ Thus:

The Support-Based View of Coherence (SBC): A set of attitudes $\{A_1...A_n\}$ is coherent to the extent that (and because) relations of structural support obtain among $A_1...A_n$.

The Opposition-Based View of Incoherence (OBI): A set of attitudes $\{A_1...A_n\}$ is incoherent to the extent that (and because) relations of structural opposition obtain among $A_1...A_n$.

Neutrality as the Absence of Pressure (NAP): A set of attitudes $\{A_1...A_n\}$ is neutral if and only if (and because) no relations of structural support or opposition obtain among $A_1...A_n$.⁶⁶

SBC and OBI are accounts of degrees of coherence and incoherence. If some attitudes in $\{A_1...A_n\}$ support others while some attitudes in $\{A_1...A_n\}$ oppose others, then $\{A_1...A_n\}$ counts as both coherent to some degree and incoherent to some degree. Facts about degrees of coherence and incoherence then determine "on balance" facts about coherence and incoherence: $\{A_1...A_n\}$ is coherent on balance iff (and because) $\{A_1...A_n\}$ is coherent to a greater degree than $\{A_1...A_n\}$ is incoherent; vice versa for on balance facts about incoherence.⁶⁷ We can further say that $\{A_1...A_n\}$ is coherent *simpliciter* iff and

⁶⁵ Can individual attitudes (or their singleton sets) be (in)coherent? On our view, this is possible just in case the relations of structural support and opposition are not irreflexive, which is an issue we leave open. For relevant discussion see Worsnip (2021: 17).

⁶⁶ While we take the intuitive notion of neutrality to be non-graded, a graded notion of neutrality can also be defined: one set might count as "more neutral" than another iff the former is closer to being (perfectly) neutral than the latter. A notion of neutrality on balance (that differs from the one captured by NAP) can be defined as follows: a set of attitudes is neutral on balance iff it is coherent and incoherent to equal (perhaps non-zero) degrees.

⁶⁷ This relation between pro tanto and overall coherence and incoherence mirrors the standard view of the relation between, for example, pro tanto and overall harm and benefit, on which an event is harmful overall iff it is pro tanto harmful to a greater degree than it is pro tanto beneficial.

because $\{A_1...A_n\}$ is coherent to some degree and not incoherent to any degree (and vice versa for incoherence simpliciter).

Suppose we take a coherent set of attitudes – say, the set {belief that one ought to φ ; intention to φ } – and add a random, unrelated attitude to it – say, the belief that p. What should one say about the resulting set of attitudes – i.e., {belief that one ought to φ ; intention to φ ; belief that p}? Intuitively, the belief that p looks like "junk" in this set, as it neither supports nor opposes any attitude in it. We can accordingly introduce the following notion of junkiness:

The Pressure-Based View of Junkiness (PBJ): A set of attitudes $\{A_1...A_n\}$ is junky to the extent that (and because) it contains attitudes that are not involved (in any way) in the provision of either support or opposition to any attitudes among $A_1...A_n$.

We can then say that $\{A_1...A_n\}$ is junky *simpliciter* just in case it is neutral (in the sense defined by NAP). The set {belief that one ought to φ ; intention to φ ; belief that p}, moreover, can be said to be *junkily coherent* (as by PBJ, it is junky to some extent, and by SBC, coherent to some extent). Consider an analogy: the state of affairs *Pekka is happy and grass is green* is standardly taken to be intrinsically good (though not *basically* so), despite the fact that it contains junk (i.e., *grass is green*). This is similar to the view that junky sets of attitudes can be coherent.⁶⁸

More generally, the degree to which a set of attitudes is (in)coherent is a complex function of the support relations, opposition relations, and perhaps others (such as undermining relations and other modifying relations; cf. §4.2) obtaining among the attitudes. Both the number of instances of these relations and their strength—i.e., how much support or opposition that is provided in a given instance—are relevant. If A_1 provides lots of support for A_2 while A_3 provides only a little support for A_4 , for instance, then { A_1, A_2 } is coherent to a greater degree than { A_3, A_4 }; and if each of A_5 and A_6 provide lots of support for A_7 (i.e., what we earlier called 'V-support') then { A_5, A_6, A_7 } is even more coherent than both { A_1, A_2 } and { A_3, A_4 }. Parallel remarks hold for incoherence: if a subject believes that p and that if p then q, for example, they make a more serious rational mistake if they are fully confident that q is false than if they are merely quite confident that q is false—and that's because their attitude toward q is opposed to a greater degree in the former case than in the latter. How to determine the degree to which a set of attitudes is (in)coherent in all cases depends on a host of issues concerning how amounts of pressure might be formally measured or represented

⁶⁸ If one disagrees with this view, other options for handling junk are available. For example, one might adjust SBC so that it requires that the set in question be non-junky, or adjust one's "measure" of coherence so that it entails that a set's degree of junkiness detracts from its overall degree of coherence. Thanks to an anonymous reviewer for prompting this discussion.

which are beyond the scope of this paper. (For example, can such amounts be perspicuously represented by real numbers? Can they be incomparable? How, if at all, can they be aggregated? And how should borderline cases be handled?) The more complex the truth about such issues turns out to be, the more complex the function from pressure to (in)coherence will likely be as well.

Coherentists about epistemic justification have proposed a variety of measures of (what they call) coherence, often by relying on tools from probability theory.⁶⁹ In such discussions, coherence is typically treated primarily as a property of sets of propositions and only derivatively as a property of sets of doxastic attitudes (in particular, full or partial beliefs) toward those propositions. One popular idea, roughly put, is that the coherence of a set of propositions depends on whether and to what extent its members increase the probability of each other.⁷⁰ While the support-centric approach is compatible with thinking that a probabilistic measure of this type captures part of theoretical structural rationality, it has the virtue of being more general. For as emphasized earlier (§2), there are a variety of non-doxastic attitudes—including preferences, intentions, fears, hopes, desires, "admirations," and the like—that also enter into relation(s) of (in)coherence of the kind we're concerned with. And it is difficult, to put it mildly, to see how probabilistic proposals could be generalized so as to cover such cases.

Moreover, while we lack the space to properly adjudicate the issue, it is doubtful that even theoretical structural rationality can be fully understood in probabilistic terms. It has been argued, for example, that phenomena like undermining and higher-order defeat are difficult to capture in a purely probabilistic framework (e.g., Christensen 1992, Pryor 2013, Weisberg 2015; see also Christensen 2010). For related reasons it is difficult for probabilistic approaches to explain why it is incoherent to both believe that *p* and believe that one's evidence does not support believing that *p* – after all, neither belief content, it seems, is guaranteed to significantly reduce the likelihood of the other.⁷¹ The support-centric approach, by contrast, fares better: if one assumes (as is plausible) that if *S* has substantive support for believing that their evidence does not support believing that *p*, then (at least normally) *S* thereby has substantive support for not believing that *p* and that the combination of those two beliefs is therefore incoherent (at least normally and at least to some degree)

⁶⁹ See, e.g., Fitelson (2003) and Hartmann and Trpin (forthcoming). We thank an anonymous reviewer for encouraging us to discuss these proposals in more detail.

⁷⁰ Other probabilistic notions of coherence are also available (cf. Schubert and Olsson 2012: 114).

⁷¹ We ignore numerous complications. For discussion of how probabilistic approaches might accommodate higher-order beliefs, see Schubert and Olsson (2012).

4.5 Revisiting the Cases

One of the most important features of the support-centric view of (in)coherence is that it's emphatically non-Schopenhauerian: it does not treat coherence as merely the absence of incoherence. This is part of what enables it to straightforwardly explain the salient differences between the combinations of attitudes presented in the introduction. In this subsection we'll show how.

First consider the (strongly) coherent combinations:

Combo 1: {belief that it is raining; belief that if it is raining, then the streets are wet; belief that the streets are wet}

Combo 2: {preference for strawberries over apples; preference for apples over cherries; preference for strawberries over cherries}

Combo 3: {intention to go for a run; belief that in order to go for a run, you have to get out of bed; intention to get out of bed}

These combinations can be represented as follows:



Figure 11: Combo 1



Combos 1-3 are all instances of structural Y-support. By SBC, they all count as coherent – and indeed as both on balance and non-junkily coherent. For in Combo 1, the first two beliefs support the third; in Combo 2, the first two preferences support the third; and in Combo 3, the first intention and the belief together support the second intention. (Applying the Proportionality Thesis: in each case, under the supposition of decisive substantive support for the first two attitudes in the set, a similar amount of substantive support for the third attitude is channeled as well.) Moreover, since no attitude in these sets opposes any other, they also count as coherent simpliciter.

Now consider Combos 4-6:

Combo 4: {belief that it is raining; belief that if it is raining, then the streets are wet; belief that the streets are not wet}

Combo 5: {preference for strawberries over apples; preference for apples over cherries; preference for cherries over strawberries}

Combo 6: {intention to go for a run; belief that in order to go for a run you have to get out of bed; intention to not get out of bed]

These combinations can be represented as follows:



Figure 14: Combo 4



Figure 16: Combo 6

In Combos 4 and 5, there are three instances of Y-opposition, as each pair of attitudes opposes the third. Applying the Proportionality Thesis: if there is decisive substantive support for believing that it is raining and that if it is raining then the streets are wet, there will be a similar amount of *opposition* channeled to the belief that the streets are not wet; if there is decisive substantive support for believing that it is raining and that the streets are not wet, there will be a similar amount of opposition channeled to the belief that if it is raining and that the streets are not wet, there will be a similar amount of opposition channeled to the belief that if it is raining then the streets are wet; and if there is decisive substantive support for believing that if it is raining then the streets are wet and that the streets are not wet, there will be a similar amount of opposition channeled to the belief that if it is raining then the streets are wet and that the streets are not wet, there will be a similar amount of opposition channeled to the belief that if it is raining then the streets are wet and that the streets are not wet, there will be a similar amount of opposition channeled to the belief that it is raining. Corresponding things can be said about Combo 5. Thus, by OBI, both combinations are incoherent (on balance, non-junkily, and simpliciter).

In Combo 6, something interestingly different happens: there are only two instances of Y-opposition. That is because, while each belief-intention pair opposes the remaining intention, the intention-intention-pair does not by itself oppose the belief. Applying the Proportionality Thesis: if there is decisive substantive support for intending to go for a run and intending not to get out of bed, there is not thereby substantive opposition to the belief that in order to go for a run you have to get out of bed. This is why there is no relation of structural opposition terminating in the belief in Figure 16. Although (by OBI) Combo 6 still counts as incoherent (on balance, non-junkily, and simpliciter), its "anatomy" differs from that of Combos 4 and 5.⁷²

Finally, let's consider the neutral combinations:

⁷² Differences in anatomy (i.e., in support and opposition relations) help explain intuitive differences in the way attitudes can be rationally revised on the basis of others, as well as differences between (structurally) good and bad patterns of reasoning more generally (see Fogal and Risberg ms for extended discussion). To take one example: there is an oft-observed "asymmetry" concerning different ways of complying with wide scope requirements of instrumental rationality. Suppose, for example, that you violate means-end consistency (§3.1): you intend to E, believe that E-ing requires M-ing, and intend to not-M. While it seems rationally permissible (at least insofar as those attitudes go) to revise either intention on the basis of the other intention together with the means-end belief, it seems rationally *impermissible* to revise the means-end belief on the basis of the two intentions. This is one of the most prominent objections to the wide scope requirement-centric account of structural rationality (for discussion and citations, see Worsnip 2021: ch. 6.4.3). But the asymmetry is readily explained on the support-centric view. In cases of means-end inconsistency, each intention together with the means-end belief opposes the other intention, and hence each such pair provides a suitable basis for revision. By contrast, the two intentions themselves (whether taken individually or together) do not oppose the means-end belief, and hence fail to provide a suitable basis for revision. Thanks to an anonymous reviewer for prompting us to make this connection.

Combo 7: {belief that it is raining; belief that if you drop an egg, it will break; belief that Washington D.C. is the U.S. capital}

Combo 8: {preference for strawberries over apples; preference for Star Wars over Star Trek; preference for talking with people over talking with parrots}

Combo 9: {intention to go for a run; belief that in order to make dinner you need to buy pasta; intention to call your mom}

As emphasized above, the attitudes in these combinations seem random and unrelated, neither clashing nor cohering. The support-centric view directly yields this result: none of the attitudes involved either supports or opposes any of the others, so by NAP these combinations all count as neutral. We only provide a graph for Combo 7, as the graphs for Combos 8 and 9 have the same structure (or, rather, lack thereof):



Applying the Proportionality Thesis: if there is decisive substantive support for any one or two of these beliefs, there is not thereby any substantive support for or opposition to the remaining one(s).

5. Concluding remarks

We have argued that the two leading views of coherence and incoherence are Schopenhauerian—they treat coherence simply as the absence of incoherence and so should be rejected. We have also defended an alternative, support-centric view of coherence and incoherence that is decidedly non-Schopenhauerian. The resulting picture is a form of non-Cartesian dualism: structural and substantive rationality are both a matter of having attitudes that there is on-balance support for having.

Even for those who are unconvinced by the support-centric approach, however, we hope to have illustrated the importance of the tripartite distinction between coherent, incoherent, and neutral sets of attitudes. We take it to be a datum that there is a structurally significant difference between strongly coherent attitudes and neutral ones, and the leading views in the debate overlook this fact. At a bare minimum, doing justice to this datum should be seen as an adequacy constraint on theories of coherence and structural rationality.⁷³

⁷³ Earlier version of this essay were presented at the ANU/Humbolt/Princeton Summer Institute on Practical Normativity, Bielefeld University, University of Cologne, University of Copenhagen, Heidelberg University, Inland Norway University, Linköping University, University of Missouri-St. Louis, New York University, NYU Abu Dhabi, Stockholm University, University of Texas-Austin, Union College, and Uppsala University. We are grateful to the participants in these events as well as to Hossein Ayati, Karl Bergman, Samuel Fullhart, Camilo Martinez, Victor Moberger, Declan Samithies, and Alex Worsnip for their valuable feedback. We owe a special debt of gratitude to two anonymous reviewers of this journal for comments on earlier drafts. More general debts of gratitude are owed to Selim Berker, John Broome, Matti Eklund, and Jim Pryor. Fogal's and Risberg's work on this paper was supported by Grant 2020-01955 from the Swedish Research Council. Risberg's work was also supported by Grant P23-0506 and Grant PF23-002 from the Bank of Sweden Tercentenary Foundation.

References

Berker, Selim. 2015. "Coherentism via Graphs." Philosophical Issues 25, no. 1: 322-352.

- Berker, Selim. 2022. "The Deontic, the Evaluative, and the Fitting." In *Essays in the Philosophy of Normativity*, edited by Christopher Howard and Rach Cosker-Rowland. Oxford University Press.
- Berker, Selim. 2024. "Is There Anti-Fittingness?" Ergo 11, no. 39: 1051-1082.
- Broome, John. 2013. Rationality Through Reasoning. Wiley-Blackwell.
- Broome, John. 2025. "Is There Reason? Are There Reason-Forces?" In *New Essays on Normative Realism*, edited by Paul Boghossian and Christopher Peacocke. Oxford University Press. (page references to penultimate draft online)
- Brunero, John. 2020. Instrumental Rationality: The Normativity of Means-Ends Coherence. Oxford University Press.
- Christensen, David. 1992. Confirmational Holism and Bayesian Epistemology. *Philosophy of Science* 59, no. 4: 540-557.
- Christensen, David. 2010. Higher Order Evidence. *Philosophy and Phenomenological Research* 81, no. 1: 185–215.
- Daoust, Marc-Kevin. 2023. "Structural Rationality and the Property of Coherence." *Pacific Philosophical Quarterly* 104, no. 1: 170-194.
- Greco, Daniel. 2023. Idealization in Epistemology: A Modest Modeling Approach. Oxford University Press.
- Fink, Julian. 2023. "The Essence of Structural Irrationality." *Journal of Ethics and Social Philosophy* 26, no. 2. https://jesp.org/index.php/jesp/article/view/1886
- Fitelson, Branden. 2003. "A Probabilistic Theory of Coherence." *Analysis* 63, no. 3: 194–199.
- Fogal, Daniel. 2020. "Rational Requirements and the Primacy of Pressure." *Mind* 129, no. 516: 1033-1070.
- Fogal, Daniel. Forthcoming. "On (Not) Fitting Things Together." Analysis.
- Fogal, Daniel and Olle Risberg. 2020. "The Metaphysics of Moral Explanations." In *Oxford Studies in Metaethics* 15, edited by Russ Shafer-Landau. Oxford University Press.
- Fogal, Daniel and Olle Risberg. 2023a. "Explaining Normative Reasons." *Noûs* 57, no. 1: 51-80.
- Fogal, Daniel and Olle Risberg. 2023b. "The Weight of Reasons." *Philosophical Studies* 180: 2573-2596.
- Fogal, Daniel and Olle Risberg. n.d. "The Nature of Good Reasoning."
- Fogal, Daniel and Alex Worsnip. 2021. "Which Reasons? Which Rationality?" *Ergo* 8, no. 11: 306-343.
- Fox, Joshua Isaac. 2023. "Does Schopenhauer Accept Any Positive Pleasures?" *European Journal of Philosophy* 31, no. 4: 902-913.

- Fullhart, Samuel, and Camilo Martinez. 2024. "Coherence as Joint Satisfiability." *Australasian Journal of Philosophy* 102, no. 2: 312-332.
- Hartmann, Stephan, and Borut Trpin. Forthcoming. "Why Coherence Matters." *Journal* of *Philosophy*.
- Hawthorne, John, and Ofra Magidor. 2018. "Reflections on Reasons." In Daniel Star (ed.), *The Oxford Handbook on Reasons and Normativity*. Oxford University Press.
- Kiesewetter, Benjamin. 2017. The Normativity of Rationality. Oxford University Press.
- Kiesewetter, Benjamin. 2024. "Review of *Fitting Things Together*." *Mind* 133, no. 532: 1229-1238.
- Kiesewetter, Benjamin and Alex Worsnip. 2023. "Structural Rationality." Stanford Encyclopedia of Philosophy (Fall 2023 Edition), edited by Edward Zalta and Uri Nodelman. https://plato.stanford.edu/archives/fall2023/entries/rationalitystructural/
- Kolodny, Niko. 2007. "How Does Coherence Matter?" Proceedings of the Aristotelian Society 107: 229-263.
- Langlois, David Joseph. 2014. *The Normativity of Structural Rationality*. PhD diss., Harvard University. Cambridge, MA.
- Lasonen-Aarnio, Maria. 2021. "Coherence as Competence." Episteme 18, no. 3: 353-376.
- Lee, Wooram. 2024. "What Is Structural Rationality?" *Philosophical Quarterly* 74, no. 2: 614-636.
- Lord, Errol. 2018. The Importance of Being Rational. Oxford University Press.
- Pearl, Judea. 2000. Causality. MIT Press.
- Pollock, John L. 1995. Cognitive Carpentry. MIT Press.
- Pryor, James. 2004. "What's Wrong with Moore's Argument?" *Philosophical Issues* 14, no. 1: 349–378.
- Pryor, James. 2012. "When Warrant Transmits." In *Mind, Meaning, and Knowledge: Themes from the Philosophy of Crispin Wright*, edited by Annalisa Coliva. Oxford University Press.
- Pryor, James. 2013. "Problems for Credulism." In *Seemings and Justification*, edited by Chris Tucker. Oxford University Press.
- Pryor, James. 2018. "The Merits of Incoherence." Analytic Philosophy 59, no. 1: 112-141.
- Przyjemski, Katrina. 2017. "Strong Epistemic Possibility and Evidentiality." *Topoi* 36: 183–195.
- Risberg, Olle. 2019. "Review of Errol Lord's *The Importance of Being Rational*." *European Journal of Philosophy* 27, no. 3: 800-804.
- Rosa, Luis. 2022. "Coherence and Knowability." *Philosophical Quarterly* 72, no. 4: 960-978.
- Roussos, Joe. 2022. "Modelling in Normative Ethics." *Ethical Theory and Moral Practice* 25: 865-889.

- Roussos, Joe. Forthcoming. "Normative Formal Epistemology as Modelling." *The British Journal for the Philosophy of Science*. https://doi.org/10.1086/718493
- Schopenhauer, Arthur. (1839) 2009. "Prize Essay on the Basis of Morals." In *Two Fundamental Problems of Ethics,* translated and edited by Christopher Janaway. Cambridge University Press.
- Schubert, Stefan & Olsson, Erik J. 2012. "On the Coherence of Higher-Order Beliefs." *Southern Journal of Philosophy* 50, no. 1: 112-135.
- Singh, Keshav. 2025. "Rationality Reunified." In *Oxford Studies in Metaethics* 20, edited by Russ Shafer-Landau. Oxford University Press. (page references to penultimate draft online)
- Siscoe, Wes. 2021. "Belief, Rational and Justified." Mind 130, no. 517: 59-83.
- Snedegar, Justin. 2018. "Reasons For and Reasons Against." *Philosophical Studies* 175, no. 3: 725-743.
- Staffel, Julia. 2019. Unsettled Thoughts. Oxford University Press.
- Sylvan, Kurt. 2024. "Review of *Fitting Things Together* by Alex Worsnip." *The Philosophical Review* 133, no. 2: 206-211.
- Tucker, Chris. 2022. "Weighing Reasons Against." In *Oxford Studies in Metaethics* 17, edited by Russ Shafer-Landau. Oxford University Press.
- von Wright, Georg Henrik. 1963. Norm and Action. Routledge and Kegan Paul.
- Way, Jonathan. 2022. "Review of Alex Worsnip, *Fitting Things Together*." Notre Dame Philosophical Reviews. https://ndpr.nd.edu/reviews/fitting-things-togethercoherence-and-the-demands-of-structural-rationality/
- Weisberg, Jonathan. 2015. "Updating, Undermining, and Independence." *British Journal for the Philosophy of Science* 66, no. 1: 121-159.
- Weisberg, Michael. 2007. "Three Kinds of Idealization." *Journal of Philosophy* 104, no. 12: 639-659.
- Williamson, Timothy. 2017. "Model-Building in Philosophy." In *Philosophy's Future*, edited by Russell Blackford and Damien Broderick. Wiley-Blackwell.
- Worsnip, Alex. 2021. *Fitting Things Together: Coherence and the Demands of Structural Rationality*. Oxford University Press.
- Worsnip, Alex. 2025. "Coherence." In *A Companion to Epistemology* (3rd edition), edited by Kurt Sylvan, Jonathan Dancy, Ernest Sosa, and Matthias Steup. Wiley-Blackwell. (page references to penultimate draft online)