Structural Rationality for Self-Governance:

Why We Ought to Have Coherent Sets of Attitudes

Jacob Ashur LiBrizzi

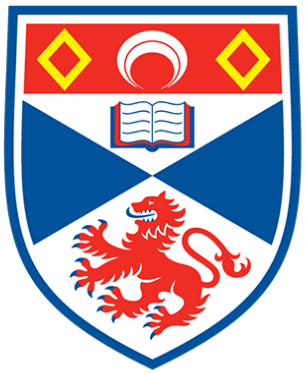

University of
St Andrews

This thesis is submitted in partial fulfilment for the degree of

Master of Philosophy (MPhil)

at the University of St Andrews

December 2020

Candidate's declaration

I, Jacob Ashur Librizzi, do hereby certify that this thesis, submitted for the degree of MPhil, which is approximately 37,000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for any degree. I confirm that any appendices included in my thesis contain only material permitted by the 'Assessment of Postgraduate Research Students' policy.

I was admitted as a research student at the University of St Andrews in September 2018.

I confirm that no funding was received for this work.

Date 26 May 2021 Signature of candidate

Supervisor's declaration

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of MPhil in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree. I confirm that any appendices included in the thesis contain only material permitted by the 'Assessment of Postgraduate Research Students' policy.

Date 26 May 2021 Signature of supervisor

Permission for publication

In submitting this thesis to the University of St Andrews we understand that we are giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. We also understand, unless exempt by an award of an embargo as requested below, that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that this thesis will be electronically accessible for personal or research use and that the library has the right to migrate this thesis into new electronic forms as required to ensure continued access to the thesis.

I, Jacob Ashur Librizzi, confirm that my thesis does not contain any third-party material that requires copyright clearance.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis:

Printed copy

No embargo on print copy.

Electronic copy

No embargo on electronic copy.

Date 26 May 2021 Signature of candidate

Date 26 May 2021 Signature of supervisor

Underpinning Research Data or Digital Outputs

Candidate's declaration

I, Jacob Ashur Librizzi, hereby certify that no requirements to deposit original research data or digital outputs apply to this thesis and that, where appropriate, secondary data used have been referenced in the full text of my thesis.

Date 26 May 2021 Signature of candidate

Table of Contents

General Acknowledgements i

Abstract ii

Dissertation Program iii

0.1: Coherence Requirements iii

0.2: My Framework—Normative Reasons, Oughts, and Rationality. iii

0.2a: Normative Reasons as Explanations of Oughts iv

0.2b: Coherence Requirements as Structural Rationality v

0.2c: Reason-Based Normativity vi

0.3: Our Roadmap. vi

Chapter One: Coherence Requirements 1

1.0: Chapter Introduction 1

1.1: Coherence Requirements as Prohibitions on Incoherent Sets of Attitudes 1

1.2: Six Coherence Requirements, Six Prohibitions 2

1.2a: The Instrumental (Means-Ends) Requirement 3

1.2b: The Intention Consistency Requirement 5

1.2c: The Modus Ponens Requirement 7

1.2d: The Transitive Preference Requirement 9

1.2e: The Doxastic Enkratic Requirement 11

1.2f: The Practical Enkratic Requirement 12

1.3: Chapter Conclusion 20

Chapter Two: Challenges for Normative Coherence Requirements 22

2.0: Chapter Introduction 22

2.1: Why Care about Coherence Requirements? 23

2.1a: Substantively Rational Attitudinal Incoherence 23
2.1b: Structuring Deliberation with Coherence-Based Reasons..... 25
2.1c: Structural Irrationality as Evaluative Criticism..... 26

2.2: Do Coherence Requirements Lead to Bootstrapping and other Controversial Outcomes? 30

2.2a: Narrow-Scope Requirements and Bootstrapping..... 30
2.2b: Wide-Scope Requirements and the Symmetry Problem 34
2.2c: Substantive Reasons Guiding Coherence Requirement Satisfaction..... 37
2.2d: Coherence Requirements and Liberal Transmission 40
2.2e: Coherence Requirements and Unalterable Attitudes..... 43

2.3: Are Coherence-Based Reasons Right-Kind Reasons?..... 48

2.3a: Right-Kind/Wrong-Kind Reasons—A First Pass 48
2.3b: Kolodny's Challenge..... 49
2.3c: Worsnip and Right-Kind Reasons for Structural Rationality..... 50
2.3d: Are Coherence-Based Reasons Action-Oriented Reasons?..... 51
2.3e: Reasons for Action Versus the Right-Kind Framework..... 52
2.3f: Coherence-Based Reasons as Right-kind Reasons for Proper Deliberation 52
2.3g: Right-Kind Reasons in an Agency-Constitutivist Framework 53

2.4: Chapter Conclusion 54

Chapter Three: Structural Rationality for Self-Governance 55

3.0: Chapter Introduction..... 55

3.1: Self-Governance, Proper Descriptions, and Authored Action 56

3.1a: Agents as Potential Self-Governors..... 56
3.1b: Self-Governance and Action 57
3.1c: Proper Descriptions as Constitutive Standards..... 58
3.1d: Action and Proper Descriptions 59

3.2: Coherence Requirements and Normative Reason-Relations 60

3.2a: Reason-Relations—Derivative and Non-Derivative 60
3.2b: Worsnip and Non-Derivative Normativity 63
3.2c: A Problem for Worsnip's Account 63
3.2d: Possible Responses from Worsnip 64
3.2e: A Way Forward 66

3.3: The Necessity of Attitudinal Coherence for Self-Governance 66

3.3a: Attitudinal Coherence as Where We Stand on Issues 67
3.3b: My Argument—Attitude Coherence as the Intelligibility of Actions 70
3.3c: Three Questions for My Position 75
3.3d: A Counterfactual Thesis for Coherence Requirements 81

3.4: The Non-Derivative Normativity of Self-Governance 85

3.4a: Self-Governance under the Non-Derivative Reason Relation 85
3.4b: The Explicability of Non-Derivatively Normative Self-Governance 86

3.5: Chapter Conclusion—Structural Rationality for Self-Governance 91

Dissertation Conclusion: 92

References 93

General Acknowledgments

My thanks go out to everyone who helped me to succeed in this project. First, I wish to thank Adam Etinson, whose annual review feedback helped me narrow my research focus to the normativity of coherence requirements. Next, I wish to thank my secondary supervisor, Justin Snedegar, for navigating me through the complex terrain of this project. Finally, to Jens Timmermann, my primary supervisor, I will always be grateful for the incredible output of time and support. His acumen challenged me to raise the caliber of my work.

Funding

NA

Research Data/Digital Outputs Access Statement

NA

Abstract

In this project I argue that, as agents, we have a normative commitment to avoid incoherent combinations of attitudes. To this end, I defend an amended version of Alex Worsnip's (2020) position on normative requirements of structural rationality. After offering a more uniform account of coherence requirements for Worsnip's project, I argue against several problem cases thought to prevent the normativity of these requirements. Finally, I advocate to recast Worsnip's project within a constitutivist account of agency and self-governance. I argue that doing so provides the necessary framework to explain sufficiently why agents are rationally bound to normative coherence requirements.

Dissertation Program

0.1: Coherence Requirements

Holding consistent intentions, pursuing means to ends, following our best judgment are all "coherent" behaviors. Our attitudes (intentions, judgments, beliefs, etc.) are "fitting together" in an intelligible way. The question, however, is this: Ought we to have coherent combinations of attitudes? Rather than mere "psychic tidiness,"[1] are there normative requirements of rationality prescribing such behavior? Recently, there has been significant pushback against the idea that there are such normative requirements.[2] The purpose of this dissertation is to meet this opposition and show that there are in fact normative coherence requirements governing the structure of our attitudes. Truth be told, I think that the necessary arguments for normative coherence requirements already exist. The problem is that these arguments are at best partial and scattered across various authors. Currently, I have yet to find anyone put the pieces together in order to defend normative coherence requirements sufficiently. That said, I understand my contribution to this topic to be modest—to bring together these necessary arguments (with a few amendments of my own) under one unified account.

Broadly speaking, I aim to defend a fusion between what I see as the two most promising approaches so far—Alex's Worsnip's (2020) prohibitional account, advocating for the avoidance of incoherent combinations of attitudes, and Michael Bratman's (2009) account, linking the normativity of coherence requirements to our self-governance. While primarily a defense of Worsnip's position, I believe that, when assimilating Bratman's connection to self-governance, we find a promising account that explains why coherence requirements are normative for us.

0.2: My Framework—Normative Reasons, Oughts, and Rationality

In order to defend the normativity of coherence requirements properly, I need to lay out the framework I will be using. Doing so will include the account

[1] See Kolodny (2007: 241).

[2] See Kolodny (2005), Benjamin Kiesewetter (2017), and Brunero (2020) among others.

of normative reasons (hereon shortened to "reasons")[3] and oughts I am working with as well as my understanding of normativity. Finally, since I will be defending coherence requirements as rational requirements, I will say a bit about the kind of rationality at issue in my work.

0.2a: Normative Reasons as Explanations of Oughts

I think that reasons and oughts are a package deal. When we offer a reason R to do or believe ϕ, we are pointing to the contents of R in order to explain *why* we ought to do or believe ϕ.[4] This relationship between oughts and reasons is commonly reflected in our use of language when we say things like: "We ought to ϕ *because* R" or "R is why we ought to ϕ."

Terms like "because" and "why" link our ought-claims to our defense of those ought-claims. If I claim that we ought to wait to walk across the street until the car passes, I am defending this claim by saying (or implying) that walking across the road beforehand increases the likelihood that we will get run over (I am assuming that getting run over by a car is bad). My ought-claim is that we ought to wait to walk across the street until the car passes. My reason explaining this ought-claim is that we might get run over, and getting run over is bad for us.

We can find this same relationship between normative reasons and oughts with moral and epistemic issues as well. We commonly hear statements like "ϕ-ing is the morally right (or wrong) thing to do." In such cases, ϕ-ing being morally right (or wrong) is a reason explaining why we ought (or ought not) to ϕ.[5] Similarly, in cases about belief, we commonly point to evidence as our reason R why we ought to believe ϕ. This relationship between oughts and normative reasons is commonplace. With this very broad sketch in place, I add the following caveats.

---

[3] The reason-dimension of my work involves normative reasons only. Therefore, unless otherwise stated, "reasons" are normative reasons.

[4] I am further relying on the idea that reasons are consideration of either facts or true propositions that count in favor of some response (i.e., adopting or dropping attitudes, performing actions, etc. This notion of "counting in favor," I borrow from Scanlon (2000: 50).

[5] For those who believe moral requirements can be reduced to (or are grounded in) other requirements (say rational requirements for example), normative reasons to follow moral requirements would explain *why* we ought to be moral in terms of the more fundamental requirements.

When I say that we have a reason to comply with coherence requirements, the reasons I am referring to are "all-things-considered" or "definitive" reasons.[6] These reasons have the last word, so-to-speak, in that they are not overridable. These reasons explain unconditional oughts. That said, I remain neutral whether reasons come in varying strengths. I am simply saying that, whatever strengths or scale-system exists for reasons, the reasons I am talking about are at the top. At the same time, I would *not* say that my reasons are "objective" in the sense that they ignore context. Let me explain.

As I see it, reasons are sensitive to our epistemic limitations in a given context. In context C, we have a reason R to ϕ if and only if we could be made aware of R's content in C. If, unbeknownst to me, my friend next door needs medical attention, it seems problematic to say that I have a normative reason to call an ambulance (since I have no knowledge of the problem).[7] Once I am made aware of my friend's situation, however, it seems natural to say that I do have a reason to call the ambulance. In this way, we should not be understood as failing to respond to reasons with content of which we have no knowledge.

0.2b: Coherence Requirements as Structural Rationality

In my work, I am committing myself to a dualist notion of rationality. Following Daniel Fogal and Alex Worsnip (forthcoming), there is both substantive rationality and structural rationality.[8] Substantive rationality is the more common notion of the two. We are substantively rational when we respond correctly to the correct reasons (whatever those might be). This kind of rationality governs the justification of particular attitudes and actions. Structural rationality, on the other hand, governs the *structure* of our attitudes—

[6] By "definitive" or "all-things-considered," I follow Broome's (2004) notion of final reasons rather than his pro tanto reasons.

[7] Perhaps we can say that "there exists" a reason, but that I do not have that reason. Either way, my response (or lack thereof) to this reason should not be criticizable if I am unaware of it, and thus unable to respond properly. See Worsnip (2019), Henning (2014), and Kiesewetter (2017) for further discussion on this point.

[8] See also Kiesewetter (2017: 45-80), Brunero (2020), and Worsnip (2020: 9-12) for similar versions of this distinction. As Fogal and Worsnip state the difference, "structural rationality is about whether one's attitudes fit together or cohere with each other, whereas substantive rationality is about whether one's attitudes are actually reasonable or justified" (forthcoming: 1).

i.e., what *combinations* of attitudes we ought to have or avoid.[9] This kind of rationality does not speak to the justification of particular attitudes, only the structural relationship between attitudes.

That said, I understand coherence requirements as requirements of structural rationality. These requirements govern how our attitudes should and should not fit together—that we *ought to consider coherent, and avoid incoherent, combinations*. In doing so, these requirements (once defended) will offer a normative basis for criticizing rationally those of us holding incoherent combinations of attitudes.

0.2c: Reason-Based Normativity

By "normative" I broadly mean that a requirement has sufficient evaluative authority over us. This notion of normativity can be spelled out in two ways. First, if a requirement *S* is normative, then we ought to follow *S*. Given the relation between reasons and oughts I've committed to (in Section 0.2a), I also take on a reason-based conception of normativity: If *S* is normative, then there exists a reason to follow *S*. Second, if *S* is normative, then facts about what *S* requires of us *constitute reasons* for us to have appropriately related responses. In this way, requirements are normative insofar as they give us reasons and oughts.[10] Later, in Chapter Three, I will flesh out the details of the reason-relation(s) for normative requirements. For now, I think it best not to work in the fine details until we have the general view on the table.

0.3: Our Roadmap

The general structure of my dissertation follows in three chapters. In Chapter One, I introduce six prohibitional requirements using Worsnip's definitional account of attitudinal incoherence. Following brief expositions of the first five requirements, I offer an alternative (stricter) conception of the practical enkratic requirement. In doing so, I aim to unify all six requirements

[9] In this dissertation I am treating structural rationality and coherence requirements interchangeably. Nevertheless, my intention is to remain neutral whether there are other kinds of structural rationality beyond coherence requirements.

[10] I adopt this framing of the relation between requirements and reasons from Worsnip (2020: 3-9).

under a more uniform understanding of incoherence and unintelligibility (more to be said soon on this point).

Next, in Chapter Two, I address several challenges brought against the normativity of coherence requirements and argue how the prohibitional wide-scope account succeeds in bypassing these issues. In addition to disarming scope-related issues, I follow Worsnip's position for how considerations of coherence function in our deliberations. Responding to Kieran Setiya's (2007) and Benjamin Kiesewetter's (2017) challenges, I offer a way forward that avoids relying on right kind/wrong kind reason distinctions. Finally, working from Worsnip's understanding of right-kind reasons, I dispel a possible worry linked to classifying coherence-based reasons as reasons for actions.

Finally, in Chapter Three, I argue that a sufficient defense of Worsnip's position requires recasting his coherence requirements under a constitutivist framework about agency. I argue that Niko Kolodny's "Why be rational?" question offers a more substantial challenge than Worsnip's position initially accounts for. Nevertheless, under this new constitutivist framework, I believe Worsnip's position can meet Kolodny's challenge. In doing so, I draw a necessary condition between complying with coherence requirements and being self-governing with regard to our actions. That said, while Worsnip offers us a place in our deliberations for complying with coherence requirements, I offer a normative defense of Worsnip's reasons.

Chapter One: Coherence Requirements

1.0: Chapter Introduction

The purpose of this chapter is to discuss briefly each of the coherence requirements I will later[11] defend as normative requirements of structural rationality. Following Worsnip's work, I understand these requirements to be wide-scope and, "fundamentally, as prohibitions on [incoherent] combinations of attitudes" (2020: 16). These requirements do not tell us whether a particular attitude is normatively justified or not. For their rational guidance functions at the level of attitude combinations only.[12] That said, this chapter will proceed in the following fashion. In Section 1.1, I introduce Worsnip's definitional account of incoherence as a foundation for coherence requirements. Using this framework, in Section 1.2, I discuss each requirement in turn. Let us begin.

1.1: Coherence Requirements as Prohibitions on Incoherent Sets of Attitudes

Much of my understanding of coherence requirements I import from Worsnip's (2018a) definitional account of attitudinal incoherence.[13] Following Worsnip, I understand coherence requirements as prohibitions on combinations of conflicting attitudes—i.e., prohibitions of incoherence. Defining structural incoherence, Worsnip writes,

> A set of attitudinal mental states is jointly incoherent iff it is (partially) constitutive of the mental states in question that, for any agent that holds these attitudes, the agent is disposed, when conditions of full transparency are met, to give up at least one of the attitudes. (2018a: 188)

As I understand Worsnip, we are coherent when, under conditions of full transparency, we are not disposed to drop any of our attitudes. Under conditions of "full transparency," we hold a combination of mental states

[11] See Chapters Two and Three for this defense.

[12] Nevertheless, in collaboration with substantive reasons, these requirements do have a role to play in determining our attitudes (see Chapter Two, Section 2.2c).

[13] Hereon I shorten "attitudinal mental states" to "attitudes."

without the aid of "self-deception, mental fragmentation, or any failure of self-knowledge" pertaining to our attitudes (188). In other words, we are fully transparent about our attitudes when we are sufficiently "clear eyed" about how those attitudes relate to one another.

This approach has intuitive appeal. At first glance, we can recognize the force of coherence standards quite clearly if, under conditions of full transparency, violating them invokes severe mental stress or dissonance. If, under full transparency, it is either impossible or at least very difficult to hold some set of attitudes together, this fact alone gives credence to the idea that this set is incoherent.[14] In what follows, I commit to a stricter understanding of "disposed" than Worsnip does. From my point of view, when we are fully transparent about a set of incoherent attitudes we hold, we *must* either (1) loose our transparent (clear-eyed) understanding of our attitudes or (2) drop at least one of our attitudes. In other words, we *cannot* transparently hold a set of incoherent attitudes. This difference in my position from Worsnip's will be most evident when discussing the practical enkratic requirement (Section 1.2f). That said, using Worsnip's framework, let us begin going through some specific requirements.

1.2: Six Coherence Requirements, Six Prohibitions

Understanding coherence requirements as prohibitions on structural incoherence, we can capture six canonical requirements. As a brief reminder, these requirements are "wide-scope," meaning that they are "stated using a deontic operator that takes wide-scope over a material conditional" Worsnip 2020: 17).[15] Furthermore, in order to emphasize their prohibitional nature, I have chosen to write out these requirements in term of what attitude combinations we *ought to not* have.[16] I believe that writing these requirements

[14] So far, I have said nothing to defend the normative status of attitudinal coherence. The related cognitive strain of holding particular combinations of attitudes offers compelling evidence that incoherence is present, not that prohibiting such a situation involves a normative requirement.

[15] Said differently, we can satisfy such requirements by either adopting consequent attitudes or revising antecedent attitudes. However, *how* we ought to satisfy these requirements will be determined by substantive reasons (i.e., moral, prudential, epistemic, etc.). As I will discuss (in Chapter Two, Section 2.2b-2.2c, see also 2.1b), these requirements *do not*, on their own, allow that adopting or dropping any particular attitude is a sufficiently rational option.

[16] Writing out the requirements in this way is a simple application of De Morgan's Laws (Chartrand et al. 2013: 112).

this way better emphasizes their governance over *attitude combinations, not particular attitudes*. Lastly, I understand these requirements to be sensitive to an agent's beliefs irrespective of the truth-status of those beliefs.[17] I will now discuss each briefly.

1.2a: The Instrumental (Means-Ends) Requirement

The instrumental requirement is commonly understood as follows: if we intend some end ϕ, then rationality says that we ought to either (a) intend what we believe to be the necessary means ψ for ϕ-ing, (b) abandon the belief that ψ-ing is necessary for ϕ-ing, or (c) abandon our original intention to ϕ. This way of writing the requirement showcases the options that our *substantive reasons* will weigh in on. Rather than "equally rational" options, we should instead understand (a)-(c) as three different ways in which this structural requirement may be *combined* with substantive requirements of evidence, morality, etc.[18] Strictly in terms of structural rationality, however, I think that a better way of writing this requirement illustrates what it requires of us *all the time and in every context.* That said, I write out the instrumental requirement as follows:

> *Instrumental Requirement* (IR): we ought to not [intend to ϕ, and believe that ψ-ing is a necessary means for ϕ-ing, and fail to intend to ψ].

Written as such, we see that satisfying this requirement is *really* about avoiding a single combinational form of incoherent attitudes—where we do not intend the means we believe is necessary for our intended end. Even though we can rephrase this requirement as a disjunctive ought, we should think of IR as a prohibition on instrumental incoherence.[19] In order to show that IR truly

---

[17] Once again, I am relegating the justification of particular attitudes (the truth status of beliefs, the correctness of judgments, the permissibility of particular intentions, etc.) to substantive rationality.

[18] I discuss this point further in Chapter Two: Sections 2.1b and 2.2c.

[19] Perhaps someone may object, claiming that a requirement of rationality should tell us what we ought to do, and understanding the instrumental requirement as such fails in this regard. However, I think such an objection is easily dismissible in the following way. The kind of rationality this requirement expresses is *structural*, not substantive. That is, if coherence requirements are about avoiding incoherent (or holding coherent) combinations of attitudes, then such requirements govern *attitude combinations*, not the justification of *particular attitudes*.

expresses incoherence (in the way defined above), however, we need to show why, under conditions of full transparency, we would be disposed to drop at least one of the attitudes in the prohibited set above. The answer is straight forward: because cases violating *IR* under conditions of full transparency become unintelligible. To illustrate this point, consider an example.

Imagine I decide that, an hour from now, I will buy groceries in town. Also imagine I believe that, given the distance to town, driving is necessary to reach the grocery store in time. Now imagine that I sincerely tell you the following: "I intend to grocery shop in town in an hour, I believe that driving is necessary to reach town in time to grocery shop in an hour, but I do not intend to drive to town." Something is amiss here. For, as Worsnip points out, in such a case I "would be confused about what it is to have an intention" (191). I think Worsnip is correct.

This conclusion stems from the common understanding that intentions are distinct from attitudes like wishes and desires (191). When I say that "I intend to ϕ," I have done more than simply express what I am motivated to do. I am expressing my decision to ϕ. Attitudes such as desires and wishes function in our decisions only antecedently, as *possible motivations*. Intentions, on the other hand, should be understood as *guiding attitudes*[20] that "tell the story" of what we choose and aim to do.

That said, it is unintelligible to talk of having an intention to ϕ *while not intending to do what I believe must be done in order to* ϕ. For, regardless of what independently constitutes ϕ-ing, my beliefs about how to ϕ set a conceptual standard I hold myself to when I decide to ϕ. If, by my own beliefs, grocery shopping in town in an hour requires driving to town (to get there in time), I cannot intelligibly claim that I am intending to grocery shop in town in an hour unless I intend to drive there. For doing otherwise violates the very conditions

---

That said, it is reasonable to stipulate that such requirements are not meant to tell us which attitudes to adopt or drop. This latter process (deciding what we should do in a given context) should be relegated to substantive rationality in the sense of responding correctly to substantive normative reasons (pertaining to issues of epistemic evidence, morality, prudence, etc.). What a requirement of structural rationality should prescribe is that, whatever path our substantive reasons lead us down, this path ought not to be a structurally irrational one. More to be said on this point in Chapter Two, Section 2.1.

[20] By "guiding" I am following Bratman (1987, 2009) and others who hold that intentions are constitutive parts of partial plans aimed toward the successful completion of ends.

that *I believe* are necessary for grocery shopping in town in an hour. And in doing so, I would effectively be claiming that I intend to do what I believe is impossible to do. Such a claim (as well as the corresponding combination of attitudes) will always be unintelligibly irrational.

Given the above, we can see how violating IR leads to an unintelligible illustration of what we choose and aim to do. When we do not follow IR, we lose a clear sense of what we are intending to do. Now, this is not to say that instrumental irrationality is unintelligible full-stop. Rather, instrumental irrationality is unintelligible *under conditions of full transparency* (2018a: 190-191). It is a common story that we intend to do something, but then get distracted or somehow end up not intending the necessary means for that something. Perhaps our attitudes are fragmented[21] such that we hold the combination of attitudes IR prohibits but do not "connect the dots" so-to-speak between our attitudes. But when we are forced to consider such a combination under conditions of full transparency (as the above example does) the picture becomes unintelligible. And this unintelligibility is what grounds our disposition either to lose our "clear-eyed" perspective on our attitudes or to change one of the attitudes. Thus, the instrumental requirement is incoherent in the relevant sense. This requirement fits the disposition account of incoherence above and functions as a prohibition on instrumental incoherence in all cases.

1.2b: The Intention Consistency Requirement

A similar, albeit far simpler, story can be told about why we ought not to hold intentions that we believe are not copossible. This requirement is commonly known as the intention consistency requirement.[22] Once again I formulate this requirement in terms of what we ought to not do:

> *Intention Consistency Requirement* (ICR): we ought to not [hold intentions i_1, i_2, …, i_n and believe that i_1, i_2, …, i_n are not copossibly achievable].

[21] See Stalnaker (1984) and Egan (2008) for discussions of mental fragmentation.

[22] Among others, Bratman (2009: 413) discusses this requirement.

Similar to the instrumental requirement, ICR instructs us that we ought to avoid attitude combinations that, under conditions of full transparency, become unintelligible to us. In the case of ICR, the attitudes in question are two or more intentions and the instrumental belief that these intentions are not copossibly achievable.

Consider once again the grocery shopping example. However, this time imagine that, in addition to holding the intention to grocery shop in town in an hour I also hold the intention not to drive to town today. Since, from the original example I believe that driving to town is necessary for grocery shopping in an hour, I also believe that the two intentions I now hold are not copossibly achievable. If I succeed in grocery shopping in town in an hour (from my set of beliefs), this result will obtain only because I drove to town in time. But if so, then I (intentionally) fail to succeed in my intention to not drive to town today. Vice versa, succeeding in not driving to town today ensures (from the vantage point of my beliefs) that I (intentionally) will not succeed in my intention to grocery shop in town in an hour.

Thus, in either case, effectively I am intending to do something I believe is impossible to do. Since, for at least one of my intentions, I believe that I am intentionally making it impossible to pursue what I aim to do. In such a case, I am still violating the instrumental requirement (since I do not intend to drive to town), but I am also violating ICR. Now, from our prior discussion of the instrumental requirement, we saw how, under conditions of full transparency, intending to do something we do not intend to succeed at doing is unintelligible. The same unintelligibility (under conditions of full transparency) holds when violating ICR—for in both cases we are intending to do what we believe is impossible to do. In either case, when we are "clear eyed" so-to-speak about our attitudes, this combination becomes unintelligible.

That said, ICR holds under Worsnip's definition of incoherence. Since, when considered under conditions of full transparency, this combination of attitudes become unintelligible to us. We can follow Worsnip in saying that, in the light of day, any claim to hold this incoherent combination of intentions and instrumental beliefs misunderstands what it means to hold an intention in the first place. Moving forward.

1.2c: The Modus Ponens Requirement

Now let us talk about modus ponens. We commonly accept that if p, and $p \rightarrow q$, then q. As a coherence requirement, modus ponens is commonly and prescriptively applied to cases of belief formation. If we believe truly that the earth is round, and we believe round things cannot be flat, then we ought to believe the earth is not flat. This sounds reasonable. However, before going further, a clarification is needed.

In the current context, the prescriptions of a modus ponens requirement should function over combinations of beliefs (and belief entailments),[23] not particular beliefs. As discussed briefly in the beginning of this chapter, the kind of rationality I am seeking to express with coherence requirements is structural, not substantive. Such requirements do not prescribe the adoption or avoidance of particular attitudes, but instead prescribe at the level of combinations of attitudes. Thus, if modus ponens is to be understood as a coherence requirement under this account, it must function within this same framework. That said, we *should not* think of the modus ponens requirement in this context as *telling us what particular beliefs* we rationally ought to (or ought not to) have given our other beliefs. Rather, modus ponens *should* be understood here as telling us what *combinations of beliefs* we rationally ought to have (or ought to avoid).

Finally, as done with the other requirements discussed, I will lay out a modus ponens requirement in a prohibitional fashion. This requirement will prescribe what combinations of beliefs we rationally ought to avoid. That said, I propose the following:

> *Modus Ponens Requirement* (MPR): we ought to not [believe p, and believe $p \rightarrow q$, and not believe q].

Here, there is no prescription of particular attitudes. For the prohibition functions at the level of belief combinations. Specifically, MPR is telling us that we ought not to have combinations of attitudes where our absence of a belief is

[23] Hereon belief entailments are implicitly included when talking about one's beliefs.

not consistent with the entailments of our other beliefs.[24] To illustrate what MPR prohibits, consider a more perplexing version of the earth example.

Imagine I have the following combination of attitudes $\{a, b, c\}$. I (a) believe the earth is round, (b) believe that round things cannot be flat, but (c) do not belief that the earth is not flat.[25] At this point, regardless of other concerns about what particular beliefs I should or should not have, it should be quite clear that there is a *structural problem*.[26] My beliefs a and b (from my perspective) entail that the earth is not flat. Nevertheless, I lack this entailed belief. More to the point, (c) my lack of believing that the earth is not flat is inconsistent with the entailments of my beliefs (a and b). So, $\{a, b, c\}$ is jointly incoherent.

Further to the point, cases like $\{a, b, c\}$ are jointly incoherent in the sense we are aiming at because, under conditions of full transparency, holding such attitude combinations become unintelligible. In cases like $\{a, b, c\}$, there is no understanding for why, given a and b, I hold c instead of the belief that the earth is not flat. We can only intelligibly describe such a situation by saying things like "I have not connected the dots somehow" or "I have not put the pieces together to see the problem." Statements like these imply cases like $\{a, b, c\}$ "make sense" only when full transparency is not present, when I am fragmented somehow. But if I were clearly to state $\{a, b, c\}$, or if I "did connect the dots" so-to-speak, it becomes unintelligible how I would hold such a combination of attitudes. In this way, like with intention cases in the prior section, clear-eyed inconsistency is always incoherently unintelligible.

---

[24] By prohibiting incoherent belief combinations, MPR lends to the common view that incoherent belief combinations signal the presence of false beliefs.

[25] Here I am understanding a lack of belief as an attitudinal mental state.

[26] As Fogal and Worsnip (forthcoming) point out, the rational failure of holding unjustified attitudes is distinctly different from the rational failure of holding incoherent belief combinations: "Consider someone – let's call him Tom – who believes that he is Superman, and believes that Superman can fly. In believing that he's Superman, Tom has a belief that flies in the face of his evidence (or so we may safely stipulate). Moreover, it's very natural to describe this belief as irrational. But now let us add another piece of information about Tom: he believes he can't fly. (Suppose he's tried several times, without success.) Intuitively, given this further piece of information, we can now see that Tom is irrational in a second respect: specifically, he fails to believe an obvious consequence of his other beliefs" (forthcoming: 6-7).

That said, MPR's naturally fits within the larger conception of coherence requirements. For its prohibition of such attitude combinations is exactly a prohibition of the kind of incoherence I am working with under Worsnip's definition. With MPR now in place, it is time to move to the next requirement.

1.2d: The Transitive Preference Requirement

The next requirement I wish to discuss is what is known as the transitive preference requirement. Simply put, this requirement is usually understood to claim that preference reasoning ought to function transitively. If I prefer a to b, and I prefer b to c, then I ought to prefer a to c. In keeping with the current framework, however, I propose the following: that preference reasoning ought not to be intransitive. Spelled out, we get the following:

> *Transitive Preference Requirement* (TPR): we ought to not [prefer a over b, and prefer b over c, and prefer c over a].

As we can see, the combination type of preferences TPR prohibits is cyclical (or intransitive). But are attitude combinations involving cyclical preferences incoherent? Yes, I believe they are. To illustrate why, first consider the following example.

I am trying to decide what I will do today. Before me are three options: (a) working on my dissertation, (b) mowing the lawn, and (c) watching the final season of *Dark* on Netflix.[27] When considering the pairwise comparisons one at a time, it is easy to imagine the following circular reasoning. Because my dissertation is necessary for both my academic and professional advancement, I prefer working on my dissertation to maintaining the lawn. Between mowing the lawn and watching television, however, television seems rather frivolous while maintaining the lawn attends positively to health concerns (limiting the presence of ticks and other disease carrying insects living in tall grass). So, I prefer mowing the lawn to watching the final season of *Dark*. Now, between working on my dissertation and watching the final season of *Dark*, I think that

[27] The spirit of this example I borrow from Worsnip (2018a: 191-192).

watching television is far less work than philosophy, and my deadline is still far off. So, I end up preferring watching *Dark* to working on my dissertation today. While seemingly commonplace, this example shows intransitive preference reasoning: (a) over (b), (b) over (c), and (c) over (a).

The problem is that reasoning of this kind fails to establish clearly any conclusion over what to do. For while there is a clear sense of preference between *any two options*, the above pairwise preferencing fails to show my preference between *all the options*. In Bratman's terms, I lack any clear standpoint on what I prefer to do today (2009: 430-431). However, reaching a conclusive decision on what I prefer to do is exactly the aim of preference reasoning. Thus, intransitive preferential reasoning is a failure to reason preferentially at all. Such reasoning is reasoning without a conclusion.

Furthermore, intransitive preferential reasoning becomes unintelligible under conditions of full transparency. As Worsnip's points out, the crucial insight behind the intelligibility of the example above is that we think about "the pairwise comparisons one at a time" (Worsnip 2018a: 192).[28] As long as I do not "connect the dots" between my preference pairs, the incoherence in play is sustainable. However, when I consider these preference pairs all together (now in a fully transparency sense), my intransitive preferencing becomes unintelligible. If, in the context of deciding what to do, I say, "I prefer dissertation work to lawn mowing, prefer lawn mowing to watching television, but prefer watching television to dissertation work," this seems nonsensical. In Worsnip's words, "this sounds like a joke" (192). Once again, we have a clear sense in which, under conditions of full transparency, the above combination of attitudes (in this case preferences) is unintelligible—and so fit Worsnip's definitional account of incoherence. That said, TPR prohibits incoherent preferences. Moving forward.

[28] The difference between *pairwise considerations* and *considerations covering the entire set* of options offers a useful illustration about conditions of full transparency: The former pairwise preferences lack full transparency because there is no "connecting the dots" between these preference pairs whereas the latter consideration of the entire set brings the preference pairs into conjunction with one another.

1.2e: The Doxastic Enkratic Requirement

Now consider the doxastic enkratic requirement. Commonly understood, enkratic failure (or akrasia) occurs when we act or believe against our own best judgment.[29] In terms of evidence and belief, then, the doxastic enkratic requirement prescribes that we ought to believe what our evidence-based judgments support. The attitudes in question, then, are our beliefs and evidence-based judgments about those beliefs.

As a prohibition of doxastic incoherence, then, I understand this enkratic requirement as follows:

> *Doxastic Enkratic Requirement* (DER): we ought to not [judge that we lack sufficient evidence for believing p and believe p].[30]

The idea here is simple. Evidence aside, if I *judge* that I lack sufficient evidence for believing the earth is flat, then by *my own standard of judgment* I ought not to believe the earth is flat. Thus, the question is whether we can intelligibly hold combinations of attitudes that violate DER. Without full transparency about these attitudes, I believe we can hold such combinations. However, once we sufficiently examine our attitudes (consider them together under conditions of full transparency), holding such combinations of attitudes becomes unintelligible.

Imagine that, contrary to my judgment about the evidence, my everyday experience of walking and driving instill in me the non-occurrent belief that the earth is flat. In this case, while I judge that I lack sufficient evidence to believe the earth is flat,[31] I still feel that my everyday experience is telling and so, deep

[29] Some believe that akrasia does not exist (See Richard M. Hare [1952] for example). If this is the case, that akrasia does not exist, then our work is done for us. We need not be concerned about structurally rational requirements prohibiting akrasia. However, where there remains serious debate on this issue, I work under the assumption that akrasia is possible (just in case).

[30] The alternative case, where we judge that we have sufficient evidence to believe p yet we do not believe p, is covered by the modus ponens requirement: We ought not to [believe a, and believe $a \rightarrow b$, and not believe b]. In this case, a is our recognition of the evidence pertaining to p, $a \rightarrow b$ is our judgment that the evidence pertaining to p is sufficient for believing p, and b is p.

[31] Say because I've seen pictures of the earth from the moon and the ISS, studied the mathematics concerning sunlight displacement over spheres, and generally believe the available scientific data.

down, I still believe the earth is flat. While undoubtedly odd, this case unfortunately doesn't seem that farfetched for some. First, because there are people that believe the earth is flat. And second, because even today we still incorporate nomenclature that seems to contradict widely accepted scientific realities (why "sunrise" and not "earthrise"?). Regardless of our feelings about my mental state in this example, at least the situation seems intelligible as long as I keep from "connecting the dots." However, the example's intelligibility quickly disappears under conditions of full transparency—for holding this same combination of attitudes now leads to an explicit contradiction.

If, under conditions of full transparency, I both believe the earth is flat and judge that I lack sufficient evidence for this belief, I am effectively holding contradicting judgments about the available evidence. Under conditions of full transparency, we hold beliefs we think are justified, and these justifications are our judgments about the evidence supporting those beliefs.[32] Thus, under conditions of full transparency, maintaining the belief that the earth is flat entails that I think this belief is justified—which is just to say that I judge that I have sufficient evidence for believing the earth is flat. But now I have a problem. For my position is now that I judge both that I do and do not have sufficient evidence to believe the earth is flat. So, I explicitly hold contradicting judgments, which is unintelligible.

That said, if doxastic akrasia is possible, it seems possible only outside conditions of full transparency. For when we explicitly "connect the dots" between a belief we hold and our judgment that this belief is unjustified, something's got to give. This unintelligibility of doxastic akrasia in explicit cases (with full transparency over the one's attitudes) shows how such violations of DER are incoherent in Worsnip's (2018a) sense. Now, let us move on to the last coherence requirement.

1.2f: The Practical Enkratic Requirement

Lastly, let us briefly talk about the pragmatic enkratic requirement. Like with its doxastic counterpart, we violate the practical enkratic requirement

[32] As Worsnip puts the point, "our beliefs, when they are formed reflectively in ways that are transparent to us, are controlled by our judgments about the evidence" (2018a: 194).

when we go against our best judgment. In the practical case, however, the general idea is that our intentions ought to fall in line with our best judgment of what we ought to do.

Following my framework, then, I propose we understand the practical enkratic requirement as a prohibition on attitude combinations where our intentions contradict our judgments about those intentions. That said, the requirement is as follows:

> *Practical Enkratic Requirement* (PER): we ought to not [judge that we ought to ϕ and not intend to ϕ].[33]

We violate PER (are practically akratic) when we fail to intend what we judge we ought to do. Simple enough as a rule, it seems. However, there remains an open question whether we can classify practical akrasia alongside other instances of incoherence. I think we can. However, to do so, I must part ways with Worsnip's (2018a) assessment of practical akrasia. Nevertheless, I believe my approach here better assimilates practical akrasia into Worsnip's (2018a) definitional account of incoherence. The initial problem is this.

It is generally held that when we are practically akratic, when we violate PER, we are "clear eyed" about our incoherent set of attitudes. If this is true, then it seems that violating PER can be done under conditions of full transparency. If so, then it seems we have a conceptual problem with our account of incoherence. For instances of incoherence under conditions of full transparency are supposed to be unintelligible. That said, if we can be reflective and clear-eyed when we violate PER, such cases do not seem unintelligible at all.

Worsnip's answer to this problem is to claim that practical akrasia is a relatively weak form of incoherence. As Worsnip writes, we "should allow that incoherence is something that comes in degrees" (2018a: 199) and that in "less incoherent cases, such as akrasia, the disposition is weak enough" that we can

[33] For alternative cases where we judge we ought not to ψ and we fail to intend not to do it, substitute $\neg\psi$ for ϕ. This substitution will also cover stronger cases where we intend to do something we judge we ought not to do since intending to ψ entails not intending to $\neg\psi$.

transparently sustain this incoherent set of attitudes. While this concession about practical akrasia does not itself make Worsnip's account of incoherence untenable, it does throw a wrench into the mix so-to-speak. Here is why.

In all other cases of incoherence, Worsnip's defense of his dispositional definition was that, when under conditions of full transparency, these cases became unintelligible. This is a compelling argument because, when viewed under conditions of full transparency, if the cases are unintelligible, we cannot reasonably think that such occurrent attitude combinations can be held in the mind. However, by saying that practically akratic attitude combinations can be transparently held together, we are saying that Worsnip's argument for prior cases of incoherence cannot apply here (or at least not in a general way). For we are saying that practically akratic people who are fully transparent about their attitudes *can be intelligible*.

That said, practical akrasia's incoherent status becomes less than obvious, and if this state really is incoherent, this claim begins to seem ad hoc. For it seems that we are left without a strong argument for why practical akrasia *is* incoherent.[34] That said, we either allow that practical akrasia is unlike other incoherent states, or we say that "somehow" practical akrasia is "just weakly incoherent enough" that it can be transparently sustained. The first option begs the question whether practical akrasia really is an incoherent state, and the second option seems very ad hoc in the current framework because it no longer seems that there is much argument for how practical akrasia fits the dispositional account.

Now, perhaps I'm being ungenerous to Worsnip's analysis of practical akrasia. However, I think there's a better option to take—an option that allows for cases of what we *think* is clear-eyed practical akrasia while also showing that such cases are in fact unintelligible. This option, I think is promising and, if adopted, adds weight to Worsnip's overall dispositional account.

[34] Of course, we can still say that, with practical akrasia, the disposition remains (albeit weakly) to drop at least one of the attitudes in question, and in this way these cases do fit the dispositional account. However, we can have dispositions to drop attitudes for other reasons than the attitude combination being incoherent (we decide to change our minds or any number of substantive reasons). Without showing the unintelligibility of a fully transparent practically akratic person, we lose a strong argument as to why practical akrasia *is* incoherent. Concerns like this are why I believe that John Brunero comments that Worsnip's "theory has difficulty allowing for the irrationality of clear-eyed akrasia" (Brunero 2020: 207, fn. 51).

As a preliminary note, we should first understand that, in cases of practical akrasia, the judgment involved must be an *all-things-considered* judgment. By "all-things-considered" I mean that, given some situation or context C, I weigh all my various reasons, concerns, etc. and judge that, all-things-considered, I ought to ϕ. This point is important because, without an all-things-considered judgment present, my lack of intention following a judgment to ϕ is not necessarily incoherent.[35] Without an all-things-considered judgment, the fact that we lack the corresponding intention to some judgment can simply be explained by further reasoning, leading to a different decision (ultimately our all-things-considered judgment).[36]

That said, if being practically akratic requires that we fail to intend what we all-things-considered judge we ought to do, *then we end up with a contradiction of beliefs about our reason(s) corresponding to this judgment*. This contradiction arises because an all-things-considered judgment to ϕ entails the belief that we have an all-things-considered normative reason R to ϕ.[37] However, if our lack of intention to ϕ really is something we choose to do (an attitude we have self-knowledge about), then this fact entails that we believe we

[35] Consider this: Based on my commitments as a graduate student, I ought to work on my dissertation. Now say I have a sick family member needing care. Most would say that I ought to take care of this person. The point here is simply that I have various reasons and judgments about what I ought to do based on different parameters, standards, and commitments. The fact that I intend to take care of my family instead of working on my dissertation tonight does not make me practically akratic. For in this case I must have further judged that taking care of family *outweighed* being a diligent student tonight. Furthermore, whatever I do intend can be linked to my final all-things-considered judgment. Hence, I can always be understood as practically *enkratic*. This same kind of explanation can be made about failing to do the seemingly "right thing." While in this case I may be *morally* criticizable, it is not clear that I am *rationally* criticizable. Rather than intend to do what I judge to be the "right thing" (based on some standard), I may also judge (based on a standard of self-interest) that doing the "right thing" is too burdensome for me or disallows me the satisfaction of my desires, etc. Granted, this case may show me to be *morally* bankrupt, but it does not necessarily show me to be incoherent. It merely shows that I judged to do something else based on a different standard from my first judgment (and I further judged that my second judgment outweighed the first).

[36] What happens if we simply ignore our judgment (one that is not all-things-considered)? Would this case be practically akratic? Not necessarily. If I judge (not all-things-considered) that I ought to ϕ, but fail to intend to ϕ there is still the question of *why I failed to intend to* ϕ. If a reason and/or judgment can be cited, then this fact undermines the idea that I am being practically akratic. For it seems that I simply weighed my judgment with an alternative one. However, if no reason or judgment can be cited, then it seems that my failing to intend to ϕ is not something I have done (or failed to do) at all. In either case we fail to arrive at practical akrasia.

[37] R could also represent a collection of normative reasons we believe collectively outweigh the alternatives. The present argument follows regardless of this fact.

have another normative reason R_0 either weighing against ϕ-ing or weighing in favor of some alternative to ϕ-ing. The crucial point is this: the fact that we are swayed by R_0 *over* R shows that we believe R_0 *overrides* R.[38] However, if we believe that R_0 overrides R, then we do not believe that R is an all-things-considered normative reason to ϕ. Now we have a contradiction. For now, it is the case that we believe both that R is and is not an all-things-considered normative reason to ϕ.

As it stands, then, holding this contradiction of beliefs under conditions of full transparency is simply unintelligible. It would be intelligible if we did not "connect the dots" about *why* we fail to intend to ϕ. If, instead of ϕ-ing, I end up watching television all night (but do not know why I do this), we could say that I am being practically akratic. However, in this case, since I fail to have self-knowledge about my attitudes,[39] I am not under conditions of full transparency about those attitudes. If, however, I am under conditions of full transparency, then I am fully aware that I believe both that R is and is not an all-things-considered (non-overridable) reason to ϕ. Holding such a contradiction of occurrent beliefs is just as unintelligible as holding occurrent intentions believed to be non-copossible.

As I have argued, clear-eyed practical akrasia is just as incoherent as the other forms of incoherence we have been discussing. However, now I am faced with a dilemma. If the above argument holds water, how do I explain why cases of practical akrasia seem clear-eyed?

Practical Akrasia and Trading on Perspective

I think an answer is this: in examples of what seems like clear-eyed akrasia, we unknowingly trade across different standards between our all-things-considered judgment and what we choose to do. Consider an example of

[38] If we did not believe that R_0 overrides R, then we would lose any sense of *why* we followed R_0 rather than R, which means we fail to have full self-knowledge about our lack of intention to ϕ. But if this is the case, we are no longer under conditions of full transparency (Worsnip 2018a: 188). Thus, believing that R_0 overrides R is necessary to stay within the full transparency requirement of the example.

[39] For example, I lack self-knowledge both about why I failed to intend to ϕ and why I instead intended to watch television.

practical akrasia by Bratman. In this example, a depressed Sam drinks and has the following conversation:

> Drinking: 'Look here. Your reasons for abstaining seem clearly stronger than your reasons for drinking. So how can you have thought that it would be best to drink?" To which Sam replies: "I don't think it would be best to drink. Do you think I'm stupid enough to think that, given how strong my reasons for abstaining are? I think it would be best to abstain. Still, I'm drinking.' (1979: 156)

At first look, this example seems to illustrate exactly the kind of clear-eyed akrasia we have come to know. That said, I will now argue for an alternative interpretation—that Sam *is not* clear-eyed about his attitudes *because* he lacks full self-knowledge about the entailments of his judgment against drinking. In doing so, I argue that Sam violates Worsnip's conditions of full transparency. Okay, how is this possible?

I think the story goes like this: Sam unknowingly trades between two different standards—one leading to his judgment against drinking, the other leading to his decision to drink. If so, then what allows Sam to act intentionally against his all-thing-considered judgment is that, due to his depression, he irrationally *fails to occupy the neutral perspective wherein* his reasons for drinking *can be overridden* by his reasons against drinking. Sam's depression, in effect, separates Sam's inconsistent beliefs:

1. That he has an all-things-considered reason R against drinking, and
2. That he has a reason R_0 that overrides R.

When Sam is asked about his reasons against drinking, he recalls that, all-things-considered, he ought not to drink. Nevertheless, referring back to his drinking, Sam unknowingly shifts perspective back to his depression wherein his reasons for drinking, *now isolated away from his reasons against drinking*, have no opposition. In this way, Sam holds both inconsistent beliefs above (albeit isolated away from one another). What is irrational about Sam, then, *is not* that

he has no reasons for drinking; rather, *it is that he fails* (*due to his depression*) *to make the connection* that his all-things-considered judgment against drinking entails that R_0 *cannot* override R.

That said, I now offer two reasons for this interpretation.[40] First, if my earlier argument against clear-eyed akrasia holds water, then understanding Sam to be "clear-eyed" about his attitudes requires taking the position that holding occurrent inconsistent beliefs about our reasons is intelligible. But this just seems false because it makes no rational sense to believe that $p \wedge \neg p$ is true.

Second, we already have a precedent for unknowingly trading between different standards when we are not clear-eyed about our attitudes. Recall my discussion of the transitive preference requirement (section 1.2d). In my Worsnip-inspired example, I was trying to decide between three options: (a) working on my dissertation, (b) mowing the lawn, and (c) watching the final season of *Dark* on Netflix. In doing so I reasoned in a circle: I preferred

1. Dissertation work over lawn mowing (because of academic and professional considerations),
2. Lawn mowing over watching season three of *Dark* (because of health concerns about ticks), and
3. Watching season three of *Dark* over dissertation work (because I am lazy and watching television is less work than philosophy).

The purpose of preference reasoning, however, is to arrive at a decision about what I prefer to do. That said, my reasoning is intelligible *only* because I do not consider all three options together. As long as, unbeknownst to me, I keep the preference pairs separated, I *fail to notice that I am reasoning by different standards* with different pairs. In order to make a decision about my three options, however, I need to consider all three under the same standard. This unity of perspective across all options is what facilitates the linearity of preference reasoning.

[40] I do want to reiterate that I am interpreting Sam as *choosing* to drink. He is not overwhelmed by his desires, nor is he being forced in any way. I understand Sam's drinking as an intentional action. For if Sam is not intentionally drinking, he already lacks self-knowledge about (the reason) why he is drinking—violating Worsnip's conditions of full transparency.

The same holds true for the case about Sam's akratic drinking. From Sam's current perspective, his drinking alleviates his depression and so, *from this perspective*, drinking is *prima facie the best thing* to do (Bratman 179: 168). When Sam is challenged on this point, he swaps back to his "neutral" perspective whereby some neutral standard[41] he once again agrees that not drinking is all-things-considered best. However, directly following this claim, Sam once again falls back into his depression-oriented perspective wherein he chooses to drink based solely on the reasons in favor of doing so.

In this way, we can understand Sam's behavior as both irrational and intentional. For as Bratman puts the point, "Sam's irrationality in the original case does not preclude it from being a case of reasoning but is, rather, a feature of that reasoning" (169-170). My claim, then, is that this irrational "feature" of Sam's reasoning derives from his unconscious swapping back and forth between perspectives and standards of judgment.[42] As Sergio Tenenbaum similarly remarks, it is possible to recall something we have judged to be true while also failing to recall the normative force behind this judgment (2007: 270).[43]

The above interpretation of Sam allows us to understand practical akrasia as being unintelligible under conditions of full transparency. For Sam fails to have full self-knowledge of his attitudes because he fails to recognize the entailments of these attitudes. Sam fails to realize that his all-things-considered judgment against drinking, together with his intentional act of drinking, entails that he believes that he has a reason against drinking that is both overridable and not.

---

[41] A question remains as to how to account for this neutral perspective where we make all-things-considered judgments. See Derek Baker (2018) for a discussion of this problem as well as skepticism for its solution. Nevertheless, I leave this issue open. I believe that having some way of arriving at all-things-considered reasons and oughts is a precondition for the discussion of practical akrasia.

[42] Bratman seems to point to the disconnection itself in Sam's reasoning (his lack of explanation) as to *why* he chooses reasons to drink over stronger reasons against doing so (1979: 167-170). I go just one step further and say that Sam's disconnection derives from his unconscious swapping back and forth between standards of judgment. For Bratman and I agree that it is possible to "explain why Sam reasons in this admittedly irrational way by, for example, appealing to the depth of his depression, the strength of his desire to get out of it, and, perhaps, to the way in which his depression dulled his appreciation of his reasons for abstaining" (170).

[43] Tenenbaum goes on to say that, in cases of practical akrasia, we can look to our emotions and desires as the interfering mechanisms that can place us in a "perspective whose inadequacy [we][…] only vaguely understand[…] right now" (277).

This is good news, I think. For if the above analysis is correct, then practical akrasia falls comfortably within Worsnip's (2018a) dispositional account alongside the other forms of attitudinal incoherence. At the same time, this interpretation also allows us to treat practical akrasia as a special case of incoherence. For in all other cases, Worsnip's position on incoherence functions on the idea that any

> apparent intention, belief, or other attitude that, given the agent's other mental states, will put her in sustained, transparent violation of a coherence requirement does not really count as an instance of that attitude: it is not intention or belief proper, but something less—for example, instead of an intention, a wish or desire; instead of a belief, a pretense or a supposition. (2018a: 194).

However, with cases of practical akrasia we can still say that Sam has a genuine all-things-considered judgment against drinking and a genuine intention to drink (since both these attitudes are determined from separate perspectives).[44] With Sam, the case is simply that he fails to have full self-knowledge of the fact that having these attitudes entails inconsistent beliefs about the status of his reasons against drinking. That said, I believe this interpretation can accommodate Worsnip's moderate position between more hardline accounts that either deny the existence of practical akrasia or deny its incoherent status.[45]

1.3: Chapter Conclusion

Working within Worsnip's dispositional account of incoherence, I have spelled out six wide-scope coherence requirements. Each of these requirements prohibit the holding of specific incoherent attitude combinations. Departing

[44] Of course, Sam's intention is arrived at within an emotionally compromised perspective, but nonetheless, his reasons for intending to drink are genuine. So, the problem is not that the intention is not genuine, but rather that the intention may be ultimately unjustified due to the normative status of the reasons that instantiate it.

[45] For the former, see Hare (1952: 19–20, 164–170). For the latter, see Arpaly (2000). It should be noted that my position treats practical akrasia as something that happens when we are *not* clear-eyed (not under conditions of full transparency) about our attitudes. If others define practical akrasia *as clear-eyed akrasia*, then I deny this phenomenon for the reasons I have given. In this case, I will speak of practical akrasia* rather than the clear-eyed conception.

from Worsnip, I have argued that, under conditions of full transparency, cases of practical akrasia are just as unintelligible as the other forms of incoherence discussed. If correct, I believe my treatment of practical akrasia offers a more unified account of incoherence for Worsnip's project.

With this account of coherence requirements on the table, I now move forward in defending their normative status. In Chapter Two, I address several arguments made against the normativity of coherence requirements. Here, I also make my case for adopting a wide-scope conception of these requirements. Finally, in Chapter Three, I will address a final challenge—that of the source of normativity for coherence requirements. In doing so, I make a constitutivist case for the normativity of coherence requirements. As a consequence of my view, intentional attitudinal incoherence (as a potential action) is not possible. I discuss this point and argue why I believe this consequence bears no problem for the normativity of coherence requirements.

Chapter Two: Challenges for Normative Coherence Requirements

2.0: Chapter Introduction

In this chapter, I defend normative coherence requirements by responding to three questions. These questions will cover various potential issues and challenges that have been raised against the idea that coherence requirements are normative. Now, this chapter will not offer a general argument for the normative source of coherence requirements; arguing this point will come in Chapter Three.[46] That said, I continue to assume that coherence requirements are normative (i.e., that we have a normative reason to comply with these requirements) and argue that this assumption does not lead to conceptual problems. That said, this chapter unfolds as follows.

In section 2.1, I will answer the question: Why care about coherence requirements? Doing so entails discussing why coherence requirements are necessary in an account of rationality as well as how these requirements are supposed to function in our deliberations. I primarily rely on Worsnip's work in this section. In section 2.2, I respond to the question: Do normative coherence requirements lead to bootstrapping and other controversial outcomes? Here I advocate for a wide-scope approach to coherence requirements. In doing so, I dismiss narrow-scope alternatives due to bootstrapping problems and defend the wide-scope approach against the well-known symmetry problem. I also address transmission problems raised against the wide-scope view.

Moving ahead, in section 2.3, I address the question: Are coherence-based reasons right-kind reasons? In this section I discuss Worsnip's position as well as how my agency-constitutivist defense of normative coherence requirements maintains the right-kind reasons distinction. Given this laundry list of challenges, my hope is that this chapter defends the idea that normative coherence requirements are unproblematic. Let us begin.

[46] Chapter Three effectively responds to Kolodny's (2005) challenge that there is no reason to adhere to coherence requirements in general.

2.1: Why Care about Coherence Requirements?

As previously mentioned, for this project I grant the existence of substantive rationality. I grant that we have moral, epistemic, prudential, etc. reasons that, if ignored, entail substantively rational criticism. That said, with these reasons on the table is there anything left for coherence requirements to do?[47]

There is an intuitive sense that, when I act contrary to my best judgment, or fail to intend (what I believe to be) the necessary means for my intended end, something is very wrong with me. In such cases, I am being irrational. This irrationality is, I believe, common ground in most discussions of rationality. What I take to be the above challenge is whether there is a need for *structural requirements* dictating that attitudinal incoherence should be avoided. The question is, then, whether coherence requirements provide a needed function, or whether substantive rationality is enough.

Addressing this issue affords the opportunity to lay out some groundwork for this project. In doing so, I follow Worsnip's lead on many of these points. Rather than offer an original position, the purpose of this section is to motivate the discussion of coherence requirements to begin with. That said, in section 2.1a, I argue that being substantively rational does not guarantee the absence of attitudinal incoherence. Next, in section 2.1b, we see that, thanks to Worsnip's framework, there is conceptual room for considerations of coherence to function in our deliberations. Finally, in section 2.1c, I discuss attitudinal incoherence as an evaluative criticism. If correct, I believe this section offers a strong position in favor of the utility and function of coherence requirements. Let us begin.

2.1a: Substantively Rational Attitudinal Incoherence

It is believed by some[48] that responding correctly to our substantive normative reasons alone guarantees the coherence of our attitudes. If so, then as long as we are not substantively irrational, we are not structurally incoherent. I

[47] This question echoes Kolodny's comment that such requirements are little more than an odd concern for "psychic tidiness" (2007: 241).

[48] See Kolodny (2005) and Lord (2014) among others.

think Worsnip shows this claim to be false. Consider the following permissive cases.

> [S]uppose that one is at a restaurant and that there are three available dishes, each delicious in its own, different way. It's plausible that in some such cases, one's reasons are permissive with respect to one's preferences: no particular pairwise preference would be substantively irrational. From that, it follows that there's no substantive irrationality in preferring dish A to dish B, or in preferring dish B to dish C, or in preferring dish C to dish A. But it's structurally irrational to have all of these preferences together. […]
>
> [S]uppose that given one's evidence, one may permissibly assign some proposition p a credence anywhere between 0.64 and 0.65 -- and therefore, that one may permissibly assign not-p a credence anywhere between 0.35 and 0.36. It follows that it's substantively rational to assign p a credence of 0.64, and substantively rational to assign not-p a credence of 0.35. But it's structurally irrational to have both of these credences. (Worsnip [2018c])

The above examples intuitively show that responding to substantive reasons alone cannot guarantee attitudinal coherence.[49] Where the first example violates the transitive preference requirement, we can understand the second example to violate the modus ponens requirement.[50] As seen in Chapter One, cyclical preferences are incoherent because they are unintelligible under conditions of full transparency. In the same way, violating the modus ponens requirement is equally incoherent. Nevertheless, it seems that, in these cases, responding to

[49] In addition to the examples above, consider Kiesewetter's (2017: 286-287) concession that substantive rationality alone cannot account for a categorical form of the instrumental requirement—i.e., one that does not allow for various exceptions where incoherence is tolerated. In Chapter Three, I will argue for the normativity of coherence requirements, which does not allow for exceptions of "rational" incoherence.

[50] Taking for granted the credence value relationship $\text{cred}(\neg p) = 1 - \text{cred}(p)$, we can say that the credence value of p implies the credence value of $\neg p$. Now, let P: $\{\text{cred}(p) = 0.64\}$ and Q: $\{\text{cred}(\neg p) = 0.36\}$. Then, $P \rightarrow Q$, $P \vdash Q$. So, it is incoherent to believe that p has a credence value of 0.64 and $\neg p$ has a credence value of 0.35.

our substantive permissive reasons alone does not guarantee that we avoid attitudinal incoherence. For substantive rationality evaluates the rational status of attitudes *individually* and not by how these attitudes "fit together."[51] In order to classify such cases as illustrations of irrationality, it seems that we do need structural norms prohibiting attitudinal incoherence.[52]

From the above, I think that cases like these provide a genuine need for coherence requirements.[53] For we need a way of accounting for such cases as instances of irrationality. However, at this point, I think it prudent to discuss briefly how I understand substantive reasons and complying with coherence requirements to work together.

2.1b: Structuring Deliberation with Coherence-Based Reasons

As I see it, coherence requirements play a necessary role in our deliberations by allowing a structural framework for us to deliberate over substantive issues. As prohibitions of attitudinal incoherence, coherence requirements prescribe that, whatever attitudes we have, they ought to be structured in coherent ways. As Worsnip writes,

> [C]onsiderations of coherence, rather than constituting reasons for individual attitudes, constitute reasons to structure deliberation in certain ways: specifically, to treat incoherent combinations as off-

[51] To say otherwise would entail claiming that substantive rationality has its own *structural* requirements. Saying that {$\neg p$ is false} *because* {p is true}, or vice versa, implicitly relies on the *structural* claim that we cannot have { $p \wedge \neg p$}. And if substantive rationality has its own *structural* requirements, then the normative coherence requirements question is moot. For even reason-response theorists would have to commit to some form of coherence requirements.

[52] One may protest that, in these cases, we still fail to respond to our instrumental reason to take the necessary means to our end. As this argument for either case will effectively run the same way, I have chosen to discuss only the first case. In cases of preference reasoning, *the purpose* of considering what dish I prefer over others *is to make a decision about what I will order*. That said, by reasoning in a circle, I fail to respond correctly to my instrumental reason to take the necessary means to decide what I will order. I think this response fails for the following reason: It might be the case that, if we have a *decisive* reason to intend some end, this argument holds water; however, as Kiesewetter argues, if (as in numerous cases) our reason to decide what to order *is not* a decisive reason, then we no longer necessarily have an instrumental reason to take the necessary means to decide what to order (2017: 263-268). I imagine the same problem will hold when discussing the credence case. Because of this issue, I find that substantive rationality still fails to guarantee attitudinal coherence in permissive cases.

[53] For a more in-depth discussion on this point including other example types, see Worsnip (2018b).

> limits, and to focus one's deliberation on adjudicating between the coherent combinations. (21)

There is a nice separation of powers here. Coherence requirements prohibit incoherent attitude combinations, and (of the coherent combinations left available) substantive reasons determine what particular attitudes we ought to have. Worsnip's idea here is a response to Kolodny's challenge that considerations of coherence have no place in our deliberations (2005: 547-8).[54] Given the examples (in 2.1a) where substantive reasons fail to account for attitudinal incoherence, Worsnip makes a strong case for the utility of coherence requirements. Coherence requirements govern acceptable attitude structure; substantive reasons govern acceptable attitudes.

Now, what I have rehearsed about Worsnip's position so far in no way *justifies* the normativity of coherence requirements. As I have said, this justification will come later in Chapter Three, where I repackage Worsnip's position within an agency-constitutivist framework. That said, I continue to use Worsnip's dualist framework throughout this chapter. Moving forward in this section, I now discuss a further question: what kind of irrationality is attitudinal incoherence?

2.1c: Structural Irrationality as Evaluative Criticism

Following Worsnip, we can think of the rational criticism associated with attitudinal incoherence as evaluative—as pointing directly to the presence of attitudes that "don't cohere, or fit together, or "make sense" in combination" (2020: 13). Understood in this way, attitudinal incoherence (structural irrationality) is not a failure to respond to reasons. Instead, this irrationality is a

[54] Kolodny argues that, since we will have substantive reasons for our particular attitudes, having a second reason (that choosing attitude X will make my attitudes jointly coherent) is superfluous. More so, this second coherence-based reason would be the "wrong kind" of reason (i.e., because of some value in attitudinal coherence rather than the substantive reasons that directly bear on the justification of X). Right away, we can see that the first issue (that we would have an unnecessary second reason for X) resolves itself under Worsnip's account. For Worsnip claims that our coherence-based reasons are reasons for coherent attitude *structures*, not for *particular attitudes*. I believe Worsnip succeeds against the second, "wrong-kind" reason challenge as well (Worsnip 2020: 24-28). I further discuss right-kind reasons in Section 2.3.

failure to *structure our attitudes properly*, and in this way is evaluative in nature. Okay, why think this?

First, criticisms using evaluative standards are well-known. Think of evaluative judgments of skill for instance. Dancing, playing music, carpentry work, philosophical research, and many more, are all activities criticizable in evaluative ways. Given that we are understanding coherence requirements to constitute reasons to structure our deliberations in coherent ways, it seems natural to understand structural rationality as an evaluative standard. For complying with coherence requirements equates to meeting a "certain mental performance" (23).

Second, how well I ϕ is a distinctly different judgment from whether or not I ought to ϕ. If I am trying to determine whether I should take up dancing, looking to my substantive reasons seems the natural choice. Once I am dancing, however, my deliberation over how well I am performing seems a very different (structural) matter.

Third, as Worsnip points out, classifying evaluative criticism in terms of failing to respond to reasons is problematic. Considering beauty, Worsnip writes,

> [E]ven if there are reasons to bring about beautiful things, it's not [the case] that what it is to be beautiful is to be responsive to beauty-reasons; and so, it's not that criticism of someone or something as ugly just is a charge of their having failed to respond to beauty-reasons. (13)

I think this point has force to it. For we can meet (and fail to meet) evaluative standards without responding to reasons. Assuming that there is a genuine standard of beauty, it is "possible to be beautiful without having made oneself so in response to a beauty-reason" (11). That said, if structural irrationality is a failure to meet some evaluative standard, then we need not associate this criticism with a failure to respond to reasons. This kind of rational criticism can stand on its own.

Disarming a Challenge to Evaluative Standards

That said, Benjamin Kiesewetter (2017: 33-36) poses a challenge to thinking that rational criticism could be accounted for by an evaluative standard. As I understand Kiesewetter, his argument runs in the following way:

1. If rational criticism is based on an evaluative standard, then this standard must be provided by facts about the well-functioning of our rational capacities (34).
2. Then our failure to meet this evaluative standard is a failure of our rational capacities (35).
3. Therefore, whenever we fail to meet this evaluative standard, we also fail to have the proper functioning rational capacity to be rationally accountable (36).

I think this argument fails.[55] The problem, as I see it, is that there is a conflation between "our failure to meet a rationally evaluative standard" and a "failure of the capacity itself." Following Worsnip's account, the evaluative standard invoked in structurally rational criticism is proper (coherent) deliberation. That said, the standard itself is not our *capacity* to meet that standard. All other things being equal, my failure in a given moment to perform simple mathematical computations does itself entail my incapacity to do so. Similarly, our failure to structure coherently our deliberations in a given moment does not entail that we lack the capacity to do so.[56] We simply made a mistake.[57] In Worsnip's terminology, we failed to be under conditions of full transparency about our attitudes, and because of this fact, we failed to deliberate properly.[58]

[55] To be fair to Kiesewetter, his criticism here is raised specifically against Raz's (2005) argument that rational criticism references an evaluative standard. That said, I leave aside whether my criticism lands on Kiesewetter, Raz, or both.

[56] Not unless our failure is due to being knocked unconscious or killed or remote controlled by a mad scientist/evil demon, etc. However, these cases are, of course, excluded because they violate the failure precondition for normative requirements.

[57] As an example of this, recall Worsnip's preference reasoning cases wherein we deliberate in circles about what food to order (Section 2.1a). With respect to any one of our preference pairs, we deliberate fine. It is just that, overall (across all the preferences), we fail to deliberate properly.

[58] Recall (from Chapter One: Section 1.1) that, under conditions of "full transparency," we hold a combination of mental states without the aid of "self-deception, mental fragmentation, or any failure of self-knowledge" pertaining to our attitudes (Worsnip 2018a: 188). In other words, we

The Consistency of Evaluative Structural Rationality and Reason-Based Normativity

As a final note, we should not worry that a reason-based account of normativity is inconsistent with the above picture of evaluative criticism. The fact that Worsnip describes the normativity of his coherence requirements in terms of reasons does not entail that we need to respond to these reasons directly. The fact that structurally rational criticism now comes apart from any failure to respond to reasons means that we can separate the question of criticism from the question of whether that criticism has genuine normative force. Think of chess (or some other rule-governed activity). When playing chess, if I move my bishop in an "L" shaped way, you can criticize me for doing so. Nevertheless, this fact does not entail that the rules of chess are normative. More to the point: even if chess were normative, the criticism I receive need not be understood as my failure to respond to the reason that a bishop is not a knight, etc. We can simply point directly to my failure to play the game correctly. Returning to attitudinal incoherence, Worsnip writes,

> on a perfectly consistent picture, we can note that charging someone with structural irrationality expresses a kind of evaluative criticism, and then go on to ask—separately—whether the fact that some state would be irrational (and, thus, subject to that kind of evaluative criticism) has any normative significance: whether it is, for example, a reason not to be in that state, or to do something that would result in one's not being in it. (2020: 13-14)

Going forward, I understand the relation between reasons and coherence requirements (if justifiable) to be explanatory of the *normative character* of structurally rational criticism—not the mechanisms that determine this criticism. We are structurally irrational when we fail to structure our attitudes in coherent ways. Structural rationality is normative when we can justifiably say that there are reasons why we ought to avoid attitudinal incoherence. That

---

are fully transparent about our attitudes when we are sufficiently "clear eyed" about how those attitudes relate to one another.

said, I think we now have a decent picture of the view I am defending. It is now time to address further potential criticisms of coherence requirements.

2.2: Do Coherence Requirements Lead to Bootstrapping and other Controversial Outcomes?

Coherence requirements predominantly fall under one of two scope frameworks: narrow and wide. Narrow-scope coherence requirements follow an *if-then* relation, requiring the adoption of specific consequent attitudes based on our antecedent attitudes. Alternatively, the wide-scope approach makes antecedent attitudes alterable, allowing various ways to satisfying coherence requirements. I have opted for the wide-scope approach. That said, I need to say a bit about why. To that end, in (section 2.2a) I discuss why I find the narrow-scope approach to coherence requirements problematic. Afterward, (in section 2.2b) I will discuss a well-known challenge for the wide-scope view and (in section 2.2c) show how a dualist perspective on rationality overcomes this problem. Finally, (in sections 2.2c and 2.2d) I address and meet two further challenges for wide-scope coherence requirements.

2.2a: Narrow-Scope Requirements and Bootstrapping

A strong case against narrow-scope coherence requirements is that they invite bootstrapping problems. Bootstrapping occurs when coherence requirements allow our present attitudes to generate oughts and normative reasons for further attitudes. These cases are problematic. For, as Benjamin Kiesewetter says, "[y]ou cannot, just by adopting some crazy belief about what you ought to do, make it the case that you really ought to do it" (2017: 83). Furthermore, as we will see, the existence of bootstrapped oughts and reasons also leads to incoherent instruction (John Brunero 2020: 38-45). How is bootstrapping possible?

Narrow-scope coherence requirements entail bootstrapping problems because of their *if-then* structure. Narrow-scope coherence requirements take the form: *if certain antecedent attitudes are present, then the adoption of certain*

consequent attitudes is required. Consider narrow-scope versions of the instrumental and practical enkratic requirements.[59]

> *Narrow-Scope Instrumental Requirement* (IR_{NS}): If we intend to ϕ, and we believe that ψ-ing is a necessary means for ϕ-ing, then we ought to intend to ψ.

> *Narrow-Scope Practical Enkratic Requirement* (PER_{NS}): If we judge that we ought to ϕ, then we ought to intend to ϕ.

Both these requirements follow the "if antecedent, then consequent" form. The problem is, of course, that satisfying the antecedents of narrow-scope coherence requirements requires only that we have specific attitudes—*not that having those attitudes is justified*. That said, we can end up with problematic cases where *unjustified* attitudes somehow *justify* the adoption of other attitudes.

Illustrating this point, consider an example by Kiesewetter. Suppose you judge

> that you ought to put all your efforts into building a machine which enables you to travel through time and meet Elvis Presley. So let us suppose that you ought not to put all your efforts into this project, and consequently ought not to intend to do so, but you sincerely believe you ought to do it. (2017: 82)

As Kiesewetter points out, PER_{NS} leads to "the conclusion that you ought to intend to put all your efforts into this crazy project" (82).[60] Furthermore, if, following the PER_{NS} prescription, an intention to build the time machine is

[59] Since my aim is to defend a unified account of all six coherence requirements described in the previous chapter, it is enough to consider potential problems one or two of the requirements in narrow-scope form. For I think a piecemeal defense of a supposed unified account of coherence requirements would be too ad hoc.

[60] While I put the practical enkratic requirement in terms of *judgments*, Kiesewetter words it more commonly in terms of *beliefs* about what we ought to do. Nevertheless, the problematic outcome remains the same.

formed, IR_{NS} then prescribes that the necessary means to build the time machine ought to be intended. In both cases the example assumes that building a time machine is not possible. Thus, any belief that building the time machine is possible, any judgment that we ought to do so, and any intention to do so, is unjustified and irrational.[61]

One potential problem with arguments like the above is that they implicitly rely on both value judgments and intuitions about what is possible. What if we are wrong that time travel is impossible, and what if it could be the case that we ought to go back and talk to Elvis? If so, bootstrapping counterexamples like the above come up short (or at least the jury is out on what these examples show). Along these same lines, we cannot say that coherence requirement instruction leads to normatively unacceptable conclusions without further arguments for why those conclusions are *actually* unacceptable.[62]

Recognizing this issue, Brunero recasts the bootstrapping challenge as the claim that normative narrow-scope coherence requirements produce incoherent advice.[63] A brief summation of one of Brunero's examples should illustrate the challenge:

> Imagine we intend to travel to Chicago by plane and only two airlines are available: United and Delta. We judge that Delta is the better option. Nevertheless, we intend to fly United. Since we judge that Delta is the better option, PER_{NS} tells us we ought to intend to fly Delta. Yet, since we intend to fly United, IR_{NS} tells us

[61] If you believe that building a time machine is possible, then start again with a different premise about something you believe is not possible. The same kind of counterexample will follow.

[62] Imagine that we believe that killing someone is the necessary means for achieving some normatively justified end. The "ought" that the narrow-scope instrumental requirement gives us to intend the means (killing someone) is unacceptable only if it is true that killing is wrong, and that truth relies on *definitively proving* a moral position where killing is an unjust action. While I do believe that killing is morally wrong, the point remains the same.

[63] As Brunero argues, formulating the bootstrapping challenge this way avoids relying "upon controversial intuitions about whether immoral intentions […] generate reasons" (2020: 33).

> that we ought to fly United. So, it seems that we ought to intend to fly Delta and we ought to fly United.[64]

As Brunero rightly points out, the above conclusion "is incoherent advice" (2020: 39). The effectiveness of Brunero's example is that it avoids referencing any potentially problematic moral positions or intuitions about what is possible. Furthermore, Brunero's bootstrapping challenge is stronger than alternative versions because it illustrates a case where coherence requirements literally tell us we ought to be incoherent!

As seen above, narrow-scope coherence requirements run into trouble. In some cases, these requirements somehow justify consequent attitudes with unjustified (or even irrational) antecedent attitudes. In other cases, these requirements actually propose incoherent advice. So, we definitely have a problem.

One response to bootstrapping problems is to say that coherence requirements prescribe subjective "oughts" only. If correct, then these requirements would not bootstrap into existence any other (namely objective) kind of ought.[65] The problem, however, is that this approach unintentionally undermines the normative status of the requirements in question. Here is why.

The problem is that we do not have a good explanation about what subjective oughts normatively amount to. First, if by "subjective oughts" we mean that coherence requirements determine what attitude combinations we *ought* to have *based on our subjective perspective*, then we lose the ability to criticize one another rationally. For having a particular set of attitudes could be rational from our perspective, yet irrational from the perspective of others.[66] In

[64] See Brunero (2020: 39) for the original formulation of this example.

[65] See Way (2009), Schroeder (2009), Parfit (2011: 33–34), among others for versions of the subjective ought strategy for select coherence requirements.

[66] This issue of rational criticism equally challenges an earlier response to bootstrapping objections by Schroeder (2004). This approach, which Schroeder has since abandoned, argues that, instead of "oughts," coherence principles tell us only that we have *a reason* to adopt certain attitudes based on the attitudes we already have. However, these reasons can be outweighed in various cases. As Kiesewetter writes, the reasons strategy "cannot account for the criticism that is commonly associated with irrationality" (2017: 85). Schroeder also concedes to this point (2009: 232).

order to have rational criticism, the "ought" in question needs to reference an objective sense of ought.[67]

Another explanation might be that subjective ought prescriptions are counterfactual claims about objective oughts. If, given our present attitudes, we subjectively ought to adopt further attitudes, then this is the case *when our original attitudes are true or justified*. However, this explanation also seems to undermine the normative status of coherence principles. For, as Stephen Finlay points out, if "'subjectively ought' simply means 'objectively ought if one's beliefs are true,' then we may doubt the normativity of the principle wherever those beliefs are false" (2010: 72). So once again, we run into trouble.

A further issue about this strategy is that it does not actually avoid the bootstrapping problem to begin with. Say the objective ought strategy somehow avoids the lack of rational criticism and generalizability above. Say that subjective oughts can be shown to represent genuine normative requirements. If so, then we have merely reinvented the original bootstrapping problem, albeit now for subjective oughts. For it will be the case, as Kiesewetter points out, that these requirements will "entail that just by adopting an attitude for no reason, we can make it the case that we stand under a genuinely normative requirement to adopt some other attitude" (2017: 87). So, we have not fixed the bootstrapping problem, only recast it.

Because of these issues, I set the narrow-scope approach aside in favor of a wide-scope account of coherence requirements. That said, there remains a well-known challenge against wide-scope accounts. Therefore, in advocating for wide-scope coherence requirements, I must address this challenge.

2.2b: Wide-Scope Requirements and the Symmetry Problem

Wide-scope coherence requirements offer an alternative to the narrow-scope framework. This framework avoids the bootstrapping issue because, rather than require that we be in some consequent state, wide-scope requirements also alternatively allow us to revise our antecedent state (Kiesewetter 2017: 88, 138). Nevertheless, this symmetrical advantage that

[67] By an "objective" ought, I mean an "all-things-considered" ought—an ought that, *in a given context*, cannot be overridden (See Dissertation Program: Section 0.2).

wide-scope requirements possess over narrow-scope versions is considered by many to be a problem in its own right.[68] As Kiesewetter's frames the issue, the

> symmetry of wide-scope requirements […] poses a problem for the wide-scope account, for it seems that there is an asymmetry between different ways of satisfying structural requirements of rationality that wide-scope requirements cannot capture. (138-139).

As I see it, the complaint is this: the symmetrical nature of satisfying wide-scope coherence requirements allows that we can satisfy these requirements in ways that are obviously irrational. And because of this fact, wide-scope requirements fail as rational requirements. Call this the *Symmetry Problem*. That said, the asymmetry lost in the wide-scope approach is the fact that not all means of achieving attitudinal coherence are created equal—there are right and wrong ways of getting there.

To illustrate this point, consider the wide-scope instrumental requirement (IR). Below are two equivalent ways the writing IR—the first as a prohibition (my preference), the second as a disjunctive ought (more common form). That said, IR states that we ought

> (1): to not [intend to ϕ, believe that ψ-ing is a necessary means for ϕ-ing, and fail to intend to ψ].
>
> (2): to [not intend to ϕ, or not believe that ψ-ing is a necessary means for ϕ-ing, or intend to ψ].[69]

Let us focus on the latter iteration of IR for a moment. Because IR is understood as an unconditional requirement, advocates of the symmetry challenge seem to

[68] See Kolodny (2005), Schroeder (2004), Bedke (2009: 687-689), and Finlay (2010: 70-71) among others for different iterations of the challenge I discuss.

[69] Recall that, by De Morgan's Laws, both expressions are equivalent (Chartrand et al. 2013: 112).

claim that the normativity of IR entails that *all means of satisfying* this requirement are *rationally permissible for all cases*.

If this claim is true, then we get problem cases like the following. Let us say that ϕ is {attend a 9:00AM event on time}, ψ is {wake up before 9:00AM}, and our instrumental belief is {waking up before 9:00AM is necessary to attend a 9:00AM event on time}. Now let us say we intend to ϕ. According to IR, we have three options. We can either intend to wake up before the event, cease to intend to go to the event on time, or cease to believe that waking up before the event is necessary for attending the event on time. According to IR, then, we are rationally permitted to drop our instrumental belief while retaining our intention to ϕ. Of course, this choice is irrational and absurd.

Alternatively, consider a moral application of the practical enkratic requirement (PER). In its disjunctive form, PER says that we ought to [not judge that we ought to ϕ or intend to ϕ]. Following the same argument as before, Mark Schroder writes that a

> problem for Wide-Scoping is that it is symmetric. It doesn't distinguish between acting in accordance with your moral beliefs and adopting moral beliefs in accordance with your actions, and as a result it fails to distinguish between following your conscience and the distinctive vice of rationalization. (2009: 227)

Changing our judgment about what we ought to do simply because we lack the intention to follow through seems an irrationally backward means of satisfying PER. And as seen above, such a means allows for immoral rationalizations.[70] Once again, there is a clear sense in which satisfying PER seems problematically symmetrical. For even in non-moral cases, ceasing "to believe that you ought to do something on the basis of your lack of intention to do it […] is not a rational process; it is wishful thinking" (Kiesewetter 2017: 141).

Okay, so this is the symmetry problem for wide-scope accounts of coherence requirements. As stated, this challenge may seem significant.

[70] An equally problematic situations arise in the same way with the doxastic enkratic requirement in cases where we revise our evidential judgments to suit our beliefs.

However, there is a straightforward way of dismissing this worry, which I now discuss.

2.2c: Substantive Reasons Guiding Coherence Requirement Satisfaction

I think the best way to dismiss the symmetry problem is to argue that wide-scope coherence requirements need not, on their own, prescribe *how we ought* to satisfy them—only that we *ought to avoid incoherent attitude combinations*. We should not think that satisfying coherence requirements alone make our actions rationally permissible. It is just the case that, when we satisfy these requirements, we avoid the further rational criticism associated with holding incoherent attitude combinations.

This point brings us back to the different ways we can equivalently write wide-scope coherence requirements. Consider again the instrumental requirement (IR), which states that we ought

> (1): to not [intend to ϕ, and believe that ψ-ing is a necessary means for ϕ-ing, and not intend to ψ].
>
> (2): to [not intend to ϕ, or not believe that ψ-ing is a necessary means for ϕ-ing, or intend to ψ].

That we ought (1) does not necessarily entail that we ought to (or that it is permissible to) do so *by any means in all contexts*.[71] All that is being said is that, in all contexts, there will be *a means* of satisfying (1) that we ought to pursue. Because of this, and that (1) and (2) are equivalent, we can see that a wide-scope account need not commit to saying that satisfying IR is rationally permitted *by*

[71] There are various ways of satisfying any one wide-scope coherence requirement. This fact shows that the total set of means for satisfying any one of these requirements is not a set of necessary means (choosing at least one option is sufficient). Therefore, the claim that we have a reason for all means of satisfying a wide-scope coherence requirement in all contexts seems to rely on a liberal transmission principle. However, as Kiesewetter and Gertken (2020) conclude, liberal transmission is untenable as a normative principle. Instead, we must settle for the more conservative transmission principle, Generic Instrumental Reason (GIR): "If A has a final reason to ϕ, then A has a reason to take means to ϕ-ing" (2020: 7). As Kiesewetter and Gertken point out, however, GIR "allows transmission to a reason for the act-type of taking a means, but not for particular means" (13).

any of the listed means in all contexts—not as long as we can supply the asymmetrical guidance (of choosing the "right" option) by further rational principles.[72] The same point follows for the other wide-scope coherence requirements.

That said, I do not propose that we offer any further structural principles. Instead, following Worsnip, I believe we should allow our substantive reasons to do the work of determining *how* we ought to satisfy coherence requirements—i.e., in what way we should avoid attitudinal incoherence in different contexts.[73] By prescribing to a dualist account of rationality, we have access to both our coherence requirements and substantive reasons. On Worsnip's account of coherence requirements, we should understand these requirements as prohibitions on holding incoherent combinations of attitudes. In this way, "facts about one's existing attitudes, and the incoherence thereof, should guide the way one structures one's deliberation" (2020: 22, fn. 61). When deliberating on some issue, coherence requirements guide in the sense of rationally allowing coherent attitude combinations only. Nevertheless, when it comes down to *deciding which* coherent combination we *ought* to choose, "the only considerations that should then appear are the substantive" ones (22, fn. 61).

---

[72] On this point, wide-scope proponents have offered additional basing principles. For example, Broome writes, we can "account for this asymmetry by means of a further requirement of rationality" (2013: 139-140). In the same vein, Way has also said that the "Wide-Scope view can allow that rationality also requires you to form and sustain your attitudes in the right sort of way, where this is specified by independent basing principles" (2011: 232). Thus, we add further structural requirements either to (1) rule out controversial coherence requirement satisfaction or (2) determine which means of satisfaction are correct. It remains debatable, however, whether accounting for asymmetrical guidance within accounts of structural rationality effectively reduces wide-scope requirements to narrow-scope requirements. At this time, I will not weigh in on this issue. For I think there is a more straightforward solution that solves the symmetry problem. See Way (2011), Kolodny (2005: 520-1), and Kiesewetter (2017: 142-146) for further discussion on this point.

[73] We may think that substantive rationality is doing too much work to justify the need for coherence requirements. I do not think so. First, insofar as we understand coherence requirements as requirements of structural rationality, it makes sense that these requirements govern *structure* and not which particular attitudes we have. Therefore, relegating which attitudes we ought (and ought not) to have in particular contexts up to substantive rationality can be seen as a natural way to go. Second, we have already seen the insufficiency of substantive requirements in permissive cases (Section 2.1a). Third, as Worsnip briefly notes, even without permissive cases, there is something problematic about deliberation that does not rule out in advance incoherent attitude combinations (2020: 30). We should not have to consider incoherent attitude combinations genuinely—their incoherence should be enough or rule them out as a possibility.

Consider a series of bridges stretching across a gorge in various directions. Some of the bridges are out, and as such, the transit authority has marked them with signs. Given these signs, it is obvious that these paths are unadvisable. But as far as deciding which available bridge to take, these signs offer no further guidance. For once we have ruled out the unadvisable routes, which available bridge to take will be decided based on where we ought to go today (not where we should not go).

I think this is an effective way of understanding wide-scope coherence requirements—as signposts prohibiting attitudinal incoherence. These requirements serve their purpose insofar as they rule out incoherent attitude combinations; they need not do more than this. Once we are limited to "structurally coherent" options, our moral, epistemic, prudential, (etc.) reasons will naturally take us the rest of the way to deciding what to do in any given context.[74]

In this way, we need not think of the disjunctive form of wide-scope coherence requirements as universally permissible options in all contexts. Rather, we can interpret these options as the ways in which wide-scope coherence requirements *combine with* substantive requirements in contextually dependent ways. In all cases we ought to avoid attitudinal incoherence. In any particular case, how we ought to satisfy wide-scope coherence requirements will be a question for substantive rationality.

That said, I think Worsnip's prohibitional account of wide-scope coherence requirements can avoid the symmetry problem as discussed above. That said, there is another potential challenge that needs addressing. I turn to this challenge now.

[74] Understood this way, I do not think that wide-scope requirements will necessarily reduce to narrow-scope versions. Kiesewetter, a proponent of the symmetry challenge, writes: "Insofar as basing asymmetries pose a problem for the wide-scope view at all, it must be because of a principled reason that structural requirements of rationality that correspond to the structural irrationality claims have to reflect such basing asymmetries" (143). Follow Worsnip, I have just shown that this is not the case if we allow for both structural and substantive requirements to work together in our deliberations. Understood as such, wide-scope coherence requirements are designed *only* to prohibit attitudinal coherence, unlike their narrow-scope counterparts that prescribe particular attitudes (given antecedent conditions). As long as we relegate asymmetrical guidance to our substantive reasons, we need not think that wide-scope requirements reduce to narrow-scope versions.

2.2d: Coherence Requirements and Liberal Transmission

Kiesewetter has claimed that coherence requirements become normatively problematic when combined with the liberal transmission principle.[75] The liberal transmission principle states that, if we have a definitive reason to intend some end ϕ, and ψ-ing is some means to ϕ, then we also have an instrumental reason to ψ.

> *Liberal Transmission* (*LT*): If we have an intrinsic [non-instrumental] reason to ϕ, and ψ-ing is a means for us to ϕ, then we have a reason to ψ (Kiesewetter 2017: 92).

Kiesewetter uses this principle to undermine coherence requirements by taking ϕ to be the satisfaction of a particular coherence requirement. Granted, on Worsnip's account we do not have a prior reason to satisfy coherence requirements. Nevertheless, Worsnip's position does say that coherence requirements constitute reasons to structure our deliberations in coherent ways (2020: 21). Furthermore, this "structuring" amounts to meeting the satisfaction conditions of these requirements. Thus, it seems that we can run Kiesewetter's argument for taking the means to structure our deliberations in coherent ways. For brevity's sake, then, I will discuss this issue in terms of having a reason to comply with coherence requirements.

Assuming that coherence requirements are normative, we have a reason to follow those requirements. Therefore, if ϕ-ing is satisfying one of these normative requirements, then we have a reason to ϕ. From here, if *LT* is true, we also have a reason to ψ—where ψ is any and all means for satisfying that coherence requirement (Kiesewetter 2017: 92-95). If this argument is sound, then the normative status of any particular coherence requirement ends up justifying *any* (including controversial) means of satisfying that requirement. If Kiesewetter is correct, then the wide-scope view has a serious problem.

[75] In Kiesewetter's defense, I am here responding to older (2017) criticism that I believe Kiesewetter himself would now retract. For Kiesewetter and Gertken (2020) abandon (and argue against) the liberal transmission principle used here. Nevertheless, I offer an independent argument against the current criticism.

Consider my practical enkratic requirement PER. Imagine that we judge that we ought to save a drowning child, but nevertheless we do not intend to follow through. In this case we are violating PER. If PER is normative, then we have a reason to satisfy it. While PER does not explicitly prescribe dropping our judgment that we ought to save the drowning child, this option does technically satisfy PER (for our attitudes would no longer be incoherent in the relevant sense). Therefore, LT says that, based on PER's normative status, we have a reason to drop our judgment that we ought to save the drowning child. However, this way of satisfying PER seems normatively unacceptable.

Nevertheless, hope remains for coherence requirements. Meeting the above challenge, I will now reject the use of LT as seen above. This rejection stems from the fact that LT's normative status is conditional on a restricted set of means for satisfying ends. And this fact, in turn, precludes the use of LT as used above. Let us begin.

First, if LT is true, it is true only when the means ψ for satisfying some end ϕ do not overturn the normative status of ϕ. Otherwise, LT would be literally saying that the ends justify any means in the most extreme and controversial way. I do not think it is a stretch of anyone's imagination to see the normative difference between helping someone and helping that person by killing someone else. The normative acceptance (or lack thereof) of the means plays a crucial role on the normative status of the end. In this way, we cannot divorce the normative status of an end from the kind of means chosen to realize that end. Helping someone and helping someone by killing someone else are normatively distinct ends.

Consider Sue and Karin. Let us assume for the argument that helping people is a normatively justified end (i.e., there is a reason to do so) and killing people is a normatively unjustified act (i.e., there is no reason to do so). Karin is trying to distance herself from a bad relationship, and Sue decides to help. Now, Sue could help by offering Karin a place to stay for a while. Alternatively, Sue could decide to kill Karin's partner, releasing Karin from the relationship.

Nevertheless, I think it is obviously false to claim, just because Sue has a reason to help Karin, that Sue in turn has a reason to kill Karin's partner.[76]

Therefore, if *LT* claims that killing Karin's partner is normatively acceptable, then we should conclude that the trouble lies with *LT* and not with coherence requirements. Alternatively, if *LT* is not claiming that it is fine to kill Karin's partner, then the means referenced by this principle are restricted to avoid such cases. But the only way to guarantee that such cases do not arise is to restrict *LT*'s application to ends that are normatively justified *in reference to a restricted set of acceptable means*.[77]

Therefore, insofar as the satisfaction of a coherence requirement *is* an acceptable (end) input for *LT*, the (means) output from *LT must not be* one that undermines that normative status of the input (the coherence requirement). Since *LT* is only guaranteed to be true when ends and means normatively line up correctly, this principle cannot be used to show that other requirements break this restricted means-end relationship without invalidating its own use in the given context. Therefore, *LT* cannot be used to generate reasons for unacceptable means of satisfying other normative requirements. On these grounds, I reject the use of *LT* in the above challenge against wide-scope coherence requirements.

That said, I think that wide-scope coherence requirements are successful in avoiding scope problems. Assuming the normative status of these requirements leads neither to bootstrapping nor to issues of symmetrical requirement satisfaction. With this class of objections out of the way, it is time to turn to a different kind of challenge.

---

[76] As discussed by Kiesewetter and Gertken, similar challenges against *LT* have surfaced under the heading of the *too many reasons problem* (2020: 4-7). See Broome (2005: 7) and Rippon (2011: 17) for further examples. Because of this problem, and other issues, Kiesewetter has since backed away from *LT*, endorsing instead the more conservative principle *Generic Instrumental Reason* (GIR): "If A has a final [non-instrumental] reason to ϕ, then A has a reason to take means to ϕ-ing" (7). This principle, however, does not claim that we have a reason to take any specific means to satisfy a normatively justified end, only that we have *a reason to act in such a way as* to satisfy that end. As Kiesewetter and Gertken write, GIR "allows transmission to a reason for the act-type of taking a means, but not for particular means." (13). As understood, GIR poses no problem for coherence requirements—since the bootstrapping objection only holds if we generate normative reasons for particular means that, in a given context, are controversial.

[77] See Kiesewetter and Gertken (2020: 4-5) for a similar point.

2.2e: Coherence Requirements and Unalterable Attitudes

Kiren Setiya (2007) claims that cases exist where the only way to satisfy coherence requirements is by adopting a controversial (potentially unjustified) attitude. In turn, these cases can also bring coherence requirements into conflict with one another. If correct, then Setiya presents a significant challenge to all formulations of coherence requirements. The key to Setiya's argument is the claim that sometimes our attitudes are *unalterable* (i.e., an intention, belief, etc. that cannot be changed). If possible, then if these unalterable attitudes are unjustified, we will be required to adopt complementing attitudes in support of the unalterable ones—a big problem for normative coherence requirements.

Introducing Setiya's Unalterable Attitude Challenge

Consider an example given by Setiya involving a supposed unalterable attitude.

> Suppose, then, that I decide to smoke, knowing that I need to buy cigarettes in order to do so. Perhaps it is true that my intention adds a further reason, along with the pleasure of smoking, for me to buy them. But these reasons are not conclusive: in the sense of 'should' which reports what there is most or decisive reason to do, I know that I should not buy a pack of cigarettes. Doing so would be akratic. (2007: 654)

> [T]here is nothing I can do to change my intention to smoke or my belief about the necessary means: these attitudes are not under my control. It follows that the only way in which I can conform to the conditional [if I intend to smoke and believe buying cigarettes is a necessary means to smoking, then I intend to buy cigarettes] is by intending to buy cigarettes. In other words, intending to buy cigarettes is a necessary means to the truth of that conditional. (660)

Above, we have two coherence requirements involved: the practical enkratic requirement (PER) and the instrumental requirement (IR). Recall that violating PER involves intending contrary to our best judgment. So, insofar as we intend

to buy cigarettes, we are violating *PER*. Nevertheless, since we do intend to smoke, not intending to buy cigarettes (what we believe is necessary for smoking) violates *IR* since doing so would be instrumentally incoherent. The correct solution to this problem should be that we drop our intention to smoke.[78] However, Setiya has removed that solution by claiming that the intention to smoke is unalterable. If correct, then Setiya has placed us between a rock and hard place to defend normative coherence requirements.

The way forward, as I see it, is to argue that coherence requirements only govern combinations of *alterable* attitudes. If correct, then I read Setiya's "unalterable attitudes" as things like compulsions or desires—attitudes not governed by rational requirements. There is nothing irrational about conflicting desires. There seems something very irrational about conflicting intentions. So, rather than fall into a possible verbal dispute or metalinguistic negotiation with Setiya over the term "intention," I will argue that the cases Setiya considers violates the normative preconditions for coherence requirements. For I argue that a normative precondition of these requirements is that all the attitudes in question are alterable.

Success and Failure Conditions for Normative Requirements

As a presupposition of normative requirements, there is a common understanding that we must have the ability to succeed and fail with respect to those requirements. Let *S* be a normative requirement. The function of *S* (or any normative requirement) is to guide us toward some behavior *X*[79] (where *X* meets the success conditions of *S*). However, guiding us towards having *X* presupposes the ability to guide us toward having *X*, which in turn presupposes *our* ability to have *X*. If we cannot succeed at having behavior *X*, then *S* cannot guide us toward having *X*. That said, *S*'s ability to function as a

[78] More controversially, if we found out that recent research had conclusively determined that smoking was no longer a health hazard, we could also choose to drop our judgment that we should not smoke. Doing so would also solve Setiya's puzzle. However, for the sake of argument, let us assume that our judgment against smoking is correct, and that dropping it would be irrational.

[79] By calling *X* a behavior, I am being purposefully general to include actions and mental states alike. In this way, my argument is meant for theoretical and practical rational requirements alike.

normative requirement partially relies on our ability to succeed with respect to S.

In a similar way, *S's* function as a normative requirement relies on our ability to fail with respect to S. If we cannot fail with respect to S, then *S's* "guidance" over our behavior is superfluous. First, if we cannot not-X, there is no need to be guided by S to X. And second, if we cannot fail with respect to S, then S-based "evaluations" can never provide new information (for the verdict will always be the same). Simply put, without the ability to fail with respect to S, this requirement reduces to a mere description.

That said, applying normative evaluation requires that the evaluated party meets the success and failure conditions of the corresponding requirement. The party evaluated by S must be able to succeed and fail with respect to S. Otherwise (as just discussed), S cannot function as a normative requirement. With this understanding in place, let us return to Setiya's example.

Dismissing Setiya's Challenge

In Setiya's example, my intention to smoke is supposedly unalterable, making it the case that I either violate the practical enkratic requirement (PER) or violate the instrumental requirement (IR). I violate PER if I intend to buy cigarettes (since I judge that I ought not to buy cigarettes) and I violate IR if I do not intend to buy cigarettes (since I intend to smoke). As said before, the rational solution is to drop my intention to smoke—something Setiya claims is impossible in this case.

Now, while Setiya's challenge is aimed against coherence requirements, this issue involves substantive rationality as well. Why? First, because coherence requirements do not govern the justification of particular attitudes (in this case, my intention to smoke). They instead govern what attitude combinations we ought to avoid. And second, because coherence requirements do not govern *how* they ought to be satisfied, only that they should be satisfied in some way. As previously discussed (in section 2.1b), reasons of substantive rationality are available to determine how to satisfy both PER and IR. For, on

Worsnip's account, our substantive reasons determine *which* coherent attitude combinations we ultimately ought to pursue.

Lastly, let S be a substantive normative requirement that prescribes that people ought not to smoke. S allows us to reference a normative requirement when we say we ought not to buy cigarettes (as it is a means to smoking). Furthermore, I take Setiya's example to rely on some version of S. For Setiya writes, "I know that I should not buy a pack of cigarettes" (2007: 654). Additionally, some version of S must be in play because, if not, then forming the intention to buy cigarettes and dropping my judgment that I ought not to buy cigarettes would be a noncontroversial solution to the situation. So, some version of S is in play.

Now, with the above setup in mind, I reject Setiya's challenge based on the fact that the conditions of his example violate my ability to succeed with respect to the relevant requirements. These requirements are PER, IR, and S. In order for me to succeed rationally in Setiya's example, I need to avoid being practically akratic and instrumentally incoherent. I also need to succeed with respect to S—i.e., I need to both *not intend to buy cigarettes* and *drop my intention to smoke*.[80] However, the latter part of succeeding with respect to S (dropping my intention to smoke) is exactly what the conditions of Setiya's challenge prohibit. That said, since I cannot succeed with respect to S, I cannot succeed with respect to the example's requirement set $\{PER, IR, S\}$.[81] Therefore, since I cannot succeed with respect to the relevant normative requirements, Setiya's example disqualifies itself. For his example violates the success condition for normative requirements. Agreeing on this point, Kiesewetter writes that "the normativity of rationality […] does, I think, require the capacity to modify one's attitudes in the light of reflection" (2017: 100). In order for Setiya's example to be consistent with normative requirements, then the intention to smoke must be

[80] Couching the example in this way is important since I need to avoid being incoherent *in the right way*, which involves responding to the right *substantive reasons*. This point is paramount in order for my answer against Setiya to work in alternative cases where only one coherence requirement is present. Substantive reasons "keep us on track" so-to-speak by satisfying coherence requirements in the right way.

[81] Since satisfying S is integral to satisfying both PER and IR, we cannot divorce my ability to succeed with respect to PER and IR from my ability to succeed with respect to S. See also the previous footnote.

modifiable.[82] That said, Setiya's example fails to undermine coherence requirements.

Dismissing Kiesewetter's Amended "Setiya Challenge"

Nevertheless, Kiesewetter believes an amended version of Setiya's challenge can work against coherence requirements. Here is Kiesewetter's version.

> Suppose that Smith intends to smoke, believes that buying a pack of cigarettes is necessary in order to do so, but does not intend to buy a pack of cigarettes. Smith is capable of revising his intention to smoke as well as forming the intention to buy cigarettes; he has the relevant capacity of self-governance necessary in order to count as violating a reason to be instrumentally rational […]. Next, suppose that Jones is exactly like Smith in every psychological respect. Jones differs from Smith only in that the following counterfactual claim is true about him: were Jones to start revising his intention to smoke, someone (perhaps a hypnotist or a neuroscientist) would prevent this from happening. Since Smith and Jones share the same psychology, and are equally guided by their deliberation, it seems that since Smith is self-governed so is Jones. Hence they are both subject to the instrumental requirement. (101)

I find Kiesewetter's position here odd. Kiesewetter seems to think that the above example captures Setiya's challenge without violating the success and failure conditions necessary for following normative requirements. Kiesewetter claims that Jones and Smith are relevantly equivalent to one another, and because of this fact, if coherence requirements are normative, they apply to both persons. However, Kiesewetter's claim that Smith and Jones are relevantly

[82] Once again, I tie the success condition for normative requirements to the set $\{PER, IR, S\}$ because only S deals directly with modifying particular attitudes (coherence requirements govern attitude combinations) and succeeding with respect to S is necessary to succeed with respect to the set $\{PER, IR, S\}$.

equivalent entails the claim that both Smith and Jones *both* have the ability to succeed and fail with respect to coherence requirements. But this last claim is obviously false.

Smith and Jones are not relevantly equivalent such that coherence requirements apply to both. Smith has the ability to change his antecedent attitudes whereas Jones does not. It makes no relevant difference that, in other cases, Jones might be able to change his antecedent attitudes or that he can "begin to change them" in the counterfactual case. In the counterfactual case, where Jones tries to drop his intention to smoke, outside interference blocks this ability. If Jones cannot *succeed* at changing his intention to smoke, he simply does not have the *ability* to change this attitude. Actuality presupposes possibility, and both Setiya and Kiesewetter have removed Jones's possibility to change his intention in the counterfactual case. So even if both Smith and Jones are relevantly equivalent in all other cases, they are not equivalent in the counterfactual case.[83] Therefore, in the counterfactual case, the relevant rational requirements are inapplicable to Jones. So, once again this kind of challenge misses its mark against coherence requirements. Moving forward.

2.3: Are Coherence-Based Reasons Right-Kind Reasons?

In this section, I will follow Worsnip's argument for why coherence requirements constitute right-kind reasons. Following this overview, I will discuss how my amendments in Chapter Three of Worsnip's account retains the right-kind reason distinction. Let us begin.

2.3a: Right-Kind/Wrong-Kind Reasons—A First Pass

There is a general idea that a "right-kind" reason is one that is the right kind for what it promotes or justifies. For example, having sufficient evidence

[83] There is a further problem for Kiesewetter's version of Setiya's challenge. For Jones should be exempt from rational criticism in the counterfactual case for another reason: *the fact that Jones is trying to make the rational decision to drop his intention to smoke while a third party improperly interferes.* It seems misplaced rationally to criticize Jones (who is trying to make the rational decision to drop his intention to smoke) while a third party is in fact the one(s) violating the coherence requirements (albeit in Jones's head). We should always evaluate people on what they do (or do not do), and in this case, the third party are the ones at fault. Lastly, since the example relies on some version of the normative requirement *S*: {we ought not to smoke}, the third party's action to keep the intention to smoke in play shows them to violate a rational requirement.

that some proposition p is true is commonly taken to be the "right kind" of reason for believing p. For our beliefs should be based on evidence for those beliefs. Similarly, a right-kind reason to intend to ϕ would be one that appropriately (normatively) counts in favor of having the intention to ϕ. In this case, the reason may be moral, prudential, rational, etc.

On the other hand, a "wrong-kind" reason is one that promotes something but does so in the wrong way (failing to justify that something normatively). For example, threats of harm or promises of financial gain would be the wrong kind of reason to believe p. While this kind of reason may be motivating, it fails to justify believing p normatively. For these considerations in no way show the truth of p. Threats of harm and promises of financial gain are just *not the sort of things* that justifies believing p.[84]

2.3b: Kolodny's Challenge

Niko Kolodny has claimed that considerations of attitudinal coherence cannot be right-kind reasons. As Kolodny writes, when a person rationally forms

> the belief that p, […] he does so on the grounds of the evidence he believes there is, not on the grounds of his recognition that, given that he believes that there is conclusive evidence, it would be irrational of him not to believe that p. (2005: 547)

> Thus, if the fact that rationality requires compliance is a reason to comply, it is a reason that we can, and typically do, ignore. In this respect, it would be an odd sort of reason. (548).

As Kolodny sees it, our attitudes (in this case a belief) must be justified directly by the right-kind substantive reasons (in this case, evidential reasons for

[84] Some (including Kolodny [2005] and Kiesewetter [2017: 96-108]) draw the right-kind/wrong-kind distinction as "object-given" and "state-given" respectively. However, see Schroeder (2012) and Worsnip (2020: 25 fn. 66) for reasons against this approach. See also Raz (2009), Heuer (2011), and Hieronymi (2013) for alternative formulations of the distinction. Following Worsnip, I draw ultimately upon Howard's (2019) framework for right-kind and wrong-kind reasons.

believing p). In this way, Kolodny argues that reasons constituted by coherence requirements—reasons to comply with these requirements—could never be right-kind reasons. For the right-kind reasons will always be the substantive reasons that directly justify our particular attitudes. Coherence-based reasons can never directly justify particular attitudes.

2.3c: Worsnip and Right-Kind Reasons for Structural Rationality

A useful aspect of Worsnip's account is that we need not debate Kolodny directly on his criticism—we can simply bypass it. For on Worsnip's account, coherence-based reasons are never reasons for particular attitudes—they are reasons for coherent deliberation (for coherent structure across attitudes). In this way, the original wrong-kind reasons argument against coherence-based reasons fails. For as Worsnip puts the point, "while right-kind reasons for belief may be exhausted by evidential consideration, this doesn't mean that right-kind reasons for coherent deliberation are exhausted by (or even include) evidential considerations" (2020: 25).

Borrowing from Christopher Howard (2019), then, we can understand the distinction between right-kind and wrong-kind reasons as "fit-related"[85] and "value-related"[86] respectively. This distinction has intuitive appeal. It makes sense that evidential reasons will be right-kind (fit-related) reasons for believing some p since they explain why we ought (or ought not) to believe p (based on the evidence). On the other hand, I may find it valuable to intend some ϕ because doing so will lead to large financial rewards. In this case, the promise of financial gain explains why intending to ϕ would be good for me.[87]

[85] As Howard writes, a "way in which a fact can count in favor of an attitude is if it explains why the attitude would be fitting to its object. For example, the fact that Sharon spends a great deal of her time doing charity work is a fact that counts in favor of admiring Sharon, since it's a fact that explains why she's admirable, and so worthy of—or fit for—admiration" (2019: 218).

[86] As Howard writes, a "second way in which a fact can count in favor of an attitude is if it explains why the attitude would be somehow valuable or good. For example, the fact that a deplorable dictator will order your execution unless you admire him is a fact that counts in favor of your admiring the dictator, since it's a fact that explains why your admiring him would be good" (218).

[87] Howard's distinction here resonates with Raz's (2009) distinction between "standard" and "non-standard" reasons. As Heuer comments on Raz's distinction, standard reasons "relate to the inherent standard that governs the forming of the attitude, whereas others [non-standard] are provided by the value of having the attitude" (2011: 177). The idea that right-kind reason references an internal standard strongly supports the idea of reasons that are "fit-related."

With this distinction in mind, Worsnip specifies two readings of Howard's conception of fit-related reasons.

> *Broad Reading if "Fit"*: "to say that it's fitting to ϕ is roughly to say that it's correct to ϕ, given the standards internal to the nature of ϕ-ing" (Worsnip 2020: 25).
>
> *Narrow Reading of "Fit"*: "some response R towards an object O is fitting just if O merits R" (Ibid. 26).

Under both these understandings of "fit," coherence-based reasons *fit* as right-kind reasons. With the broader reading, it seems natural that attitudinal coherence is an internal standard for proper deliberation. This point seems natural considering that, when under conditions of full transparency, cases of attitudinal incoherence become unintelligible (Chapter One: Section 1.2). Further, under the narrow reading, coherence-based reasons also seem to fit. For as Worsnip argues, "incoherent combinations merit being treated as off-limits (in virtue of their incoherence), and that the coherent combinations merit our deliberative attention or focus (in virtue of their coherence)" (26). In line with Worsnip, I think both readings of "fit" favor coherence-based reasons (reasons to structure our deliberations in coherent ways) as right-kind reasons. That said, there is now a possible issue that needs addressing.

2.3d: Are Coherence-Based Reasons Action-Oriented Reasons?

Under the above readings of "fit," if O can be understood as an action, then coherence-based reasons can be interpreted as being reasons for action. On this point, Worsnip acknowledges that understanding structural "deliberation as a 'mental performance'" further lends to this reading (2020: 26 fn. 68). Furthermore, as we will see in Chapter Three, I classify deliberation as a mental action. Because of this fact, my account also encourages the interpretation that coherence-based reasons are reasons for action. Okay, why worry about this?

Truth is an internal standard for beliefs, and so evidence that *p* is true constitutes a fit-related reason for *p*.

Both Ulrike Heuer (2011) and Pamela Hieronymi (2013) argue that reasons for action cannot function under the right-kind reason framework. Nevertheless, I believe this challenge can be met. To do so, I will first lay out Heuer's and Hieronymi's positions on this matter. Following this exposition, I will argue why coherence-based reasons retain the right-kind reasons classification.

2.3e: Reasons for Action Versus the Right-Kind Framework

As I see it, the problem expressed by both Heuer and Hieronymi is that the structure involved in following right-kind reasons is fundamentally different from action-oriented reasons. Because of this fact, we cannot apply the right-kind distinction to reasons for action. Were reasons for action right-kind reasons, Hieronymi writes, "we would say the right kind of reason bears on a question, the settling of which amounts to acting. But settling a question does not amount to acting—it amounts to intending" (2013: 118). In similar stride, Heuer claims that the "standard/non-standard [right-kind/wrong-kind] reasons distinction is really concerned only with reasons for attitudes" (2011: 179). And if coherence-based reasons are reasons for attitudes, then we revive Kolodny's challenge that these reasons (if they are reasons at all) cannot be right-kind reasons.[88] That said, I will now address this "problem."

2.3f: Coherence-Based Reasons as Right-kind Reasons for Proper Deliberation.

The above problem is not really a problem because, under Worsnip's account, coherence-based reasons need not be reasons for action. Instead, they should be understood as reasons for meeting a *deliberative standard*.[89] The point is this: Reasons to structure our deliberations in coherent ways have no bearing on the ends of our deliberations. These reasons are not reasons for what we ought to deliberate about. Rather, in keeping with the evaluative nature of

[88] Others sharing this position include Broome (2013) and Kiesewetter (2017). However, these positions assume that coherence-based reasons must be reasons either for particular attitudes or the particular attitudes options within a disjunctive set. In this way, the concerns these positions express fail to affect Worsnip's position because his coherence-based reasons are not for particular attitudes, but instead for the coherent structuring of deliberation.

[89] On this point, I think Worsnip would be best not to consider coherence-based reasons potentially as reasons for "a special kind of action" (2020: 26). Doing so, I think, is unnecessary and can lead to the problem being discussed.

structure rationality, coherence-based reasons are reasons for deliberating *in the right way*. That said, it is natural to consider coherence-based reasons as right-kind reasons because they directly reference an internal (constitutive) standard of deliberation. Coherence-based reasons are not fitting for *this* deliberation or *that* deliberation—they are *fitting* for *deliberation*, period.

In this way, coherence-based reasons are not reasons for having an intention to deliberate properly. These reasons are reasons for deliberating properly about what we choose (for other reasons) to deliberate about. Thus, deliberations can be mental actions, and coherence requirements can be right-kind reasons for meeting the standard of mental action in general. Understood this way, I think Worsnip's coherence-based reasons retain the fitting nature of right-kind reasons.[90]

2.3g: Right-Kind Reasons in an Agency-Constitutivist Framework

As a bit of foreshadowing, in the coming chapter I argue for repacking Worsnip's position within a constitutivist account of agency. In doing so, I move to justify the normative character of coherence requirements by showing that they function as constitutive norms of self-governance. Given this framework, we may wonder if coherence-based reasons remain right-kind reasons. The short answer is yes.

The slightly longer answer is this: Worsnip's coherence-based reasons remain "fit-related" with respect to the same constitutive standard internal to deliberation. If Worsnip's coherence requirements are constitutive norms of self-governance, then the same internal standard for deliberation (coherent deliberation) is also an internal standard for self-governance. I am *not* arguing that coherent deliberation is a *separate* means used to *achieve* self-governance.

[90] Given that Worsnip's coherence-based reasons are right-kind reasons, I set aside challenges advocating for scenarios wherein we ought, or it is beneficial, to hold incoherent attitude combinations (i.e., eccentric billionaires offering money for attitudinal incoherence or evil demons threaten harm unless we are attitudinally incoherent, etc.). See Brunero (2020: 131-136) and Kavka (1983) for versions of this argument. This challenge relies on value-based reasons, and as such, are classic examples of wrong-kind reasons. Furthermore, the structure of these cases treats attitudinal incoherence as an action (a means toward and end). Because of this fact, as Heuer (2011) and Hieronymi (2013) argue, these reasons (if not wrong-kind reasons) cannot be right-kind reasons. Therefore, to what extent these arguments succeed in showing we have a reason to be attitudinally incoherent, this reason is not one that threatens the normative status of coherence requirements. Taking a page from Kolodny, such reasons "we can, and typically do, ignore" (2005: 548).

We do not deliberate coherently *because* we intend to self-govern. Rather, as agents, we simply self-govern. Coherent deliberation *just is* a constitutive part of this phenomenon.

Furthermore, my position *in no way* relies on *valuing* self-governance.[91] No. Again, my position is that, as agents, we are always already committed to the internal standards of self-governance—it is just what we do. To this end, my position rejects any claim that self-governance is an action[92] or anything we chose to intend.[93] In this way, my position also steers clear of the issue Heuer (2011) and Hieronymi (2013) express above.

In this way, I believe that Worsnip's coherence-based reasons remain right-kind reasons for structuring our deliberations in coherent ways. These reasons reference the same standard and object as intended by Worsnip. Because of this fact, insofar as Worsnip's reasons succeed as right-kind reasons, they remain so under my account.

2.4: Chapter Conclusion

Throughout this chapter, I addressed various challenges leveled against normative coherence requirements. As I have argued, coherence requirements are able to circumvent bootstrapping scenarios while remaining consistent with the intuition that there are right and wrong ways of obtaining attitudinal coherence. Following Worsnip, considerations of coherence occupy a proper place in our deliberations, working in concert with substantive reasons. Finally, I have argued that coherence requirements avoid controversial commitments and can be further classified as right-kind reasons.

Taken together, this chapter aims to offer a sense of security that normative coherence requirements need not entail problematic normative outcomes. With that said, there remains the following question: Why ought we to adhere to coherence requirements? It is this question I now turn to in the following chapter.

---

[91] If so, then coherence-based reasons would be wrong-kind reasons (grounded in our value of self-governance).

[92] I understand self-governance with respect to actions, but this governing state is separate from the action itself. See Chapter Three: Section 3.1.

[93] For questions on how my position deals with normative failure, see Chapter Three: Section 3.3c (Question 2).

Chapter Three: Structural Rationality for Self-Governance

3.0: Chapter Introduction

In the previous chapters, I have relied on a central assumption—that coherence requirements are normative. It is time to back up that claim. In doing so, I am effectively responding to Kolodny's (2005) "Why be rational?" challenge—that no sufficiently normative explanation can be made for complying with coherence requirements of structural rationality in all cases.[94] I disagree. However, following Worsnip (2020), I do not think that this "normative explanation" amounts to having a *prior* reason to follow coherence requirements. For Worsnip argues that coherence requirements are brutely normative.[95] While I am sympathetic to this position, I find that, as it stands, Kolodny's challenge remains a significant problem for Worsnip's account. Therefore, the purpose of this chapter is to offer a way forward that circumvents this problem.

My solution is this: rather than offer a prior reason for why we ought to follow coherence requirements, we should instead recognize that, as agents, we *all already* commit ourselves to following these requirements. I defend this position by arguing that coherence requirements are constitutive norms of self-governance and that, as agents, we are constitutively committed to being self-governing. Like Worsnip's position, I think coherence requirements are brutely normative. Unlike Worsnip's position, I argue that we must understand this brute normativity within the larger normative framework of self-governance.

I believe that framing Worsnip's coherence requirements within a constitutivist account of agency provides three benefits. First, we are able to disarm Kolodny's challenge by showing it to be unintelligible.[96] Second, when rejecting Kolodny's position, we can do so in a way that explains why self-

---

[94] "I have tried to examine those [reasons] that seem initially most promising. The inadequacy of some of these proposals, and the implausibility of others, is at least some further ground, in addition to the bootstrapping problem, to doubt that we have reasons to comply with rational requirements" (Kolodny 2005: 547).

[95] Worsnip's terms these requirements as "non-derivatively" normative in the sense that they constitute reasons to structure our attitudes in coherent ways (2020: 7). In this way, reasons *flow from* these requirements rather than the requirements themselves being justified by a prior reason.

[96] Worsnip (2020) also makes this claim. However, in sections 3.3c-3.3d, I argue why I think Worsnip's strategy is problematic.

governance is brutely normative in a definitive sense.[97] Third, as an entailment of the second point, we can also explain why coherence requirements are normative *for the kind of creatures that we are* (agents).[98]

That said, this chapter unfolds as follows. As preliminaries to my argument, (in section 3.1) I lay out what I mean by self-governance, agency, and action. Next, (in section 3.2) I discuss two ways in which coherence requirements could be normative. In doing so, I raise a dilemma and propose repackaging coherence requirements under an agency-constitutivist framework as a way forward. From here, (in section 3.3) I argue that complying with coherence requirements is constitutive of our self-governance. Finally, (in section 3.4) I argue that we should take our self-governance as primitively (brutely) normative. In doing so, coherence requirements (as a constitutive aspect of our self-governance) are normative as well. Let us begin.

3.1: Self-Governance, Proper Descriptions, and Authored Action

The purpose of this section is to introduce key concepts I will use in the following discussion. I will specify my understanding of self-governance as well as agents and action. In the subsequent sections following these preliminaries, I believe the argument for coherence requirement compliance falls out naturally as a necessary condition for self-governance.

3.1a: Agents as Potential Self-Governors.

Very broadly, I understand agents to be potentially self-governing creatures. Following Connie Rosati, an agent is a creature "capable of engaging in practical and theoretical reasoning and capable of self-governance on the basis of its reasoning" (2016: 194). That said, I understand self-governance to be state wherein, as agents, we choose our actions for reasons and impose some rational structure over our actions in order to govern them as they unfold. That

[97] What I mean is that Worsnip's (2020) coherence-based reasons to structure our deliberations coherently can be understood as definitive, or all-things-considered, reasons. This claim is in opposition to Bratman's position where self-governance-based attitudinal coherence remains *conditional* on our desire to live our lives (2009: 432, fn. 60).

[98] This point refers to criticisms of the kind that Schroeder raises against wide-scope accounts of coherence requirements—that the unconditionality of these requirements leaves *why* these requirements *apply to us* a mystery (2004: 349).

said, I also take self-governance to be directional in the following way: Self-governance is always self-governance *with respect to some action* ϕ. Therefore, insofar as we rationally determine and govern our actions, we are manifesting our self-governance. Given this account of agency and self-governance, I now need to clarify what I mean by actions.

3.1b: Self-Governance and Action

When I embarked on this particular dissertation, I did so because I chose to. Furthermore, building this dissertation is something that I did, not anyone else. The events that collectively add up to this dissertation were dependent on me non-accidentally bringing them into being. Said differently, building this dissertation is my action insofar as I authored it under my own volition.[99]

Along these lines, I very broadly define an *action* as an event (or events) understood under a unified description that is caused and carried out by an agent. That said, an action is *self-governed* when the agent causing and carrying out the action chooses (for some reason) to act and rationally directs the action as it unfolds. In this way, an action is *my action* when I am self-governing with respect to that action—it is something authored by me.

Furthermore, since we commonly attribute actions to an agent by pointing to the agent's corresponding attitudes, I will say that self-governance partially involves the governance of one's action-related attitudes. I govern over my action ϕ (am self-governing with respect to ϕ) insofar as I choose to ϕ and rationally direct my ϕ-related attitudes.

Finally, I understand action broadly, as both physical and mental events. Insofar as my mental events are purposeful—oriented toward an end or goal—

[99] I am granting that agents have the capacity to choose what they do of their own volition. Admittedly, there are entire subfields of philosophy dedicated to the metaphysics of how, or even if, such volition is truly possible. While I agree that such inquiries are undoubtedly important, I believe these discussions take us too far outside the scope of my dissertation. For my work concerns whether we *ought* to follow coherence requirements. However, if one day we learn that "ought" is in no way distinct from "is," that what I ought to do is in fact reducible to what I will do, then there may not be much for ethics and metaethics to say outside of applying physics, neural biochemistry, etc. to our first person perspectives. In this case, I will not need to wonder what I ought to do. For any answer would seem rather meaningless against the inevitability of what, deterministically, I will do. Therefore, for the purposes of my dissertation, I grant agents the ability to choose what they do meaningfully. For any conversation over what reasons or requirements we ought to follow presupposes the ability for us to make a meaningful choice between live options.

these mental events are my actions as well. If I am thinking about X in service of my end E, then my thinking of X is my action. For I am deliberating about X in service to some goal I have chosen. Whether I am purposefully moving my hands or solving a math problem in my head, these are both actions.[100] I am causing, and rationally directing, some event to occur in service of what I have chosen to achieve.[101] In this way, I treat internal (mental) actions on par with external (physical) actions. Under my account, then, actions are those events that agents—when self-governing—purposefully cause and rationally direct. Now, in the discussion so far, I have been implicitly relying on a concept that will become more important as this chapter unfolds. I will discuss this concept now.

3.1c: Proper Descriptions as Constitutive Standards

Say that I am making a cup of coffee. Making this coffee is my action, and I am self-governing with respect to this action as its author. In order to talk this way, I, and others observing me, need to be able to identify similarly what I am doing as making coffee. If I am skydiving, and there is no water within my reach, chances are that I am not making coffee. So there needs to be some external standard of "coffee-making" to point to for reference. More generally put, we need a way of identifying and distinguishing particular actions as they are authored (self-governed) by agents.

In order to do this, I propose using what Jeremy Fix calls "proper descriptions." As Fix writes,

> An exercise of a capacity is by nature subject to a normative principle in that the nature of the capacity determines the content of that principle. [...] An account of the nature of a capacity, then, is just an account of the nature of the activity in question. In my terminology, for each capacity there is a principle that describes

[100] There is a potential criticism here: By allowing certain mental events to be actions, we may look to Heuer's (2011) and Hieronymi's (2013) analyses to argue that the reasons for these (mental) actions cannot be the "right kind" of reasons. I have attended to this potential criticism in Chapter Two: Section 2.3.

[101] In this way, while somewhat descriptively redundant, rationality can be applied to acts of rationality.

> its nature whose content is identical to the *proper description* of a certain activity. And that principle is normative for exercises of that capacity. (2020: 37, italics added).

As I understand it, a proper description is a constitutive standard that allows us to judge *that* and *how well* we are doing something. I take "proper" to mean ideal, that is, a description of the ideal instantiation of some action or activity. Now, in terms of the normative character of these proper descriptions, a few caveats are needed. First, I am not saying that, given the proper description of some ϕ, we ought to ϕ.[102] When we judge how well people dance or build chairs, these judgments are not concerned with the motivations for, and permissibility of, dancing or chair building—the *why* and *should* of it. Rather, we are merely judging *that* and *how well* people perform these actions. That said, the normative character involved in proper descriptions must be conditional on our reasons to perform the corresponding action. If we are ϕ-ing, then we are subject to the standard of ϕ-ing.[103] In turn, we are released from the standard of ϕ-ing when we no longer have a reason to ϕ (or a reason against ϕ-ing). With that said, let us return to idea of self-governance as authorship.

3.1d: Action and Proper Descriptions

We now have a clearer sense of what it means to author an action. I am authoring the action of making the cup of coffee when making that coffee is the object of my self-governance. However, in order for this relation between author and action to hold, I must direct my attitudes to correspond to the events we describe as my action. And in order to do this, the content and structure of my directed attitudes must correspond to the proper description of coffee-making. Otherwise, we cannot say that I am making coffee. That said, self-governing with respect to ϕ—understanding ϕ-ing as *my* action—relies on

[102] This would be absurd as it would entail that we ought to perform every possible action, take part in every possible activity.

[103] There is one exception to this conditionality, one proper description that is not conditional, that I will discuss in Section 3.4b).

the proper description of ϕ-ing as the constitutive standard determining that my attitudes correspond to ϕ-ing.

Taken together, the above offers a brief articulation of the terms and concepts I will use in the following argument. As agents, when we self-govern, we do so in terms of our actions. In turn, these actions—defined under proper descriptions—manifest our self-governance.

3.2: Coherence Requirements and Normative Reason-Relations

The purpose of this section is to explore how coherence requirements could be understood within a reasons-based conception of normative rationality. Under this framework, the general consensus is that any set of requirements *S* is normative only if *S* has a direct relation to the presence of reasons. Following Worsnip, we naturally "cash out" normativity "in terms of the notion of a reason" (2020: 2). Similarly, as Connie Rosati points out, we "tend to understand normativity in terms of reasons" (2016: 205).

That said, my argument proceeds in the following way. Following Worsnip's work (in section 3.2a) I discuss two versions of ordering a relation between normative requirements and reasons. In doing so, I discuss a dilemma I find between the versions. Next, (in section 3.2b) I survey Worsnip's position on this dilemma. Moving forward, (in section 3.2c) I point out a problem for Worsnip's account and (in section 3.2d) discuss possible responses internal to Worsnip's position. Finally, (in section 3.2e) I briefly introduce an alternative way forward for understanding the normativity of coherence requirements. This alternative understanding will comprise the remaining sections of this chapter.

3.2a: Reason-Relations—Derivative and Non-Derivative

Reason-based normativity implies that there exists a relation between normative requirements and reasons.[104] Now, given that I am working to

[104] Here I must acknowledge a limitation of my argument. I will not address particularist accounts of reasons such as Dancy's (2004) position and others. Addressing particularist positions would involve a much larger dissertation, taking us too far afield from the subject at hand. That said, I leave this oversight intentionally as a potential limitation for my view. That said, please see McKeever and Ridge (2006) for a reply to particularism in ethics.

defend wide-scope coherence requirements,[105] my discussion of this reason-relation will be limited to those applicable to unconditional rational requirements. Following Worsnip, this relation orders reasons and requirements in two ways:

> *Derivative Reason-Relation* (DRR): There exists at least one reason R, corresponding to some fact F, for any agent A to comply with S.
>
> *Non-Derivative Reason-Relation* ($NDRR$): Facts about what S requires of some agent A constitutes a reason R for A to have an appropriately related response θ.[106]

Under DRR, reasons are ordered prior to, as justifications of, the normative authority of S.[107] In this way, these prior reasons normatively explain why we ought to follow S. Nevertheless, these prior reasons themselves are left unexplained. Alternatively, $NDRR$ reverses the reason-relation such that S, now non-derivatively (foundationally) normative, constitutes reasons.[108] In this case, S is normatively "explained" by the fact that S constitutes reasons.[109] However,

---

[105] Recall my dismissal of the narrow-scope approach for coherence requirements (Chapter Two: Section 2.2).

[106] See Worsnip (2020: 2-7) for the original formulations. Worsnip does consider a third version of the reason-relation but dismisses it as being too weak to supply a useful account of normative rationality. As Worsnip writes, this approach allows that "our having reasons for these [rational] responses may have absolutely nothing to do with the fact that they are rationally required. This doesn't seem like it would be a vindication of the normativity of rationality in any interesting sense" (6). Given the problems this third option incurs, I have chosen to omit it as a viable alternative.

[107] When we say you ought to follow societal laws because doing so maintains social order, we are giving a prior normative reason R to comply with S (societal laws), where R corresponds to the fact F (following the law maintains social order). So, if we say that there is a reason to follow coherence requirements, then we are understanding coherence standards as derivatively normative. We are looking outside of coherence standards to "find some external goal that they serve" (Worsnip 2020: 4).

[108] See Prichard (1912) for a canonical case or moral requirements being non-derivatively normative.

[109] My argument is running under a framework where reasons are explained by requirements. That said, I will also consider (or "repackage") constitutivist positions on reasons in terms of $NDRR$. By constitutivist positions, I mean any position that explains reasons definitionally in terms of their function. Schroeder's (2005: 16-20, 2007: 212-217) position is one such example. For Schroeder, what it is to be a reason is to be a consideration that explains why ϕ-ing would promote my desires (2007: 217). Now, while the subjective/agent-relative nature of Schroeder's

this approach remains silent on justifying why S constitutes reasons—i.e., this approach rejects the question: Why is S normative?[110] Now, I agree with Worsnip that both orderings of the normative reason-relation meet the criteria for reason-based normativity. Both orderings offer agents reasons that make the content of the requirements guiding.

Nevertheless, there is a problem: DRR requirements alone entail a regress problem, but $NDRR$ requirements leave normativity ultimately unexplained. If S is derivatively normative, its normative status is easily explainable by an external fact F generating a reason R to follow S. However, this approach leads to a regress issue unless S is grounded somewhere by a further $NDRR$ requirement (or constitutive standard).[111] So, it seems that, even if some requirements are derivatively normative, all requirements cannot be. Alternatively, if S is non-derivatively normative, we avoid the regress issue (since there is no talk of prior reasons). However, in this case the normative justification of S is left unexplained[112]—somehow S *just is* normative. So, we either incur a regress or accept that some requirements are brutely normative, leaving the normativity of those requirements unexplained (or at least incompletely explained).

---

reasons are inappropriate for discussing wide-scope rational requirements (see Chapter Two: Section 2.2a), any constitutivist position offering objective/agent-neutral reasons is fair game. That said, in the case of objective reasons, I think the difference (significant to my argument at least) between $NDRR$ requirements and constitutive standards is merely verbal.

[110] See Kolodny (2005, 2007) and Raz (2005) for canonical challenges to normative requirements by posing such questions. My position seeks to give an answer to this question.

[111] Since, under my framework, why F generates R must be explained by a more fundamental normative requirement also needing a prior reason, etc. Again, this point follows because I am classifying constitutivist accounts of reasons under $NDRR$ requirements (See fn. 109 above). When I refer to a DRR requirement, I am following Worsnip (at least as I read him) on the point that this requirement presupposes a further normative structure. In terms of epistemic cases, Worsnip writes, "someone could hold that evidential considerations constitute reasons for belief only in virtue of such a further, general reason (such as the fact that evidential supported beliefs are more likely to be true). But such a person would, ipso facto, not count as thinking that evidence is non-derivatively normatively significant: she would be treating truth, not evidence, as fundamental" (2020: 5). The point here is that, as I am using the idea of DRR requirements, they alone cannot ground normative explanations.

[112] Since there is the outstanding question of *why* facts about S constitute reasons, the "Why is S normative?", or "Why ought we to follow S?", question remains.

3.2b: Worsnip and Non-Derivative Normativity

Now, Worsnip opts for the non-derivative reason relation in order to argue for the normativity of coherence requirements. I understand Worsnip's motivations for this to be twofold. First, Worsnip recognizes that any explanation of normativity must avoid a regress. That said, we must hold some requirement *S* as brutely normative. When faced with the question, "What reason do we have to comply with *S*?", we "should refuse to try to answer that question" (Worsnip 2020: 5). As Worsnip writes, on "pain of regress, eventually you must reach something that you take to be non-derivatively normatively significant" (6).

Second, following H. A. Pritchard's (1912) lead, Worsnip seems to claim that, in order to argue that coherence requirements are "fully" normative, and not parasitic on some other normative structure, we need to keep the normativity of coherence requirements "in house" so-to-speak. Since coherence requirements are understood as requirements of rationality, Worsnip seeks to defend the normativity of these requirements within the domain of rationality (4-8, especially 6). I do think Worsnip is correct that, in arguing for the normativity of coherence requirements, we must avoid a regress of normative requirements and reasons. The buck has to stop somewhere.

3.2c: A Problem for Worsnip's Account

Nevertheless, I find Worsnip's approach problematic. For I do not think we can reject the question, "Why comply with coherence requirements?" in the way Worsnip suggests. Granted, there is something paradoxical in being asked to justify the normativity of coherence requirements with a reason external to rationality.[113] Nevertheless, we should be able to say something about why the normativity of *these particular requirements* should be understood *non-derivatively*. Asked more generally: when choosing between normative requirements S_1, S_2, S_3, etc., why choose S_1 *rather than* S_2 or S_3 to be non-derivative?[114]

[113] There is a nonparadoxical form of this question for Worsnip's account, which I will discuss shortly.

[114] Consider debates in ethics and metaethics: If we are allowed to ignore the "Why?" (or the "Why this rather than that?") question, what do the Humeans, Kantians, and Substantive Moral

These questions foreshadow Rosati's insight that our understanding of normativity remains incomplete if we cannot account for how our normative judgments about our normative requirements hold. Rosati writes,

> I take it to be a desideratum on analyses of normative properties that they make it explicable why these inferential relations would hold among judgments about those properties. We might say that an account of normative properties captures their normativity when it makes the ordinary inferential relations among our normative judgments (our judgments about these properties) explicable. (2016: 209)[115]

It no longer remains reasonable simply to reject the question: Why comply with coherence requirements? For now, as long as we reject this question, we undermine these requirements by making their justification (at least partially) inexplicable.[116] Now, I imagine Worsnip could respond in one of three ways.

3.2d: Possible Responses from Worsnip

First, Worsnip could try to argue that, treating coherence requirements as anything other than non-derivatively normative, makes these requirements problematically weak. However, I think this response would be tenuous at best. For Worsnip himself acknowledges that "views on which rationality is derivatively normatively significant are still views on which rationality is normative in some good sense" (2020: 6). As long as the reasons and oughts associated with coherence requirements are both "all-things-considered" and

---

Realists have left to talk about? As long as each side can reasonably say that their set of requirements *S* gives us reasons, there seems little left to discuss as to *which* approach is the *correct* one. Worsnip acknowledges that different positions take on different requirements as non-derivatively normative (2020: 4-6). Nevertheless, I do not see that he offers a solution for this difficulty.

[115] On this same point, Schroeder argues that a problem for wide-scope accounts of coherence requirements is that *why* these requirements *apply to us* remains unexplained (2004: 349).

[116] Rosati further writes, "Non-naturalism [in ethics] might seem to meet this condition. After all, ethical properties can be analyzed in terms of the ethical property that is unanalyzable and brutely normative, and this would seem to be enough to explain the holding of inferential relations. But this leaves the normativity of the basic ethical property utterly mysterious, and in that respect, the explanation of the holding of these inferential relations is incomplete" (2016: 209-210).

occurring in all rationally governable contexts, these requirements will be strong enough.

Second, Worsnip could respond by arguing that Rosati's requirement of normative explicability is unobtainable. For, in order to meet Rosati's requirement, we would need to answer what Worsnip claims is an "inherently nonsensical or confused" question: What reason do we have to comply with coherence requirements? (5). It is a question that has no informative answer. However, as far as I understand it, Worsnip's account of normative coherence requirements actually allows a sensical form of this question. For Worsnip treats coherence requirements as requirements of *structural rationality*, not rationality in general. It seems then that we can ask the question: what *substantive* reason do we have for complying with coherence requirements? Now, the reason we're looking for lies within the domain of reasons, but outside the domain of structural rationality.

Third, Worsnip could try to claim that the reasons constituted by coherence requirements somehow justify the non-derivative reason-relation. I think this response fails as well. Either this justification is problematically circular (that the reasons coherence requirements give us somehow justify the requirements in the first place), or these reasons point to a further purpose (making coherence requirements derivatively normative). With the latter, I am thinking of Worsnip's claim that coherence requirements give us reasons to structure our deliberations in coherent ways (20-24). Now, I agree with Worsnip that coherence requirements should give us reasons to structure our deliberations in coherent ways. However, these reasons cannot be Worsnip's *explanation* for why we ought to follow coherence requirements—for doing so would seem to make these requirements derivatively normative as a means toward a certain "mental performance" (23).[117]

[117] That said, I do not think it is Worsnip's intent to use these reasons to *justify* coherence requirements under the non-derivative reason relation. Rather, here Worsnip is responding to Kolodny's (2005: 547-8) challenge that there is no essential function in our deliberations for reasons to structure our attitudes coherently. To this end, Worsnip (2020: 16-24) is offering significant progress toward vindicating normative coherence requirements. The point I am making is that this approach would not be successful in meeting Rosati's challenge of normative explicability.

3.2e: A Way Forward

I believe the problem raised is a significant one for Worsnip's defense of normative coherence requirements. Now, I do think Worsnip (2020) is correct that *NDRR* meets the criteria for placing *S* within a reason-based conception of normativity. And Worsnip is also right to be wary that questions of the form, "Why is *S* normative?", invite derivatively normative responses (leading to potential regress problems). Nevertheless, Rosati's point still needs addressing.

As I see it, there are two ways going forward. The first option is to argue for coherence requirements under *DRR*. This approach makes these requirements derivatively normative—i.e., in service of some further non-derivatively normative structure. The second option is to argue for coherence requirements under *NDRR*, like Worsnip does, *but in a way that meets Rosati's challenge of normative explicability*. I opt for this second option. I argue that, once understood as constitutive norms of self-governance, coherence requirements function as part of the larger *non-derivatively normative structure of self-governance*.

That said, my task going forward is twofold. First, I need to show why coherence requirements should be understood as constitutive norms of self-governance. And second, I need to argue that self-governance functions under *NDRR* as a non-derivatively normative structure—one that meets Rosati's challenge of normative explicability. To commence:

3.3: The Necessity of Attitudinal Coherence for Self-Governance

In this section I aim to show that coherence requirements are constitutive norms of self-governance. To do so, I will argue that complying with these requirements is a necessary condition of self-governance. While Bratman (2009) argues for a similar position, I find his argument problematic in various ways. That said, in the spirit of Bratman's project, I offer a new way to show the necessity of attitude coherence for self-governance. To begin, (in section 3.3a) I briefly introduce Bratman's position and discuss a challenge raised by Kiesewetter (2017). Next, (in section 3.3b) I offer an alternative argument for why complying with coherence requirements is a necessary condition for our self-governance. Moving forward, (in section 3.3c) I clarify my position by

responding to three potential questions. Finally, (in section 3.3d) I discuss how my position moves past Kiesewetter's criticism of Bratman. In doing so, I hope to offer a compelling argument for understanding coherence requirements as constitutive norms of self-governance.

3.3a: Attitudinal Coherence as Where We Stand on Issues

Bratman's Account

Bratman argues that, in order to be self-governing, our attitudes must represent where we stand on issues. As Bratman writes, we need a way for "an agent to identify with a certain thought or attitude—of what it is for a thought or attitude to speak for the agent, to be part of where the agent stands" (Bratman 2009: 430). Consider an akratic case. Say we are considering whether or not to join a protest march. We judge we ought to join the protest, but we lack the intention to do so. In this case, Bratman would say that our attitudinal incoherence shows that we remain undecided on where we stand about joining the protest (431). Until we either revise our judgment or follow through with the corresponding intention, we lack "the kind of unity of stance needed for there to be a clear answer" as to how we will guide ourselves (431). In this way, Bratman attests that holding coherent attitude combinations is necessary for self-governance.[118]

A Problem for Bratman's Account

Bratman's argument has an intuitive appeal to it. Since our attitudes correspond to our actions, it seems reasonable that our attitudes should represent where we stand on issues we engage with. Nevertheless, by connecting our *determinations* of where we stand on issues to our *compliance* with coherence requirements, Kiesewetter raises a challenge for Bratman. Kiesewetter argues that a necessary entailment of Bratman's view is that we have a reason to determine where we stand *on all possible issues*—something heavily overdemanding (if not impossible).

[118] Bratman runs this same argument for inconsistent intentions and instrumental incoherence (2009: 431).

The purpose of Bratman's position is to defend the normativity of wide-scope coherence requirements by arguing that we have a reason R that vindicates these requirements. That said, in order to vindicate coherence requirements, R must exist throughout the domain of coherence requirements—i.e., in all possible contexts where structurally rational criticism is possible.[119] Now, here is the rub according to Kiesewetter.

The problem is that Bratman's reason R to comply with coherence requirements transmits from this reason R_0 to determine rationally where we stand on issues.[120] Because of this fact, R_0 stands in an injective relation to R.[121] We have R insofar as we have R_0. Given this transmission relation, if R is to vindicate the normativity of coherence requirements, R_0 must exist at least in all possible contexts within the domain of coherence requirements. And this fact entails that we have a reason R_0 to determine where we stand on all possible issues where structurally rational criticism is possible. Illustrating this problem, Kiesewetter writes,

> Consider the case of withholding belief with respect to a certain issue. Withholding belief with respect to p means that one has not determined where one stands with respect to a certain domain in which rational requirements apply. But this is often the most rational response we can give, and it seems just false to say that we always have a reason to settle the matter whether to believe p or ¬p for any proposition p (consider "the number of stars is equal"). (2017: 106)[122]

[119] This point follows because wide-scope coherence requirements are unconditional requirements.

[120] Ultimately, the reason to comply with coherence requirements traces from self-governance to a need to determine rationally where we stand on issues to following coherence requirements. Nevertheless, the problem is that the transmission of this reason *passes through our rational determinations over where we stand on issues*.

[121] Otherwise, there are cases where Bratman's defense of coherence requirement normativity fails.

[122] Kiesewetter extends his criticism in the following way: "A similar point applies to intentions. It seems clear that we do not have a reason to determine our stance with respect to intending an action for each possible action. But if having determined where one stands in a given domain counts as a prerequisite for governing oneself in that domain, it would follow that a reason to

Not only does it seem unintuitive that we have a reason to determine where we stand on all possible issues; having this reason is seriously overdemanding! With the former, it seems problematic to say that we have a reason to determine where we stand on issues in cases where, due to the available evidence, the most substantively rational course of action is to withhold judgment about a belief or intention. With the latter, determining where we stand on all possible issues is just too demanding (if not impossible) for agents like us. Now, one response might be to claim that our reason to determine where we stand on issues is overridable (or pro tanto).[123] However, if this strategy solves Kiesewetter's issue, it does so by invalidating the normative status of coherence requirements themselves (or at least in the robust sense in which I think these requirements are defendable). Since, if our reason R to follow coherence requirements transmits from our reason R_0 to determine where we stand on all issues, then whenever R_0 is overridable, so is R. Therefore, this strategy entails the existence of cases where holding coherent attitude combinations would not be the most rational thing to do.[124]

I think we must do better; we must show why, in all cases, we are definitively committed (in an all-things-considered way) to following coherence requirements.[125] In order to amend Bratman's argument, then, I have three tasks before me. First, I must show that we can be self-governing in cases where we abstain on some issue (withholding belief, decision, etc.). Second, I must offer

govern oneself in each domain does indeed give rise to a reason to determine for each possible action whether to intend it or not" (2017: 106).

[123] See Broome (2004) for a full discussion on pro tanto reasons.

[124] If so, then in such cases there would be an overriding reason for attitudinal incoherence. See Chapter Two: Section 2.3 (end of section) for my rebuttal. As an additional point, we should recognize that cases where we have reasons to forgo our self-governance (needing to sleep at a certain time, agreeing to be put under anesthesia for medical reasons, etc.) do not entail that the reasons self-governance transmits to coherence requirement compliance are pro tanto. First, because we could say that our reason to forgo self-governance in these cases favors maintaining our self-governance later or in the long run (i.e., resting to work later, medical treatment to carry on afterward, etc.). So, we need not say there is a reason outweighing our self-governance. Second, (following "ought implies can") the presence of self-governance is a necessary condition for normative requirements—i.e., the only cases where reasons will apply are those where self-governance is present. This point applies to cases of suicide as well. For if self-governance necessarily requires attitudinal coherence, *choosing* to commit suicide will also require this same coherence.

[125] Furthermore, even if Bratman's argument could resist Kiesewetter's criticism, Bratman's position makes coherence requirement normativity *conditional* on our desire to be self-governing (2009: 432, fn. 60). Coherence requirements need a stronger normative foundation.

an alternative argument that (1) constitutively links coherence-requirement compliance to self-governance and (2) avoids the issue of abstention. And third, I must show how my position avoids Kiesewetter's criticism.

I will now attend to each issue in turn. The first two tasks I will deal with in section 3.3b. Afterward, in section 3.3c, I will clarify my argument in three ways. Finally, in section 3.3d, I will defend my position against Kieswetter's criticism.

3.3b: My Argument—Attitude Coherence as the Intelligibility of Actions

I agree with Kiesewetter that there will be cases where the most rational thing to do is to withhold belief or judgment. Nevertheless, I do not think these cases preclude our self-governance. Furthermore, I maintain that our self-governance still requires that we structure our deliberations in coherent ways. I will now argue both these points in turn.

Self-Governance without Determining Where We Stand

I think we can reasonably say that it is possible to be self-governing in cases where we abstain with regard to issues. Recall Kiesewetter's example using the proposition about stars. Say I am asked: Is the number of stars even? When asked, I realize that I lack the sufficient evidence to answer the question, and so I decide to withhold my belief on the matter. Even without specifying the minimal conditions of self-governance,[126] I think it is reasonable to say that my determination to withhold my belief in this case can be described as my self-governed action as an agent. I chose to engage with the question, determined the necessary evidence for belief was lacking, and because of this fact decided to withhold my belief on the matter.[127]

Furthermore, similar arguments can be run for any case where our rational engagement with some issue leads to our abstention on that issue. For example, say that I cannot decide whether to have apple or cherry pie. It seems

[126] What the minimal conditions of self-governance are remains an open question. See Buss and Westlund (2018) for an overview. Because of this fact, my argument argues for a necessary condition—that I structure my deliberation in coherent ways.

[127] This case seems fine as an instance of self-governance. For if not, then we are saying that we can only be self-governing when we believe we have "gotten to the bottom of things," when we believe we have definitive answers. I think this is far too strong a condition for being self-governing.

overdemanding to say that my indecision here violates my self-governance with regard to this issue. For, I can simply *decide* not to have pie at all. As long we have a reasoned process between our engagement with some issue and our decision to abstain on that issue, we can describe this reasoned process as our self-governed action.

Now, there is an important caveat here. There will be cases where we must either determine where we stand or abstain in a way that decides for us. With the former, think of cases akin to driving towards a fork in the road. Assuming I do not turn around, I will have to decide whether to go right or left. With the abstention case, recall the protest march example. If I am deliberating on whether or not to attend the march, and I choose to abstain on the issue, in reality it seems that (whether or not I want to frame the situation this way) I have decided not to attend the protest.

In this way, I do not think that finessing Bratman's argument to account for our abstention on issues avoids Kiesewetter's challenge.[128] However, this amendment is a step in the right direction for making Bratman's position more palatable. That said, I propose the following argument for showing the necessary connection between self-governance and complying with coherence requirements.

Self-Governance and Action Intelligibility

Let us start with an uncontroversial premise: If I am self-governing over my action ϕ, it must be at least the case that I am ϕ-ing. Otherwise, I lose the object of my self-governance. So far, so good. However, I now claim that, in order to determine that *I am the one who is* ϕ-ing, it is necessary that *I structure my* ϕ*-related attitudes* in a coherent way.[129]

This point is what I now will argue. That said, as a disclaimer, I want to acknowledge first that my argument runs at a very "course-grained" level, meaning that I intentionally state the argument without attending to every fine-

[128] Restricting the wide-scope coherence requirement domain will be how we ultimately avoid Kiesewetter's conclusion for Bratman's position. To be discussed in section 3.3d.

[129] Doing so removes scenarios where my ϕ-related attitudes are jointly coherent, but I am not self-governing (i.e., my desire or some compulsion takes over, a mad scientist or evil demon remote-controls me, etc.).

grained nuance along the way. I chose this initial presentation for brevity's sake—to get the idea on the table first—in order to discuss the details further. That said, my argument formally runs as follows:

1. I self-govern insofar as I govern my actions.[130]
2. I do not govern any specific action ϕ if I do not choose to ϕ and direct my ϕ-ing.
3. I do not direct my ϕ-ing if I do not direct my ϕ-related attitudes to correspond to a proper description of ϕ-ing.
4. I do not direct my ϕ-related attitudes to correspond to a proper description of ϕ-ing if I do not structure my ϕ-related attitudes in an intelligible way.
5. I do not structure my ϕ-related attitudes in an intelligible way if I do not structure my ϕ-related attitudes in a jointly coherent way.
6. Therefore, (by 1-5) I do not self-govern over my ϕ-ing unless I structure my ϕ-related attitudes in a jointly coherent way.

Being attitudinally coherent in and of itself does not guarantee my self-governance. For example, my attitudes might jointly cohere by accident while I sleep, or a mad scientist may coherently structure my attitudes remotely, etc. Nevertheless, without it being the case that *I am the one* directing this attitudinal coherence with respect to ϕ—I cannot be self-governing with respect to ϕ-ing. In this way, complying with coherence requirements functions as a necessary condition of my self-governance. Given this argument, I will now expand on the steps in more detail.

Action Explanation

Let us start with a common idea. When I act, I do so for a reason, and this reason explains my action. Now, this explanation is important for understanding me as self-governing because it shows that I chose to ϕ. I am ϕ-

[130] With respect to the question of where we stand on some issue X, ϕ can either be (1) an action referencing a determination about X or (2) an action referencing an abstention from X.

ing because of *R*. Nevertheless, *R* cannot function as the entire explanation of my action. For *R* explains only the impetus for my action—the *why* of it. *R* does not explain *how* I ϕ or *that* I am ϕ-ing. Let us focus on the latter—*that* I am ϕ-ing.

When I am self-governing with respect to ϕ, it must be the case that I am (directly or indirectly) causing and directing the events that constitute my ϕ-ing. First, because I cannot govern that which I have no causal influence over. And second, because my self-governance over my ϕ-ing requires that I both chose to ϕ and direct (manage the steps of) my ϕ-ing. If a friend is making coffee, then I cannot be self-governing with respect to making the coffee. The fact is that I am not the one making the coffee. At the same time, if my friend is a mad scientist who directs my body to make coffee, I am still not self-governing with respect to making coffee. The fact is that I am not directing my action of making coffee. In this way, *that I am directing my* ϕ-ing is necessary for my self-governance with respect to ϕ-ing. Now, as far as I am aware, what I have said so far is more or less uncontroversial. Nevertheless, there is another, more controversial aspect of my directing ϕ-ing that now needs addressing.

Attitudinal Coherence and Proper Descriptions of Action

In order to attribute an action ϕ to an agent, it is *necessary that the agent's ϕ-related attitudes are coherent with one another*. Consider an example.

> I know that Tim intends to drink a beer, and I see him heading for the fridge (manifesting his intention to open the fridge). On that basis, I attribute the belief that there is beer in the fridge to Tim. I am assuming that Tim's intentions, desires, and beliefs fit together coherently here. If I didn't think that, I would have no reason to favor attributing the belief that there is beer in the fridge over the belief that the fridge is empty and that the only available beer is in the garage. The latter interpretation literally doesn't make sense of Tim's behavior, namely his heading for the fridge rather than the garage. An assumption of coherence is thus needed to attribute mental states to Tim[.] (Worsnip 2018a: 197-198)

As I read Worsnip, the assumption of relevant coherent attitude combinations is needed to attribute attitudes between agents and their actions. Just as with Tim, without attributing jointly coherent ϕ-related attitudes to me, we cannot say that I am self-governing with respect to ϕ-ing. For we cannot say that I am ϕ-ing at all. Now, what I have just said is not simply a point about interpretation. That we must assume I am attitudinally coherent in order to interpret me as acting is symptomatic of a deeper fact—*that actions come with proper descriptions, and these descriptions have coherent content*.

As discussed earlier (section 3.1c), we identity actions by their proper descriptions. In order to be playing chess, I need to do what is required of playing chess—i.e., trying to win, waiting my turn, moving chess pieces according to the rules of chess, etc. In order for this chess playing to be my action, my ϕ-related attitudes need to correspond to a proper description of chess playing.[131]

However, in order to do this, two things must be true. First, the relevant proper description of chess playing must be intelligible (I assume this is always true[132]). Second, my ϕ-related attitudes must intelligibly correspond to the content of the proper description. For if my ϕ-related attitudes are unintelligible, they will not correspond to the proper description of chess playing. Simply put, I cannot match something that makes no sense with something else that does make sense. I would not even be comparing apples with oranges; rather, I would be comparing "I-do-not-know-what" with oranges. Therefore, in order for my ϕ-related attitudes to correspond to chess making, my ϕ-related attitudes must at least be intelligible.

Self-Governance, Action Intelligibility, and Attitudinal Coherence

Now, recall from Chapter One that, whenever a combination of our attitudes is incoherent, those attitudes are jointly unintelligible (Section 1.1-1.2).

[131] Of course, proper descriptions can become more fine-grained. For instance, there will be as many proper descriptions of chess playing as there are versions of chess. Equally so, if we are teaching someone to play chess, "trying to win" may not be part of the proper description of "teaching chess." The point is that proper descriptions need not become over-demanding as we can account for subtleties with more fine-grained proper descriptions.

[132] Without intelligible proper descriptions, there can be no identifiable actions.

We literally cannot make sense of the attitude combination as anything determinate. This unintelligibility is why, when we are "clear-eyed"[133] about this attitude combination (about its incoherence), we are disposed to drop or change at least one of the attitudes involved (it is our way of making sense of ourselves). That said, in order to avoid my ϕ-related attitudes from being unintelligible, I have to structure these attitudes in a jointly coherent way.

That said, in order to self-govern with respect to my ϕ-ing, I must structure my ϕ-related attitudes in a coherent way. For, in order to direct my action ϕ, I must do so in identifiable ways that constitute ϕ-ing. Otherwise, I am either not ϕ-ing or not directing my ϕ-ing. Therefore, to govern my ϕ-ing, I have to structure my ϕ-related attitudes in a way that corresponds to a proper description of ϕ-ing. And in order to facilitate this correspondence, I must structure my ϕ-related attitudes intelligibly (which is just to say that I must structure these attitudes in a jointly coherent way). I cannot self-govern my action of "going to town" if either (1) I do not structure my attitudes to intelligibly correspond to a proper description of going to town or (2) I am not actually going to town. In order to self-govern my action, then, my action must be coherently directed, and I must be the one directing it. Therefore, my self-governance over my actions requires that I structure my attitudes in coherent ways (that I comply with coherence requirements).[134] In this way, I understand coherence requirements as constitutive norms of self-governance.

3.3c: Three Questions for My Position

Given the above argument, I now want to expand on a few points. For, based on what I have said so far, it seems that my argument could entail an overdemanding account of self-governance and action. I do not think this is the

[133] Following Worsnip's terminology, being "clear-eyed" about our attitudes amounts to being under conditions of full transparency, meaning that we hold a combination of mental states without the aid of "self-deception, mental fragmentation, or any failure of self-knowledge" pertaining to our attitudes (2018a: 188). See also Chapter One: section 1.1.

[134] This point may seem to entail an over-demanding position on action—that all imperfect action is impossible because action requires *perfect* attitudinal coherence. In the following section, I address this worry.

case. Defending this point, then, I now attend to three questions concerning my argument.

1. Does My Position Rule Out Imperfect Action?

From my argument so far, it would seem that I have made action incredibly hard to accomplish. For when I fail to structure my ϕ-related attitudes coherently, it seems that I am not ϕ-ing at all.[135] If so, then it seems that, in order to self-govern over my actions, I need to structure my action-related attitudes in *perfectly* coherent ways. That said, I plan to stave off this interpretation by expanding on my position to allow for imperfect action.

As I see it, I need to offer two things. First, I need a way to account for agents ϕ-ing when those agents are not perfectly ϕ-ing. Second, I need a means to describe mistakes with respect to ϕ-ing. I will address each issue in turn.

Attending to the first issue, I can think we can say that, given an action ϕ, there will have to be some core subset of *essential* ϕ-related attitudes X such that, if we fail to structure X's attitudes coherently, we fail to self-govern over ϕ. In terms of an action, we can refer to the broad strokes of ϕ-ing that, if not carried out, entail that we fail to ϕ. If my action is boarding a bus, there can be both essential and nonessential sub-actions associated with my bus-boarding. For example, intending to and actually boarding the bus are essential; dancing to Gloria Gaynor's "I Will Survive" into the bus is not, but it might get the job done, nonetheless. The point is that there are essential and nonessential sub-actions associated with any complex action.[136]

Given this nested understanding of actions, we can now allow that the intelligibility of our actions runs along a spectrum, that meeting some minimal threshold of essential attitudinal coherence is what is necessary. Said differently, if I am ϕ-ing (or trying to ϕ), it must be the case that some minimally sufficient subset of my essential ϕ-related attitudes jointly cohere in

[135] In order to discuss this issue more straightforwardly, I will here focus exclusively on self-governed action. For if I am forced by some external force to perform some action, we commonly do not attribute this performance (nor the criticism/praise associated with it) to me.

[136] However small a set of attitudes we choose to associate with basic actions, my account will have to claim that attitudinal coherence with respect to these actions will be all or nothing—we either perform these basic actions or we do not.

order to correspond minimally to the proper description of ϕ-ing. Now, where this threshold exactly lies is perhaps hard to specify and may even differ across different adaptations of this view. Nevertheless, needing our attitudes to correspond intelligibly to a proper description of our action does not seem, in and of itself, a hard sell. For we need to be at least intelligibly doing something in order to be governing over that something (there has to be an intelligible object for our self-governance). That said, even without perfectly specifying this threshold, I think that positing its existence is reasonable. Moving to the second issue.

With what I have said so far, I think accounting for action-related mistakes is more or less straightforward: I make a mistake with respect to ϕ-ing when a subset Z of my ϕ-related attitudes is incoherent with respect to the essential subset X of ϕ-related attitudes. Now, I can think of two kinds of cases.[137] The first kind is when Z, being internally coherent, corresponds to a sub-action in conflict with my ϕ-ing. When playing chess, I move my bishop in an "L"-shaped way, accidentally violating the rules of chess.[138] The other kind of case I can think of is when Z is not internally coherent. Say that we are playing speed chess and, intending to move my pawn, I get distracted[139] and fail to take the necessary means (moving the pawn before the timer expires). The main point is that, if we have nonessential ϕ-related attitudes, these attitudes could potentially interfere with how well our essential ϕ-related attitudes function while we ϕ.[140]

---

[137] There can, of course, be other factors leading to imperfect action performance. For example, if, when playing chess, a gust of wind sweeps away the board and pieces, we have imperfectly played chess in the sense that we failed to complete the game. While this failure was not due to a mistake on our parts, we nevertheless failed to succeed at playing the whole game of chess. More importantly, however, such cases would not distract from the intelligibility that we were playing chess. So, these kinds of imperfections I leave aside.

[138] Perhaps I get distracted and mix up the rules in my head, or that I thought my bishop was a knight, etc. The point is that there are ways of describing an intentional action as a mistake. I do not think, however, that this mistake is intentional, which I will discuss soon.

[139] Either by some desire I have or by some other external factor.

[140] Here I would like to note that the above description of imperfect action plays well with Worsnip's (2018a) definitional account of incoherence. If we are genuinely trying to play chess, but our overall chess playing involves some attitudinal incoherence, it makes sense that, when this incoherence is pointed out, we will be disposed to change our attitudes in some way that removes the incoherence. Recall Chapter One: Section 1.1 for my discussion of Worsnip's

Lastly, what about intentional mistakes? At this point, I am comfortable claiming that there are no intentional mistakes. That said, I can think of two ways this question could be raised. The first version asks whether, under my view, intentional deviance from a proper description can be an action. The second version asks whether intentional attitudinal incoherence is possible under my view. I will address each version in turn.

First, if minimally coherent, my action, insofar as it intentionally deviates from the proper description of ϕ-ing, should signify that either I am not really trying to ϕ or I am trying to do multiple things at once and managing poorly. With the former, think of a magician appearing to saw a person in half, when in fact this is an illusion.[141] With the latter, think of my playing chess really badly because I am also talking on the phone, etc. In either case, the intentional deviance can be captured by referencing alternative proper descriptions.[142]

Second, intentional attitudinal incoherence has no place in my account simply because, by definition, this incoherence could not intelligibly correspond to any proper description of an action.[143] For, in order to do so, we would have to claim that the proper description of the action is itself both unintelligible *and* unintelligible in a way that distinguishes it from other proper unintelligible descriptions of actions. At this point, I think we lose the very meaning of description. For we lose the descriptive standards of actions that

account of attitudinal incoherence. Also see Worsnip (2018a: 188-190, 194-198) for the original content regarding this point.

[141] Alternatively, if I intend to move chess pieces around purposefully in illegal ways, then I am not playing chess. I am learning chess or teaching chess (in terms of what not to do) or some other action.

[142] Another option is that various actions are taking place that, when the corresponding attitudes are mistakenly interpreted together, appear very close to a proper description of a different (more complex) action. Either way, the same point follows.

[143] In cases with eccentric billionaires offering us money to be incoherent, the intentional action can still correspond to a set of jointly coherent attitudes. For, in these scenarios, we usually press a button, take a potion, etc. to allow for the incoherence. That said, the action (and corresponding jointly coherent attitudes) can refer to pressing the button, drinking the potion, etc. We need not say that the *effect* (the attitudinal incoherence itself) is what we intentionally do. See Kavka (1983) for the original toxin puzzle example and Gillessen (2018) for a discussion on high-order attitudinal coherence wherein normative coherence requirements account for induced attitudinal incoherence as a means to an end. See also Chapter Two: Section 2.3 (end of section footnote) for why having reasons to proceed in the above cases do not threaten the normativity of coherence requirements.

coherence allows for. Because of this issue, I disregard intentional attitudinal incoherence as an action.

2. Does Our Inability to Hold Intentionally Incoherent Sets of Attitudes Invalidate the Normativity of Coherence Requirements?

We should not think that my denial of intentional attitudinal incoherence removes our ability to fail with respect to coherence requirements, only that we cannot *intentionally fail*. As discussed above, we can clearly make mistakes wherein some combination of our attitudes is jointly incoherent. In these cases, while we do not intentionally violate coherence requirements, we can still fail to meet their standards as I have shown above. We can fail to succeed with respect to standards of some action—either in terms of incoherent attitude combinations or in terms of coherent attitude combinations that are incoherent with one another. In either case, when we fail at some level with respect to an action, we fail to meet the standards of coherence requirements at that same level. In turn, we fail to self-govern with respect to the action at that level. So, we can be more or less coherent across our attitudes as we are more or less self-governing at any point in time. I think this kind of failure is enough.[144]

Furthermore, this limitation of my view does not impede us from immoral cases or possibly even cases of self-deception (if considered an intentional act). With respect to particular actions, we can cohere our attitudes toward selfish acts or hurting others, etc. Similarly, if self-deception is an intentional act, then it seems reasonable that, with respect to that action, we can coherently cherry pick evidence that supports our false beliefs. Furthermore, we need not be structurally irrational in order to be substantively irrational. For we can certainly be attitudinally coherent with respect to ends we choose for the wrong reasons (financial reward, self-conceit, unjustified desire satisfaction, etc.). All these kinds of cases run on my view.

3. Does Self-Governance Require the Absence of all Attitudinal Incoherence?

Given the above description of imperfect action, I can now address a larger potential misreading of my view. For I am *not* saying that, in order to be

[144] For a more in-depth constitutivist defense of this position, see Fix (2020).

self-governing, we have to be *globally* self-governing—that we have to be perfectly coherent across *all* our attitudes. Certainly not. Here, I think sudoku works as a useful metaphor. When playing sudoku, we can have isolated pockets of numbers that jointly cohere within those pockets respectively. Nevertheless, while these pockets *should* cohere with one another across the board, this global coherence need not be the case to play the game (only to win it). What I am getting at is that, while we are probably almost never globally coherent (across all our attitudes), we can certainly have coherent combinations here and there. To continue the metaphor, we are all most likely unfinished sudoku puzzles, self-governing in our respective pockets but not across the board.

Summary

I have argued that compliance with coherence requirements is a necessary condition of self-governance. By structuring combinations of our attitudes in coherent ways, we are able to make sense of our actions as our own.[145] In turn, these (coherently structured) actions are the objects of our self-governance. If this argument is correct, then I have successfully reached Bratman's conclusion. That said, I think it is prudent now to say a bit about why I think my argument does not fall prey to the same problem as Bratman's argument.

[145] David Velleman (2000, 2009) argues similarly, pointing out that actions performed without our awareness can hardly be intentional actions. For Velleman, our actions must be intelligible to us. While Velleman relates this intelligibility of action to a constitutive aim of action based in a sub-agential desire for self-knowledge (2000: 21-26, 139), I make no such particular commitments. For example, Kantians may make similar arguments while relying on pure reason-based motivational models rather than Velleman's desire-based schema (See Uleman [2016] for a conative approach to practical reason). Rather than rely on sub-agential desires or claims of constitutive aims of action, I argue that coherence requirements offer a conceptually light approach to the intelligibility of action. For, rather than pinpoint a special "governing" attitude, coherence requirements structure attitude combinations in ways that instantiate intentionality. Furthermore, my approach allows us to remain neutral on the cognitivist/non-cognitivist debate over intentional action (whether knowledge is or is not involved respectively). For cognitivists like Velleman, the intelligibility requirement I argue for can be incorporated into claims about self-knowledge over our actions. For non-cognitivists like Bratman, intelligibility can be a structural condition on agents identifying with their actions.

3.3d: A Counterfactual Thesis for Coherence Requirements

Like Bratman, I have argued that complying with coherence requirements is necessary for being self-governing. Now, while I am not arguing that our self-governance gives us a prior reason to comply with coherence requirements, I am saying that coherence requirements are constitutive norms of self-governance. That said, when (in section 3.4) I argue for the normativity of self-governance, I will be making the claim that this normativity *is* the normativity of coherence requirements. Therefore, it is reasonable to assume that Kiesewetter's criticism against Bratman can still reach my argument. That said, I will now reexamine Kiesewetter's criticism with respect to my account.

As we have seen (in section 3.3b) there are numerous cases where we abstain on issues while remaining self-governing with respect to those issues (our self-governing decision to abstain on *X*, etc.). These cases put a dent in Kiesewetter's criticism since he can no longer claim that my account entails that we have a reason to determine where we stand on these issues. Nevertheless, a large portion of Kiesewetter's challenge remains. For I still need to deal with all the possible cases where our self-governance does require us to determine where we stand on issues (driving towards a fork in the road, determining whether we will attend an event, moving a bishop, etc.). That said, I now move to meet this challenge.

Self-Governance—One Action at a Time

I think self-governance (if normative) justifies the normativity of coherence requirements one agent and one context at a time. It is the case that there are countless issues where, in order to be self-governing with respect to them, we must determine where we stand. Nevertheless, I do not think that we need to be self-governing with respect to *all these issues at the same time*. There is no need to validate the normativity of coherence requirements for all contexts up front as long as these requirements are validated in all contexts in which we find ourselves. What I am advocating for is the following:

> *Counterfactual Coherence-Based Reasons*: For any possible (mental or physical) action ϕ, *were* an agent to engage with ϕ, *then* that agent would have a coherence-based reason R to structure ϕ-related attitudes in coherent ways.[146]

In other words, (if self-governance is normative, then) coherence requirements will constitute reasons to structure *different* attitudes coherently for *different* agents. Depending on what we are doing or trying to do, coherence requirements will constitute reasons to structure *those attitudes we have*. Collectively, across all possible agents in all possible contexts, I think this claim measures up against the idea that we always have coherence-based reasons to structure our attitudes coherently. As long as we conceive of wide-scope coherence requirements to "kick in" with respect to what we do or think about, we avoid Kiesewetter's claim that we have reasons to determine where we stand on all possible issues. We would only have reasons to determine where we stand on those issues that (1) we engage with and (2) we do not abstain on. This claim, I think, is neither unintuitive nor overdemanding.

A Challenge for My Account and a Domain Restriction for Structural Rationality

Given my position above, one more Kiesewetter-esque challenge remains. Consider the following example.

> *Absent Marshmallow Tower*: If I never have and never will aim to build a 4-foot-tall tower out of marshmallows, how do we explain my reason to comply with the instrumental requirement: [not intend to build the marshmallow tower, or not believe that buying marshmallows is necessary for building the tower, or intend to buy marshmallows]? I guess I do comply with it, since I do not intend to build the tower. But what reason do I have to comply with it? There is no self-governance involved here.

[146] Now, by "engage" I mean that the agent either intends to ϕ, performs ϕ, thinks about (considers) ϕ-ing, etc. I am also assuming here that the agent's potential ϕ-related attitudes are modifiable (See Chapter Two: Section 2.2e).

This worry seems to stem from the fact that the coherence requirements I aim to defend are wide-scope and so, if normative, unconditionally apply for all possible actions. So, what about those actions (or issues) I never have and never will engage with?

I think the proper response here is to restrict the domain of wide-scope coherence requirements to accommodate *Counterfactual Coherence-Based Reasons* above. Since it is not possible for me to self-govern with respect to an action (or issue) I will never engage with, my view disallows the constitution of reasons to structure attitudes relating to these actions and issues *for me* coherently. However, *were* I to engage with building a 4-foot-tall marshmallow tower, *then* I would have the relevant coherence-based reason.

Understood in this way, the "unconditionality" of wide-scope coherence requirements amounts to the claim that, in all cases where agents hold or entertain rationally governable attitudes, coherence requirements *constitute reasons for those agents* to structure those attitudes in coherent ways. Under my account, then, wide-scope coherence requirements are "unconditional" in the sense that they apply regardless of *which rationally governable attitudes an agent has*. Nevertheless, the coherence requirement instruction is still "conditional" in the sense that it relies on the existence of *agents with rationally governable attitudes* in the first place. Okay, so what is my justification for the above position on wide-scope coherence requirements? My answer is twofold.

First, I think that cases unaccounted for by *Counterfactual Coherence-Based Reasons* violate the preconditions for normative coherence requirements. Consider *Absent Marshmallow Tower*. In order for coherence requirements to have authority over how I ought to structure my ϕ-related attitudes, I must be able to succeed or fail at structuring my ϕ-related attitudes in the way coherence requirements prescribe. In the case of marshmallow-tower building, my success amounts to avoiding an incoherent combination of marshmallow-tower attitudes; my failure amounts to having an incoherent combination. However, since "I never have and never will" have any marshmallow attitudes, I cannot fail with respect to the instrumental requirement (or any other coherence requirement) in this case. For I simply never hold an incoherent

combination of marshmallow attitudes.[147] Because of this fact, *Absent Marshmallow Tower* violates the preconditions of normative coherence requirements.

Furthermore, any case that is similar in form to *Absent Marshmallow Tower* will be ruled out for the same reason. If the conditions of an example are that I have and never will have a particular set of attitudes, then I will be unable to fail with respect to coherence requirements about these attitudes. For as long as I can never have these attitudes, I can never have an incoherent combination of them. Of course, if there is a substantive reason to aim to build a 4-foot-tall tower out of marshmallows, then I can be substantively irrational if, other things being equal, I do not respond to this reason. But this is a different issue. Structural rationality is not about *what* attitudes or ends we ought to adopt or drop; it is about *how* we ought to structure the attitudes we have. There is no structuring if there are no attitudes to structure.

Second, even if cases like *Absent Marshmallow Tower* somehow avoid violating the preconditions of normative requirements, we will *never need* a reason to structure attitudes we will never have. Under my account, the reasons constituted by coherence requirements (to structure attitudes coherently) "kick in" with respect to the actions and issues agents engage with. Because of this fact, our reasons to structure our attitudes will always be available for us for any possible action or issue with which we engage.[148] If no possible agent ever considers building a 4-foot marshmallow tower, we do not need a reason to structure these absent attitudes. If some agent does consider this endeavor (or another involving the tower), then that agent will have the relevant coherence-based reason. That said, there exists no *possible (rationally governable) context* where an agent has attitudes $\{x, y, z, \ldots\}$ and there is *no reason for that agent* to structure $\{x, y, z, \ldots\}$ in a coherent way. This conclusion, I think, is enough.

[147] I am, of course, assuming that I do not acquire these attitudes in any other way (by a mad scientist or because I am dealing with a similar issue and these attitudes tangentially arise, etc.). However, in cases where I do have these attitudes, their modifiability is either outside of my control wherein I am not responsible for them (Chapter Two: Section 2.2e) or within my control, and I do have a reason to structure them coherently. In either case, however, the presence of these attitudes takes us outside the example currently discussed.

[148] Since (in Section 3.1b) I include deliberations as mental actions alongside physical actions, any possible agent will have a reason to structure their attitudes coherently for all rationally governable attitudes they will ever have.

Moving Forward

I believe that the above discussion suffices to show that we can constitutively connect attitudinal coherence with self-governance in way that avoids the pitfalls of Bratman's argument. In doing so, I claim that coherence requirements are constitutive norms of self-governance. For, the intelligibility of what we are doing when we self-govern our actions necessarily depends on our action-related attitudinal coherence. Therefore, our capacity to structure our attitudes in coherent ways is part and parcel of our capacity to self-govern. With that said, it is finally time to argue for the normativity of coherence requirements by arguing for the normativity of self-governance.

3.4: The Non-Derivative Normativity of Self-Governance

In this section I aim to show that our self-governance functions as a non-derivatively normative structure. In order to do this, (section 3.4a) I argue that self-governance functions naturally under the non-derivative reason-relation. Next, (in section 3.4b) I show that, under the non-derivative reason-relation, self-governance meets Rosati's condition of normative explicability. That is, I argue that the normativity of self-governance is self-justifying. If successful, then the normativity of our self-governance will vindicate the normativity of coherence requirements.

3.4a: Self-Governance under the Non-Derivative Reason Relation

Recall (from section 3.3a) that a set of requirements S is non-derivatively normative if facts about what S requires of agents constitute reasons for agents to have an appropriately related response θ. Now, let us plug in self-governance for S and see if the relation makes sense.

> *Non-Derivative Reason-Relation* (for Self-Governance): Any fact about what self-governance requires of some agent A constitutes a reason R for A to have an appropriately related response θ.

Does self-governance make sense under the non-derivative reason-relation? I think it does. First, our self-governance constitutively requires of us that we are

able to respond to considerations and facts as potential reasons. This means that, when I am self-governing with respect to some action ϕ, the facts about what self-governance requires of me (having the ability to respond to ϕ-related reasons in this case) constitute[149] reasons for me to ϕ. This point is more or less definitional since the principle of "ought implies can" itself is built into the very concept of self-governance.[150] Second (and this point is really more of an extension of the first point), my appropriately related response θ will just be my (mental or physical) ϕ-ing. Since the function of self-governance transparently manifests as our particular actions, θ will always be my ϕ-ing based on the reasons to ϕ. As such, I think self-governance easily fits within the non-derivative reason-relation.

3.4b: The Explicability of Non-Derivatively Normative Self-Governance

Now, let us return to Rosati's point about normative explicability. Following Rosati, "an account of normative properties captures their normativity when it makes the ordinary inferential relations among our normative judgments (our judgments about these properties) explicable" (2016: 209). Now, in cases of derivative normativity, this explicability is straightforward. We ought to follow requirement S because of some reason R.[151] However, as we have already seen (in sections 3.2a-3.2d), when it comes to non-derivatively normative structures, dealing with this question is no longer easy or straightforward. Therefore, in order to vindicate the normativity of our self-governance, I need to show how self-governance allows us to respond properly to the question: Why be self-governing? Furthermore, I need to respond to this

[149] Here, I am using "constitute" in a liberal sense. What I mean by "constitute" is that my capacity for self-governance constitutes ϕ-related reasons as *my* ϕ-related reasons. The idea here is similar to mental representations "taking in" what is external and constituting it internally. This point conforms with an externalist view on reasons. However, from an internalist perspective on reasons, we could alternatively take "constitutes" more literally, that, based on external facts, our self-governance constitutes reasons internally "from the ground up," where otherwise these reasons would not exist. The point is to stay neutral here on the reasons-externalism/internalism debate. For I do not think that a victory for either side in this debate will seriously affect my argument here.

[150] What I mean is that we typically understand our capacity for self-governance in terms of our ability (the "can" part) to respond to reasons (the "ought" part).

[151] Of course, following Worsnip (2020), this strategy is successful only if whatever R transmits from eventually bottoms out in some non-derivatively normative structure.

question in a way that is both informative and final—in a way that justifiably terminates the regress of further "Why?" questions.

As it turns out, I think Worsnip hits on this point when he references Christine Korsgaard's (1996) work. Worsnip writes, "following Korsgaard, let's call anything that brings the regress to an end—whatever is non-derivatively normatively significant—a 'source of normativity'" (2020: 6).[152] Korsgaard's insight here is that the only satisfactory answer that "will bring the reiteration of 'but why must I do that?' to an end […] [is] one that makes it impossible, unnecessary, or incoherent to ask why again" (1996: 33). As I have said, Worsnip is correct that the regress of reasons and requirements does have to stop somewhere. My point is that I think Worsnip stops in the wrong place.

For, as I have said previously (in sections 3.2c-3.2d), the problem for Worsnip's account is that it does not seem "impossible, unnecessary, or incoherent" to ask: Why structure our attitudes coherently? First, this question does not seem incoherent or impossible to answer. If we treat substantive and structural rationality distinctly (which Worsnip and I do), we can coherently ask if there is a substantive reason to follow structural forms of rationality. Second, the above question seems very necessary. For, as Rosati's aptly points out, in order to capture a requirement's normativity, we need to be able to explain why that requirement is normative (2016: 209).

That said, I think the only available place to stop the regress is with self-governance. Of course, stopping the regress with self-governance *is* stopping the regress with coherence requirements—since I have already argued that the latter are constitutive norms of the former. However, as we will see, my repackaging of Worsnip's coherence requirements into self-governance allows for a more robust defense in terms of self-governance. This is where we need to make our stand. Why? Because only with self-governance can we both reject further "Why?" questions *and* respond to the question "Why be self-governing?" in an informative way.[153] How is this possible? The answer lies in the relation between self-governance and agency.

---

[152] Here, Worsnip is referring to Korsgaard's 1996 monograph, *The Sources of Normativity*.

[153] Bottoming out the normativity of coherence requirements in other structures risks conceiving of attitudinal coherence as *valuable* for achieving other ends (effective desire-satisfaction, etc.). Doing so risks making coherence-based reasons wrong kind reasons (see Section 2.3). Furthermore, approaches like Korsgaard's (1996) project rely on recognizing the "value" of self-

Self-Governance as the Proper Description of Agency

The key to understanding the normative force of self-governance is realizing that this rationally self-guiding state *is* the *proper description of agency*. Recall (in section 3.1c) that a proper description is a constitutive standard that is normative for the exercise of what that description describes. Thus, understanding agents as potential self-governors opens the door for understanding self-governance as the proper description of agency. With this proper description in place, we can now both identify agents *as agents* and make judgments about how *well we are functioning* as agents.

In this way, self-governance is constitutively normative for *how to function* as an agent. Unlike proper descriptions of actions, however, I think that self-governance is an *unconditionally normative activity*. Why? Because insofar as we operate in the domain of reasons—insofar as normative requirements apply to us—we *are functioning according to the constitutive standard given by the proper description of agency* (i.e., self-governance). Our capacity to respond to reasons, to structure and guide our actions rationally, is part and parcel of being an agent.

Two Perspectival Forms of the Same Question

Returning now to the "Why?" question, how can I both reject further "Why?" questions *and* informatively respond to the question Why be self-governing? Because, when asked about self-governance, the "Why?" question implies two possible vantage points: *externally* by "something" considering whether to be an agent (and adopt the norms of self-governance), and *internally* by an agent (already functioning under those norms). That said, I reject the external version and will show why the internal version is self-validating.

The external version of the "Why be self-governing?" question is literally unintelligible. This is so because, outside of our capacity for self-governance, there are no standards determining (1) what is being asked and (2) what merits an acceptable answer. You cannot coherently ask a question outside the space

---

governance (agency in Korsgaard's terms). I think these kinds of positions argue for self-governance "too far down the road" so-to-speak. For the introspective regress that Korsgaard relies on *is itself* an act of self-governance. That said, we must already be committed to the norms of self-governance to determine its value for us. That said, I think we can argue more directly by looking at the structure of self-governance itself.

of reasons. Since, no matter what account of reasons is on the table, reasons to be self-governing will have to be considered within the agential domain of self-governance.[154] In this way, I reject the external version of the "Why?" question along with any further external "Why?" questions. I have ended the regress from this perspective.

However, I have done so in an informative way. I have rejected further "Why?" questions based on the fact that there is no intelligible prior vantage point from which we can take up these further questions. So, while *I do not offer an external prior reason* for being self-governing, I do informatively say *why* the regress *ought* to stop here (rather than alternatives like desire-satisfaction, brute rational normativity, universal values, etc.).[155]

The internal version of the "Why be self-governing?" question, alternatively, is self-validating. Taking a page from Matthew Silverstein, the "normative authority of practical reasons is not predicated on your having good reasons to remain an agent. Rather, it is predicated on the fact that you are already an agent" (2014: 1135).[156] The point here is that, insofar as we are able to consider the "Why?" question, we are already "playing by the self-governing rules of agency" so-to-speak. In this way, any consideration of the "Why be self-governing?" question relies on our compliance with the very norms about which this question asks.[157] In this way, we cannot choose not to validate the norms of our self-governance.[158] Thus, as self-governing agents, when we are

---

[154] Remember that, on my account, our deliberation of an issue is itself an act of self-governance. That said, we must entertain this question from within a self-governing perspective.

[155] I find Worsnip to be making a similar point, only about coherence requirements themselves: "there are some particular questions of the form 'why be *X*?' that are inherently nonsensical or confused" (2020: 5). As I have said before, however, the problem is that I think the unintelligibility of such questions does not occur until you try to escape the entire space of reasons—until you try to escape agency itself.

[156] Here Silverstein (2014) is responding to Enoch's (2006, 2011) Shmagency challenge that asks constitutivists, Why be an agent?

[157] Of course, various reasons for suspending agency exist, including taking a nap or avoiding pain (rendering oneself unconscious) (Silverstein 2014: 1134). Equally, given overwhelming situations, we may decide to take breaks from thinking and decision making to optimize our performance overall. Regardless of particulars, the normative force of reasons arrived at within self-governing agency derives from the proper description of agency (the activity of self-governance).

[158] Once again, since normative error is understood as something that happens to us or about us (as a defect of our action) rather than something we intentionally do (as a perfection of our action), the inability to choose against this validation is not a problem for my position (see section 3.3c).

faced with the internal "Why be self-governing?" question, we should not reject it. Rather, it is prudent for us to answer: "Because being self-governing allows us to lead our lives."

Non-Derivative Normative Explicability

I think this answer should be sufficient for Rosati's requirement of normative explicability. First, because this approach allows an informative response for why self-governance is an appropriate normative structure for ending the regress of "Why?" questions. Second, because this response (to the internal version of the "Why?" question) allows us to *illustrate* the necessary normativity of self-governance without relying on *a reason* to for this normative authority. Third, I think this response captures a key aspect of normativity that Rosati was driving at: that "[n]ormative judgments can carry recommending force only to [...] beings capable of self-governance" (2016: 207), and this is why "a property can have a normative character only if it 'essentially or analytically involves reference to an agent'" (2016: 209).[159] To reiterate, my answer is neither a rejection of the (internal) question asked, nor a further external reason to be self-governing. It is an expression of the constitutively normative framework under which we are already functioning. Now, this answer maybe seem circular, and in fact to a degree it is. Nevertheless, I think this answer is not viciously circular—for its circularity offers a clear explanation of our commitment to self-governance. First, because self-governance is a necessary condition for all normative discussions. And second, because our commitment to self-governance is simply what is functionally constitutive of the kind of creatures that we are. Apple trees apple; agents self-govern.

[159] This point solves a further issue raised by Schroeder (2004) against wide-scope formulations of coherence requirements. As Schroeder writes, wide-scope accounts posit "a basic, eternal, agent-neutral requirement rationally binding on every agent, no matter what they are like. As with all of the others, *it offers no explanation of this requirement*. Narrow-Scope accounts can explain the obligations or reasons that they postulate. After all, these obligations or reasons only exist given a certain condition—so we can use that condition to explain them. But not so for the Wide-Scopers" (2004: 349, italics added). With my account, we can now explain the obligations of wide-scope coherence requirements: Meeting these obligations is part and parcel of properly functioning as an agent (*being self-governing*).

3.5: Chapter Conclusion—Structural Rationality for Self-Governance

As I have argued, coherence requirements are constitutive norms of our self-governance. For the intelligibility of our governed actions depend on the coherence of our action-corresponding attitudes. Furthermore, I have claimed that self-governance is just what it means to be a properly functioning agent—to correspond to the proper description (the constitutive standard) of agency. Defending this claim, I have argued that self-governance holds under Worsnip's non-derivative reason-relation for normative requirements. Furthermore, I have argued that self-governance (unlike coherence requirements) allows us to meet full normative explicability—to explain *why* this requirement is non-derivatively normative *for the kind of creatures that we are*. Taking self-governance as the normative standard for agency, insofar as we remain agents, we commit ourselves to this standard—we continue to "play the game of agency," which requires our adherence to its norms. And as I have argued, coherence requirements are among these norms. Therefore, our commitment to the normativity of self-governance *is* our commitment to the normativity of coherence requirements.

In this way, I claim to vindicate Worsnip's account of coherence requirements against Kolodny's (2005) "Why be rational?" challenge. Kolodny asks us for a substantive, and unconditional, reason to follow coherence requirements (2005: 542-543). While I have not offered a substantive reason, I have shown that we are unconditionally committed to following coherence requirements. Furthermore, under this agent-constitutivist framework, we can now see the unintelligibility of Kolodny's "Why?" question as it takes the form: "Why be an agent?"[160] For, in all possible contexts we engage ourselves in as agents, we in turn exercise our self-governance, committing to coherence requirements as constitutive norms.

[160] In this form, Kolodny's question becomes the same as Enoch's question: "Why be an agent, rather than a shmagent? (2006, 2011). Aimed uniquely at the issue of coherence requirement normativity, my argument against Kolodny is an application of Silverstein's (2014) response against Enoch.

Dissertation Conclusion

So, where have we come to? I have aimed to show that wide-scope coherence requirements are genuinely normative requirements of structural rationality. Working from Worsnip's framework, I believe a strong case has been made. More specifically, I have worked to amend Worsnip's (2020) position on normative wide-scope coherence requirements in several ways. Beginning in Chapter One, I laid out six coherence requirements under Worsnip's dispositional framework. In order to provide a better unification of Worsnip's dispositional account of incoherence, I aimed to provide a stricter (more attitudinally incoherent) conception of practical akrasia. In Chapter Two, I explored Worsnip's position further and argued against various challenge in order to sustain the idea that normative coherence requirements do not entail controversial conclusions.

Reaching Chapter Three, I worked to repackage Worsnip's coherence requirements under a constitutivist framework about agency. Pulling inspiration from Bratman's work, I reconceived coherence requirements as constitutive norms of self-governance. That said, I believe now that Worsnip's position can effectively meet Kolodny's "Why Be Rational?" challenge. For all agents, insofar as they agents, are always already committed to the norms of self-governance. This constitutivist fact allows us to ground coherence requirements in a definitive way while also illustrating *why these requirements are normative for us*.

References:

Arpaly, N. (2000) On Acting Rationally Against One's Better Judgment. *Ethics*, 110(3): 488–513.

Baker, D. (2018). Skepticism about Ought *Simpliciter*. In R. Shafer-Landau (Ed.). *Oxford Studies in Metaethics* (Vol. 13). 230-252. Oxford UP.

Bedke, M. (2009). The Iffiest Oughts: A Guise of Reasons Account of End-Given Conditionals. *Ethics*, 119 4): 672–98.

Bratman, M. (1979). Practical Reasoning and Weakness of the Will. *Noûs*, 13(2): 153-171.

—. (2009). Intention, Practical Rationality, and Self-Governance. *Ethics*, 119(3): 411-443.

Brunero, J. (2020). *Instrumental Rationality: The Normativity of Means-End Coherence*. Oxford UP.

Broome, J. (2004). Reasons. In R. Wallace, et al. (Eds.). *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*. 28-55. Oxford UP.

—. (2005). Have We Reason to Do as Rationality Requires?—A Comment on Raz, *Journal of Ethics and Social Philosophy*, 1(1): 1-10.

—. (2013*). Rationality Through Reasoning*. Wiley-Blackwell.

Buss, S. and Westlund, A. (2018). Personal Autonomy. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. https://plato.stanford.edu/archives/spr2018/entries/personal-autonomy/.

Chartrand, G., Polimeni, A., and Zhang, P. (2013). *Mathematical Proofs: A Transition to Advanced Mathematics* (3rd Edition). Pearson.

Dancy, J. (2004). *Ethics Without Principles*. Oxford UP.

Egan, A. (2008). Seeing and Believing: Perception, Belief Formation and the Divided Mind, *Philosophical Studies*, 140(1): 47–63.

Enoch, D. (2006). Agency, Shmagency: Why Normativity Won't Come from What Is Constitutive of Action. *The Philosophical Review*, 115(2): 169-198.

—. (2011). "Shmagency Revisited." In M. Brady (Ed.), *New Waves in Metaethics*. 208-233. Palgrave MacMillan.

Fogal, D. and Worsnip, A. (forthcoming). Which reasons? Which rationality? *Ergo: An Open Access Journal of Philosophy*. Accessed From: http://www.alexworsnip.com/research/#publications.

Finlay, Stephen. (2010). What Ought Probably Means, and Why You Can't Detach It. *Synthese*, 177(1): 67–89.

Fix, J. (2020). The Error Condition. *Canadian Journal of Philosophy*, 50(1): 34–48.

Gillessen, J. (2018). Rationality, Virtue and Higher-Order Coherence. *Dialectica,* 72(3): 411–436.

Hare, R. (1952). *The Language of Morals*. Oxford UP.

Hieronymi, P. (2013). The Use of Reasons in Thought (and the Use of Earmarks in Arguments). *Ethics*, 124(1): 114-127.

Howard, C. (2019). The Fundamentality of Fit. In R. Shafer-Landau (Ed.). *Oxford Studies in Metaethics* (Vol. 14). 216-236. Oxford UP.

Kavka, G. S. (1983). The Toxin Puzzle. *Analysis*, 43(1): 33-36.

Kiesewetter, B. (2017). *The Normativity of Rationality*. Oxford UP.

Kiesewetter, B. and Gertken, J. (2020). How Do Reasons Transmit to Non-Necessary Means? *Australasian Journal of Philosophy*: 1-15. DOI: 10.1080/00048402.2020.1745252

Kolodny, N. (2005). Why be Rational? *Mind*, 114/455: 509–563.

—. (2007). How Does Coherence Matter? *Proceedings of the Aristotelian Society*, 107(1.3): 229–263.

Korsgaard, C. (1996). *The Sources of Normativity*. Cambridge UP.

Lord, E. (2014). 'The Coherent and the Rational,' *Analytic Philosophy*, 55(2): 151-175.

McKeever, S, and Ridge, M. (2006). *Principled Ethics: Generalism as a Regulative Ideal*. Oxford UP.

Parfit, D. (2011). *On What Matters* (Vol. 1). Oxford UP.

Pritchard, H. (1912). Does Moral Philosophy Rest on a Mistake? *Mind*, 21(81): 21-37.

Raz, J. (2005). The Myth of Instrumental Rationality. *Journal of Ethics and Social Philosophy*, 1(1): 2-28.

—. (2009). Reasons: Practical and Adaptive. In D. Sobel and S. Wall (Eds.). *Reasons for Actions*. 37-57. Cambridge UP.

Rippon, S. (2011). In Defense of the Wide-Scope Instrumental Principle. *Journal of Ethics and Social Philosophy*, 5(2): 1-21.

Rosati, C. (2016). Agents and "Shmagents": An Essay on Agency and Normativity. In R. Shafer-Landau (Ed.). *Oxford Studies in Metaethics* (Vol. 11). 182-213. Oxford UP.

Scanlon, T. (2000 [1998]). *What We Owe to Each Other* (3rd Edition). Harvard UP.

Silverstein, M. (2014). "The Shmagency Question." *Philosophical Studies,* 172: 1127–1142.

Schroeder, M. (2004). The Scope of Instrumental Reason. *Philosophical Perspectives,* 18(1): 337–64.

—. (2005). Cudworth and Normative Explanations. *Journal of Ethics and Social Philosophy*, 1(3): 1-28.

—. (2007). The Humean Theory of Reasons. *Oxford Studies in Metaethics* 2: 195–219.

—. (2009). Means-End Coherence, Stringency, and Subjective Reasons. *Philosophical Studies*, 143(2): 223-248.

—. (2012). The Ubiquity of State-Given Reasons. *Ethics*, 122(3): 457-488.

Setiya, K. (2007). Cognitivism about Instrumental Reason. *Ethics*, 117(4): 649-73.

Stalnaker, R. (1984). *Inquiry*. MIT Press.

Tenenbaum, S. (2007). Evaluation and Motivation Part Company? The Problem of Akrasia. *Appearances of the Good: An Essay on the Nature of Practical Reason*. 257-282. Cambridge UP.

Heuer, U. (2011). Beyond Wrong Reasons: The Buck-Passing Account of Value. In M. Brady's (Ed.). *New Waves in Metaethics*. 166-184. Palgrave Macmillan.

Uleman, J. (2016). Kant and Moral Motivation: The Value of Free Rational Willing. In I. Vasiliou (Ed.). *Moral Motivation: A History*. 202-226. Oxford UP.

Velleman, D. (2000). *The Possibility of Practical Reason*. Oxford UP.

—. (2009). *How We Get Along*. Cambridge UP.

Way, J. (2009). Two Accounts of the Normativity of Rationality. *Journal of Ethics and Social Philosophy*, 4(1): 1-9.

—. (2011). The Symmetry of Rational Requirements. *Philosophical Studies* 155(2): 227–39.

Worsnip, A. (2018a). What Is (In)coherence? In R. Shafer-Landau (Ed.). *Oxford Studies in Metaethics* (Vol. 13). 184-206. Oxford UP.

—. (2018b). The Conflict of Evidence and Coherence. *Philosophy and Phenomenological Research*, 96(1): 3-44.

—. (2018c). Review of Benjamin Kiesewetter's The Normativity of Rationality. *Notre Dame Philosophical Reviews*. https://ndpr.nd.edu/news/the-normativity-of-rationality/.

—. (2019). 'Ought'-Contextualism Beyond the Parochial. *Philosophical Studies,* 176 (11): 3099-3119.

—. (2020). Making Space for the Normativity of Coherence (2nd Draft). http://www.alexworsnip.com/research/#almostdone.