

Revisiting Normative Reasons from a Naturalistic Point of View: Replies to My Critics

MARKO JURJAKO

University of Rijeka, Faculty of Humanities and Social Sciences,
Department of Philosophy and Division of Cognitive Sciences, Sveučilisna avenija 4,
51000 Rijeka, Croatia
mjurjako@ffri.uniri.hr

ORIGINAL RESEARCH PAPER – RECEIVED 24. 3. 2026. ACCEPTED 7. 5. 2026.

ABSTRACT: In this paper, I respond to the commentaries by Attila Tanyi, Yifan Sun, Voin Milevski, and Regan Reitsma on my book *Normative Reasons from a Naturalistic Point of View*. The discussion is organized around several topics, including how methodological naturalism favors subjectivist theories of reasons; how to justify a response-dependence theory of normative reasons; why one should adopt idealized subjectivism rather than actual-desire subjectivism; and the upshot and targets of evolutionary debunking arguments in metanormative debates.

KEYWORDS: affordances, evolutionary debunking arguments, naturalism, normative reasons, realism/antirealism, subjectivist and objectivist theories of reasons

1. Introduction

This paper discusses the commentaries by Attila Tanyi, Yifan Sun, Regan Reitsma, and Voin Milevski on my book *Normative Reasons from a Naturalistic Point of View*. I would like to thank them for their careful and critical contributions. In what follows, I respond to their papers by organizing their comments according to the main themes they address.

The responses are mostly organized around three considerations. First, some criticisms are due to misunderstandings and are addressed by clarifying the view expounded in the book. Second, some responses, although not explicitly mentioned in the book, are consistent with it and can be considered extensions of its original arguments. Third, some responses go beyond the claims made in the book and concern genuine points of disagreement rather than issues of clarity.

The paper proceeds as follows. Section 2 addresses the issue of what kind of naturalism I adopt in the book. Section 3 discusses my favored account of normative reasons, which is cast in terms of a response-dependence theory with an idealization component. Section 4 further develops this

account and defends it against several objections. Section 5 explains why I prefer idealized subjectivist accounts over actual-desire views of normative reasons. Section 6 revisits the incoherence argument against subjectivism. Section 7 discusses how I understand evolutionary debunking arguments and the role of naturalistic and evolutionary considerations in justifying metanormative accounts and genealogical explanations of our normative capacities.

2. What Kind of Naturalism?

2.1. Methodological versus Ontological Naturalism

In the book, I adopt methodological naturalism as the guiding framework. This perspective involves both methodological and ontological commitments (Papineau 2023). The methodological aspect refers to a general deference to the authority of our best scientific theories when devising an account of normative reasons. The ontological aspect refers to the broader aim of providing an account of normative reasons that can be accommodated within the world of facts as revealed by our best scientific theories (Harman 2012). In the book, I do not go into these methodological issues in much depth. This has prompted several commentators to ask for further clarification, especially regarding the ontological aspects of my view of the relation between naturalism and normativity.

Tanyi begins his essay by trying to pin down what kind of naturalism is expounded in the book and what kind of metanormative positions it is supposed to justify. Similar questions emerge in Milevski's discussion as well. Tanyi usefully frames this issue with several propositions which might be used to determine the kind of realism about normativity someone accepts. They include the following claims when talking about normative reasons:

- 1) Normative properties and facts exist.
- 2) Propositions about normative reasons can be true or false.
- 3) Some of those propositions are true.
- 4) Propositions about normative reasons express beliefs.
- 5) Propositions about normative reasons are irreducibly normative, in the sense that they cannot be derived from or be defined in terms of non-normative propositions.
- 6) Normative properties are metaphysically robust or weighty. (Tanyi this issue: 19)

According to Tanyi, robust realism is the position that endorses all six of these propositions. In this regard, he reasonably indicates that while David Enoch (e.g., 2011) might be interpreted as accepting all of them, Derek Parfit

(e.g., 2011, 2017) would accept all but (6), since he rejects strong ontological views about the existence of normative facts and understands them simply as true normative propositions (for recent discussion, see Karimi 2024). Moreover, some traditional naturalists in metaethics appear to endorse all six of these propositions; for example, the so-called Cornell realists, who defend non-reductionist naturalism (e.g., Brink 1989; Sturgeon 1985). In contrast, analytic naturalists would reject only (5), since they hold that truths about normative reasons can be conceptually derived from putatively non-normative truths about rational thought and behavior (e.g., Smith 1994). Thus, according to Tanyi's classification, even naturalists seem to endorse a more robust conception of normativity than Parfit.

This characterization clashes with my own use of "robust realism", under which I include authors such as Parfit and Enoch. Admittedly, in the book I do not define robust realism by means of a precise set of criteria. My classification of Parfit's realism as robust is based rather on the fact that he understands normative truths as objective, mind-independent, irreducibly normative, necessary in a way analogous to mathematical truths, universal in scope, and largely discoverable through *a priori* reasoning. And from the perspective of the methodological naturalism I presuppose in the book, this type of Parfitian normative realism appears rather robust.

But leaving these terminological issues, aside, Tanyi asks where my view would be positioned in his framework. To be honest, I am still unsure where I stand on this list regarding normative properties. It would probably depend on what kind of normativity we are discussing. In the book I only focus on normative reasons, and the discussion there is compatible with endorsing (1)–(5). As for (6), I am less sure, because I am not entirely clear about what is meant there by metaphysical robustness. If the claim is that the property of being a normative reason can play explanatory roles in some minimal sense, then this might be compatible with my view. If it is something else, then it would be necessary to see what this else is. However, and crucially, the list does not include the mind-independence criterion, which is central to the book, since I argue there that methodological naturalism is not compatible with mind-independence about normative reasons.¹ This is the main reason why I am interested in discussing Parfit's arguments, especially the incoherence argument, which purports to show that subjective theories—typically associated with the mind-dependence criterion—are false. If that argument succeeds, then naturalism, as I understand it, cannot

¹ In the book and throughout this paper, I use the terms mind-, stance-, or response-dependence to express roughly the same view, namely that facts about normative reasons are not wholly independent of people's behavioral, affective, and reasoning capacities. In Section 3 below, I discuss my concrete account.

be made to work. I come back to the incoherence argument in Section 5 below.

Tanyi (2026: 19) asks whether I am a substantive naturalist, where this is understood as an approach that interprets normative concepts and properties in terms of natural properties. It depends on what is meant by substantive naturalism. Throughout the commentary, Tanyi ties naturalism to metanormative realism as provided in his list of 6 propositions. However, it is not clear what this realism really amounts to because he does not include mind-independence as a criterion there. As I argue in Chapter 5 of the book, and in line with authors such as Sharon Street (2006), I reject the view that normative truths are completely mind-independent. So, if substantive naturalism presupposes mind-independence, then I am not a substantive naturalist. If it does not, I would be inclined to say that I am a substantive naturalist, though only in a qualified sense. This is because I am not sure whether a sufficiently clear and precise definition of a natural property can be provided, especially one that is entirely untainted by anything normative. This becomes clear when attention is paid to the special sciences, including biology and cognitive science, which employ normative notions such as rationality without being explanatorily worse off, even when these notions cannot be fully accounted for in purely natural or descriptive terms (e.g., Bermúdez 2011; Okasha 2018). Accordingly, the aim of my methodological naturalism is not to pursue such formidable reductionist projects, but to offer a naturalistically friendly account of normative reasons grounded in the explanatory resources of the relevant sciences.

2.2. Naturalism, Subjectivist, and Objectivist Theories of Normative Reasons

In the book, I argue that a naturalistic methodology would favor subjectivist theories of normative reasons when these are contrasted with objectivist accounts. I follow Parfit (e.g., 2011: vol. 2, 489–90) in understanding objectivism as the view that truths about normative reasons are necessary and mind-independent, akin to true mathematical propositions. I understand subjectivism as the negation of this position: normative truths are not mind-independent and necessary as putative mathematical truths, but rather contingent on features of the human condition.

An initial argument provided in the book (Sec. 2.2) for adopting a subjectivist view is a puzzle that objectivist views generate. The puzzle is based on Christine Korsgaard's (2011) remark that normative reasons are supposed to be facts that tell us what we should do and how to act. In other words, they are considerations that count in favor of performing some action or adopting

some attitude (see also Rast 2024: 162). Not only that, but reasons come pre-packaged with weights, and their balance is supposed to determine our choices in particular situations. If these reasons are understood as completely mind-independent and necessary, like mathematical truths, it becomes puzzling how they could be equipped with such normative weights and authority as to require our rational obedience. Precisely because they are mind-independent and impose authority from the outside, as it were, they leave open the question of why we should care about such facts, even if they existed. In contrast, adopting a subjectivist view—according to which facts about normative reasons reflect contingent facts about agents, their concerns, and their felt ideals—does not give rise to the same difficulty.

2.3. Are Objectivist Views Really Puzzling?

Sun (2026: 40–41) asks for further clarification regarding the nature of the puzzlement and why it poses a problem for objectivists. He rightly points out that if the puzzlement concerns how normative beliefs can motivate, then it does not genuinely threaten objectivists who maintain that normative facts are mind-independent. However, if the puzzle instead concerns the causal efficacy of normative facts themselves and their capacity to motivate agents, then objectivists can simply dismiss the problem, since they typically regard normative facts as causally inert.

The puzzlement, as I understand it, is not primarily about the motivational power of normative beliefs, nor about the causal capacity of normative facts to move us. The difficulty concerns their content, which seems to differ significantly from the standard realist understanding of mathematical facts. Unlike mathematical facts, truths about normative reasons are about us. The question is: how can necessary, mind-independent facts determine what we—agents with contingent features, capacities, and histories—ought to do, want, or think? On the Parfitian picture, reasons are supposed to bear relatively determinate weights and balances, as if pre-calibrated to guide our deliberation and resolve the practical problems we face as contingent beings. Yet their existence does not depend on those problems, nor on our contingent practical standpoint (see Korsgaard 2011). Thus, the puzzlement is about how could there be such a correspondence between a necessary normative structure and the shifting practical situation of finite agents?

Sun is not convinced that even this formulation presents a problem for objectivists, since the objectivist need not deny that normative reasons can depend on one's contingent situation. To illustrate this, Sun (2026: 41) offers the example of having a reason to call his mother when she expects the call, but not when she is taking a nap. The objectivist need not deny that the

reason depends on contingent facts about the person's situation, such as whether the mother expects the call or is asleep. However, it seems to me that precisely such examples generate the initial puzzlement. The relevant question is what makes it the case that he, rather than I, has a reason to call his mother, if it is not some fact about him, his mother, or their relationship. A natural answer would be that his caring about his mother gives him a reason to call her at a certain time of day. For me, the plausibility of this answer points away from an objectivist, necessitarian, view of the grounds of normative reasons and toward a subjectivist account.

Of course, as Sun suggests, the objectivist might respond by claiming that even though the facts about one's situation are contingent—for example, that this is my mother and that she is expecting my call—once one is in such a situation, there is a necessarily true normative proposition that one has a reason to call her. I can grant that this is a position an objectivist could adopt. Still, I find it puzzling what ultimately justifies this view. If we press further and ask what makes the fact that this person is my mother, and that she is in a certain situation, a reason for me to act, the answer will likely appeal to the fact that I care about her. But imagine instead that my mother had always treated me badly and emotionally damaged my life. In that case, we might think I have reason not to call her—or even to call precisely when she is taking a nap.

At this point, it seems to me, the Parfitian objectivist would have to say that the fact about caring is what makes it necessarily true that one has reason to perform a specific action. But now it seems that a fact about ourselves determines that other facts count as reasons. This prompts the further question: what makes this fact, namely, that I care about someone, such that it necessarily generates a reason for something? Here again, the objectivist appears to face a difficulty. They might say that this is a primitive necessarily true normative proposition. Or they might attempt to derive its necessity from more fundamental normative truths that according to them constitute the normative bedrock. Whatever the proposed normative bedrock, it seems to me it will still prompt the question how is it that these objective normative facts so neatly align with the considerations we already take to provide us with reasons, such as the fact that we care about certain people, and that such reasons have pre-calibrated weights? It seems to me that this brings us back to the original puzzlement, and I am not sure what the objectivist can offer by way of further explanation.

Subjectivists, on the other hand, have a straightforward answer to this issue, since they claim that normative reasons are essentially grounded in facts about cognitive agents and their desires, concerns, and cares (see also Reitsma 2026). However, Sun worries that simply saying this places sub-

jectivists in the same puzzling position as objectivists, since claiming that facts become reasons by responding to our concerns under certain conditions might appear to be an unmotivated, *ad hoc* stopping point. I do not share this concern. My preferred version of subjectivism does not aim to justify the claim that, by formulating a subjectivist theory, we thereby establish necessary truths of any kind. Rather, it claims that some normative facts begin with us as contingent beings. This view is sufficiently simple and coherent, and it makes it unsurprising why things we care deeply about provide reasons for us, and why, if we were different, this would not be the case.

Subjectivist views come in different forms. A basic distinction is between actual-desire views, which ground reasons in our actual desires, and idealized views, which ground them in what we would desire under certain conditions. My preferred account belongs to the latter group. I return to this preference in Section 5, where I respond to Reitsma's comments. In the next section, I set out my preferred version of this view.

3. A Response-Dependence Account of Normative Reasons and Naturalism

3.1. An Outline of the Account

In chapter 3 of the book, I offer a response-dependence account of normative reasons as an instance of an idealized subjectivist account, and defend it against objections. According to it:

(RD) The fact that p is a normative reason for X to F iff X is disposed to F on the basis of p , in conditions where X is rational. (Jurjako 2024: 69)

Several commentators raise issues with this account. I will start with those that raise issues about (RD) and its relation to naturalism.

Tanyi (2026: 20–21) claims that (RD) can express a naturalist view, but that it need not do so. He raises this issue in the context of asking whether my position amounts to substantive naturalism. As noted earlier, my view might be considered substantively naturalist in one sense of that term. However, Tanyi's concern seems to be that merely endorsing something like (RD) does not by itself justify substantive naturalism. I would agree with that point.

As already mentioned, the aim of the book is not to provide an independent defense of naturalism. Rather, it proceeds from a commitment to methodological naturalism and asks which account of normative reasons would be most defensible within that framework. My claim is that something along the lines of (RD) can be justified from a naturalistic perspective.

Tanyi (2026: 23) complains that even if subjectivism—and its variant captured by (RD)—is compatible with naturalism about reasons, it is not obvious that it accounts for the “counting in favor” relation itself, rather than merely providing the grounds of reasons. After all, many authors, including Reitsma in his commentary, primarily understand subjectivism as offering an account of the grounds of reasons, for instance in an agent’s desires or concerns, and not as explicating the “counts in favor of” relation (for further useful discussion, see Tanyi 2026: footnote 21 and the surrounding text).

However, even if subjectivism is primarily understood as an account of the grounds of reasons, it has sufficient resources to offer a more ambitious account of the counting-in-favor relation. I believe that the viability of this ambition can be demonstrated by unpacking (RD) more explicitly. According to (RD), a fact p is a normative reason for X to F iff X is disposed, under conditions of rationality, to F on the basis of p . The relevant disposition may take different forms: it may consist in acting on p , in incorporating p into one’s deliberations, in treating p as motivating, or something along these lines. In this regard, (RD) can offer and is intended as an account of the reason-relation itself, rather than just to identify the grounds of that relation. In fact, the grounds may vary: they may consist in features of the agent (for example, their desires), features of the world (facts not explicitly involving agents), or some combination of both. What unites these grounds as reasons for something is their role in determining a rational agent’s thought and behavioral processes.²

A further complaint is that if the view is to qualify as naturalist, then the notion of rationality invoked in (RD) must itself be fully naturalized. However, as noted across the Précis and these responses, I operate within a methodological form of naturalism. On this approach, there is no requirement that rationality should be reducible to entirely non-normative facts. Rather, I presuppose a relatively thin conception of rationality, which in its most basic form approximates instrumental rationality and is broadly continuous with its use in the sciences of mind and behavior (see, e.g., Bermúdez 2011;

² I like to think of this point in terms of a simple Bayesian agent who regulates beliefs in accordance with evidence and acts in accordance with expected utilities. In the epistemic case, reasons are probability-raising (or lowering) pieces of evidence. In the practical case, they are factors that affect the expected utility of available actions. The counting-in-favor relation is thus captured by probability- and expected-utility-affecting considerations. What makes these considerations reasons is their role in regulating the agent’s cognitive and behavioral processes, with normativity deriving from the rationality of those processes not the facts that provide inputs. Importantly, the agent need not represent reasons in those terms; this is a theoretical model, not a phenomenological claim about how reasons appear to us. I believe structurally similar views can be found in various strands of practical philosophy (see, e.g., Gaus 2011; Smith 1994). For a more substantive view of this kind, see Section 3.2. in the main text.

Dohrn 2024; Jurjako 2025; Okasha 2018: Ch. 6). This also means that I do not assume that rationality guarantees convergence in normative judgments across all agents regardless of their contingent beliefs, desires, motivations, and other psychological features. At the same time, I deliberately leave the notion of rationality somewhat open-ended, following Bernard Williams's (1981) idea of a sound deliberative route. In a constructivist spirit, rationality—like our conception of reasons—can be articulated, refined, and negotiated in light of agents' lived experiences, social contexts, and interpersonal relations.

3.2. Subjectivism and Constructivism: Unnatural Allies?

Here seems to be a good place to address a concern raised by Reitsma (2026: 76–77). In several places in the book, I express sympathy for Korsgaard's constructivism and her idea that reasons should be explained in terms of rational capacities, and not the other way around. However, it is well known that Korsgaard's metanormative views are Kantian, whereas my subjectivist account does not admit of universally shared reasons that can be reached on purely *a priori* grounds. Does this create a problem for combining these views? I would say that it does not.

The part of Korsgaard's view that I find congenial is the constructivist idea itself. However, there are different forms of constructivism. In more standard terms, we might distinguish between Kantian and Humean forms of constructivism (see Street 2010). I see my subject-based account of reasons as compatible with a kind of Humean constructivism, according to which reasons emerge from our capacity for reasoning, but where that capacity, in its most basic form, is relatively thin and must be filled out by the motivational states that contingent beings like us typically possess.

In this respect, my view is also compatible with Williams's (1981) internal reasons doctrine. According to Williams, reasons are grounded in an agent's motivational set, which includes, among other things, dispositions to act and different forms of practical reasoning. On this picture, reasons are grounded in procedures of reasoning that themselves operate within an agent's motivational set. Although it makes sense to think of these reasoning procedures as relatively stable structures that filter and organize the elements of the motivational set, the forms of practical reasoning we ultimately accept as valid—and the substantive reasons that emerge from them—cannot always be sharply separated from the motivational elements on which they depend. The reasons we find compelling will thus be the outcome of the contingent developmental paths of individuals and their dispositions.

This contrasts with Korsgaard's Kantian view, which sharply distinguishes the procedures of rational reasoning that are shared by all rational agents simply in virtue of their rationality. From a naturalistic point of view, I do not think that the existence of such universally shared rational procedures can be justified. For this reason, I find Williams's more contingent view of motivational sets and their associated standards of reasoning more congenial to the overall view and methodological naturalism I expound in the book.

4. Response-Dependence and Idealization

4.1. Brief Summary of the Argument

In the book, I treat the justification of the idealization component as a central issue for subjectivist accounts of normative reasons, such as (RD). This connects with the influential objection raised by Enoch (2005), who argues that subjectivist accounts face the difficulty of explaining why we should think that, in order to determine what our reasons are, we need to idealize. Idealization makes sense, Enoch suggests, in cases where the subject matter we are trying to determine is independent of our responses or of the procedure we use to determine it. However, subjectivist accounts that ground the nature of normative reasons in our idealized responses to certain facts appear to collapse the distinction between the nature of reasons and the procedure used to identify them.

In the book, I offer two routes for justifying (RD) and its idealization component. The first is to grant Enoch that commonsense discourse about normative reasons may be objectivist, while arguing that naturalistic considerations nevertheless favor a response-dependence account. This involves accepting a form of revisionism about our commonsense ontology of normative reasons. To make this argumentative move more salient, I draw an analogy with color perception. There is an influential line of argument claiming that, although it may intuitively seem that colors are objective properties of external objects, the science of color vision provides reasons to revise this objectivist picture and to adopt a response-dependence account of colors, according to which colors are properties instantiated only in relation to agents capable of color vision (see, e.g., Miščević 2007; for a more recent and critical discussion, see Watkins & Shech 2025).

The second route is to deny Enoch's contention that the commonsense ontology of normative reasons presupposes objectivism. In support of this claim, I appeal to empirical studies suggesting that ordinary people's judgments about normative issues—including morality—are often not, or at least not clearly, interpretable as presupposing normative objectivism (see,

e.g., Pözlner 2023). On this view, the response-dependence account of normative reasons can be presented as offering an interpretation of people's normative judgments that is not revisionary.

Finally, I argue that the main criterion for justifying idealization in articulating an account of normative reasons is not the assumption of objectivism, but rather the account's capacity to distinguish appearance from reality—that is, to make sense of the possibility of true and false judgments about one's reasons. I contend that both revisionary and non-revisionary interpretations of (RD) can meet this condition. For that reason, the account does not face insurmountable difficulties in justifying its appeal to idealization.

4.2. Error and Normative Authority

Before moving on to objections to these lines of reasoning, I will address Milevski's concerns, which help clarify both the general contours of the account and the structure of the discussions that was just summarized. Milevski (2026: Section 4) raises questions about how an account of the (RD) type can sustain the distinction between error and correct normative judgments. He rightly construes my view as cognitivist, in the sense that judgments about reasons express propositions that are truth-evaluable. However, he raises three related concerns about my view.

First, he suggests that my view, as expressed in (RD), cannot genuinely be cognitivist because it may conflate the facts that a rational agent would endorse under suitable conditions with the facts that are normative reasons. Second, he argues that this conflation might contribute to "blurring the normative boundary" (Milevski 2026: 58), by which he seems to mean that determining what reasons there are would be reduced to descriptive facts about agents and their responses. Third, he worries that the view risks an illegitimate form of circularity: if correct judgments about normative reasons are simply those endorsed by a rational agent under ideal conditions, then there may be no independent ground for distinguishing correct from incorrect judgments, and the distinction itself threatens to collapse.

The second interpretation of the concern appears to be related to the issues associated with the triviality objection, which will be discussed below in Section 6.4. The first and third versions of the worry, however, seem to share a common ground; namely, the idea that thinking about what reasons there are in terms of rational endorsement collapses the distinction between truths about reasons and truths about rational agents. In one sense, this is correct, since the core idea of the response-dependence account proposed in the book is that, ultimately, the facts that rational agents would find relevant

for their decision-making and belief-forming processes are the facts that determine which considerations count as normative reasons. However, this view is not inconsistent with cognitivism about normative judgments, just as response-dependence about color is not inconsistent with there being truths about which colors objects have.

The appearance of vicious circularity here may arise from a failure to sufficiently distinguish the ontological account that (RD) is meant to provide from the epistemological question of how we are to determine which facts are reasons for what. (RD) is primarily an ontological account of the nature of normative reasons. It tells us that, ultimately, what counts in favor of what depends on the responses a rational individual would have to certain actions in certain contexts. Importantly, as mentioned in several places throughout this paper, this ontology is defended on the basis of naturalistic considerations and the difficulties facing objectivist accounts. Once we have settled on the general shape of the ontological account of normative reasons, we can then turn to the epistemological procedures by which we determine what those reasons are.

Here, the analogy with color perception is instructive. A response-dependence account of color can be supported by naturalistic arguments about the nature of visual perception. Suppose that the color green is identified with what an ordinary perceiver would judge to be green when presented with, say, lettuce or grass under normal lighting conditions. The relevant naturalistic considerations give us reasons to formulate such an account. At the same time, they also provide the resources for thinking about how to epistemically justify particular color judgments in specific cases. But this does not render the account viciously circular. Similar considerations apply, I suggest in the book, to an analogous account of normative reasons.

4.3. Objections to Revisionary and Non-revisionary Routes to Response-dependence

Several commentators put pressure on the revisionary and non-revisionary argumentative strategies discussed in Chapter 3. Regarding the revisionist route to response-dependence, Tanyi (2026: 31) is skeptical of the analogy with the way our account of color ontology might be revised on the basis of naturalistic considerations. The difference, he suggests, is that our discourse about color seems more amenable to revision concerning its underlying ontology. If we revise our view of colors from being intrinsic properties of external objects to response-dependent properties associated with our visual systems, nothing of great importance appears to change. However, the worry is that revisionism about normative reasons is different. If normative reasons

are not irreducibly objective features of the world, then they might lose their normativity. In that case, instead of merely revising our account of their nature, we might end up eliminating them altogether.

While I agree that one possible upshot of revising our reasons-discourse might be some form of ontological normative nihilism, I am not overly worried about this consequence.³ Methodological naturalism does not necessarily lead to the elimination of normativity. As already mentioned, subjectivism in the form of (RD) allows that reason-talk is explicated in terms of rationality, without attempting to analyze the concept of rationality in entirely non-normative terms. Of course, one might object that the thin notion of rationality I adopt is insufficient to sustain the conception of normativity as involving necessary and mind-independent truths, as envisaged by, for instance, normative realists such as Parfit. But that is precisely why the account would be revisionary: it does not aim to accommodate any kind of understanding of normativity. Crucially, however, it would not eliminate normativity as such.

Sun (2026: 45) raises the concern that when we conclude that some part of our discourse is flawed and in need of revision, any successful revision should preserve the reasons that sustain the practice in the first place. Otherwise, we risk undermining the very justification for engaging in it.⁴ More specifically, Sun worries that if we abandon the idea that specifically moral reasons are grounded in mind-independent normative facts—facts that provide us with categorical reasons for action—we may lose or undermine the point of retaining the concept of moral reason at all.

There are several things I would like to say in response to this concern. It should be noted that in Chapter 6 of the book I am interested in offering an account of how facts about categorical reasons—or at least our perception of them as categorical—may emerge from a naturalistic perspective, given the kinds of creatures we are, the capacities we possess, and the social relations we engage in. However, my primary concern in the book is not with moral reasons *per se*. One reason for this is that it is not clear that we can draw a sufficiently useful distinction between moral and non-moral reasons that a revisionary response-dependence account would need to preserve (see also Tanyi 2026: 34–35).

³ One reason I am not worried is that there are interesting recent error-theoretic accounts of normativity that aim to show why eliminativism need not generate insurmountable philosophical problems. For recent discussions, see, e.g., Khuramy and Schulz (2026), Olson (2014), and Taccolini (2024).

⁴ For a more general discussion of how practices can function as arbiters of ontological commitment, see Gan (2024).

Moreover, even if there were a defensible way to define reasons as distinctively moral—such that part of what makes them moral is that their authority is categorical and independent of the contingent perspectives of individual agents—it would remain an open empirical question whether discovering that there are no such reasons would render the practice of employing the concept of moral reasons pointless. I am not aware of specific studies that would provide evidence on this point, but my hunch is that realizing there are no such mind-independent moral reasons would not undermine, beyond repair, our moral practices or the notion of reason employed within them (see also Ventham 2023). This is because, from a naturalistic perspective, morality as a system of rules governing how people ought to behave toward one another has clear adaptive value, which can be explained in evolutionary terms (Sterelny and Fraser 2017). Of course, my view does not imply that all rational individuals, after sufficient reflection, will conclude that they have an undisputed reason to accept the authority of prevailing moral norms. However, given our social and institutional arrangements and our nature as social creatures, we can expect that most of us, most of the time, will find it reasonable to subordinate ourselves and align our conduct with widely shared moral norms, and will thus think that moral reasoning has a point within those practices (see Gaus 2011).

Concerning the non-revisionist line, both Tanyi (2026: 33–34) and Sun (2026: 46–47) put pressure on my use of empirical studies to undermine the claim that our ordinary conceptions of normativity rest on objectivist foundations. Tanyi emphasizes criticisms that have been directed at empirical studies of moral and, more generally, normative judgments. Some philosophers have argued that the results of these studies are not especially significant, since they often yield conflicting findings and rely on the unconsidered intuitive responses of ordinary people. In contrast, we are typically more interested in considered judgments formed through reflection, of the kind often emphasized in philosophical inquiry (see, e.g., Kauppinen 2007). Sun suggests that the studies I mention in the book, which typically aim to test people's judgments about the objectivity of morality, do not directly support response-dependence about normativity. Instead, he argues that the better interpretation of these studies is to regard them as metaethically neutral, by which he means that they do not have clear implications for people's intuitive conception of the objectivity of morality and thus do not count against objectivism.

While I cannot here embark on a wholesale defense of empirical studies investigating folk metaethical and metanormative commitments, I would like to point out several things. Regarding Tanyi's objection, it should be emphasized that in response to early criticisms, several philosophers and

philosophically informed researchers have refined their methodologies for investigating people's moral intuitions. In particular, they have developed studies designed to elicit more reflective judgments from ordinary participants about normative issues (see, e.g., Pölzler and Cole Wright 2023; Pölzler, Zijlstra, and Dijkstra 2024). Based on these results, they argue that it is far from clear that ordinary people hold straightforwardly objectivist metaethical commitments (for more general discussion, see Pölzler 2023). For that reason, I would reject as unfounded the criticism that empirical studies cannot yield interesting or philosophically relevant results concerning ordinary people's metaethical commitments.

Furthermore, even if we agree with Sun that the best interpretation of the studies testing people's intuitions about the objectivity of morality is that they are metaethically neutral, that is entirely consistent with the argument offered in the book. I invoked those studies primarily to undermine Enoch's claim that our normative discourse essentially presupposes an objectivist ontology and would not make sense without it. If the relevant studies instead suggest that people's intuitions are neutral in this respect, that still supports my case against Enoch. If those intuitions are neutral, then they do not necessarily support objectivism. This leaves open the possibility that a response-dependence account can offer a legitimate non-revisionary interpretation of the ontological commitments underlying people's normative judgments.

5. What Kind of Subjectivism to Adopt?

5.1. Why Adopt an Idealized Subjectivist View of Normative Reasons?

At this point, we might ask: if idealized subjectivism of the (RD) type involves such issues, why not opt for a subjectivist account that avoids idealization? This connects to Reitsma's (2026: Secion IX) discussion, where he asks why I adopt a subjectivism with an idealization component rather than a straightforward subjectivism according to which actual unmotivated or basic desires underpin normative reasons.

Idealized subjectivist accounts are typically motivated by examples of the Williams-style gin-and-tonic sort (1981). To refresh our memories, this example involves a person who wants to drink a gin and tonic, but unbeknownst to her, the bartender gives her a glass full of petrol. The question is whether she has a reason to drink from the glass. The intuitive answer seems to be that, although if she were to drink from the glass we could make rational sense of her behavior, she does not in fact have a normative reason to do so. And if we ask why she lacks such a reason, the natural answer is that, if she knew the relevant fact about the contents of the glass, she would not

have the desire to drink from it. What we do here, then, is to look at the situation from an idealized point of view and ask what this person would desire if she were rational and informed about the relevant facts concerning how to satisfy her desire.

Before we turn to Reitsma's reasons for claiming the superiority of the actual-desire view of normative reasons, a clarification is in order about how to understand idealization in this context.

5.2. The Advice Model and the Example Model of Idealization

There are at least two ways of thinking about how the idealized point of view is supposed to determine what normative reasons an agent has. Michael Smith (1995) usefully distinguishes between the advice model and the example model. According to the advice model, what an agent has a normative reason to do is what their rational counterpart would suggest or advise them to do in the relevant circumstances. According to the example model, what a person has a reason to do is what their rational counterpart would want to do in similar circumstances and her desires and overall behavior provide an example that the less fully rational self should follow.

Although in some cases both models will have the same verdict about what reasons a person has, in other cases they will diverge, and when they do it becomes clear that the advice model is superior at capturing how we think about normative reasons. To illustrate this point, Smith uses the following example. Imagine that in a game of squash I am defeated so badly that I feel humiliated and filled with rage to the point of wanting to hit my opponent on the head with a racket. We can also imagine that from an idealized point of view we would conclude that, given that the opponent beat us fairly and squarely, we do not have a reason to hit them. However, we can still ask what exactly we should do in such a situation. This is where the two models diverge.

According to the example model, we should consider what our rational and informed counterpart would do in that situation and follow their lead. However, this seems problematic, because our rational and fully informed self would presumably conclude that what they should do is go to the opponent and congratulate them by shaking their hand. Yet this appears to be the wrong verdict about what to do, because in our actual situation we are enraged, and approaching the opponent might in fact lead us to hit them on the head with the racket. In contrast, according to the advice model, our rational counterpart would take into account our situation and the fact that we are not fully rational at that moment, and would suggest that in this particular case it might be better simply to leave the squash court politely.

Now, even though I am not entirely explicit about this distinction in the book, I nonetheless emphasize the advice model as part of what structures our thinking about reasons (see Jurjako 2024: 29–35). Moreover, in the book I do not assume that the rationality that fixes the reasons applying to a person is always determined by that person being fully rational and *explicitly* reasoning about her situation. Instead, the view of rationality developed in the book treats rationality as involving certain dispositions to avoid self-defeating patterns of action, thought, and desire, which of course do not and cannot always be manifested (see Jurjako 2024: 69–70). These features of my view will be relevant for addressing the objections raised by Reitsma.

5.3. Is Actual-Desires Subjectivism Superior to Idealized Subjectivism?

Reitsma asks why, instead of adopting an idealized subjectivist view, we should not adopt a more straightforward actual-desire subjectivism that grounds normative reasons in any desire a person might have. Or, to put it more precisely, “to treat, any given (satisfiable) basic, unmotivated desire as (perhaps defeasibly) rationally potent” (Reitsma 2016: 81). Reitsma claims that this view, apart from being able to capture everything we want from a subject-based theory of reasons, also does not face the problems confronting versions of subjectivism that involve some kind of idealization condition.

To demonstrate the superiority of his view, Reitsma first explains that actual desire subjectivism can accurately capture our intuitions about reasons in gin-and-tonic-type examples. Following Reitsma (2026: 82), let us call the person in the gin-tonic example Sippy. Although it would be rational for Sippy to drink from the glass given her beliefs, this is not something he has a decisive reason to do. Reitsma argues that this verdict can be explained by distinguishing between two perspectives from which we can evaluate Sippy’s case. One he labels the perspective of rational advisability, and the other the perspective of rational virtuosity (see Reitsma 2026: 79). In a nutshell, these two perspectives answer different questions. The former addresses the question of what it is most advisable for the agent to do, while the latter concerns how virtuous an agent is in performing practical reasoning. If I understand him correctly, of the two perspectives rational advisability is the more fundamental, because it answers the question of what agents have normative reasons to do, whereas rational virtuosity concerns how well agents exercise their reasoning capacities when trying to determine what they have reason to do.

When this framework is applied to the case of Sippy, the perspective of rational advisability tells us that Sippy ultimately ought not to drink from the glass, because doing so would not satisfy his original desire to drink a gin

and tonic. At the same time, we can say that from the perspective of rational deliberation, Sippy is perfectly rational, since he has no reason to believe that what is in the glass is not what he actually wants to drink. However, given that the rational advisability perspective provides the more fundamental consideration about what Sippy ought to do, actual desire subjectivism can give the right verdict in this case.

Now, with certain caveats, I believe Reitsma and I share a similar perspective on the Sippy case. The perspective of rational advisability seems to correspond to my idea of the perspective that determines what we have normative reasons to do, namely what would advance our desires, needs, concerns, and so on under conditions of rationality.

However, Reitsma (2026: 85–86) appears to disagree on this point. There may be two possible explanations for this disagreement. One is that he does not realize that his own account actually involves a form of idealization that comes quite close to my view. The other is that he may be interpreting my view as committing me to something stronger than I intend. Let me begin with the first possibility.

Reitsma takes his account not to involve any form of idealization when determining what reasons Sippy has. His view is that what ultimately determines Sippy's decisive reason is his "basic, unmotivated desire for a gin and tonic" (Reitsma 2026: 84), together with an instrumental principle according to which the normative force of reasons is transferred from the desire that generates the reason to the most (objectively) effective means of fulfilling that desire. I am not entirely sure what Reitsma means when he says that Sippy has a basic and unmotivated desire for a gin and tonic. In terms of its complexity, the desire does not seem particularly basic. In terms of being unmotivated, perhaps this could be the case, but if a desire for a gin and tonic can be unmotivated, it seems that many other desires could be as well. In any case, here the important point is that Reitsma (2026: 81) allows that even unmotivated desires can be defeasible in providing reasons. I take this to mean that with further information we may come to think that an unmotivated desire no longer generates a reason.

In this regard, it seems to me that Reitsma's view corresponds closely to the view elaborated in the previous subsection. If Sippy were rational—in the sense captured by the advice model, according to which his counterfactual rational self would advise him not to drink from the glass—then this is what Sippy should do. Indeed, I would say that Reitsma's own account implicitly involves an idealization of this kind. This is suggested by his language of rational *advisability*, which determines what a person should do given their desires, and the *defeasibility* of basic desires that ground reasons. It is hard to see what it could mean to say that something is rationally advisable

except in terms similar to Smith's advice model, namely that it is what a rational individual, who is in a better epistemic position than the actual agent, would advise them to do.

The other reason that might explain why Reitsma does not see the correspondence between our views is that he understands the idealization condition in more substantive terms. This becomes clear when he discusses an objection often directed against counterfactual analyses of normative reasons. The worry is that making an agent too rational or too informed, relative to how the individual actually is, may undermine the very desires that are supposed to ground the relevant reasons. For instance, Reitsma (2026: 85) suggests that if Sippy were to engage in a process of rational deliberation about his situation and begin considering the possibility that a waiter might be dishonest, this could frighten him and make him suspicious of waiters in general, leading him to lose any desire to drink in a bar again. Reitsma then asks why we should think that whatever desire Sippy might counterfactually acquire through such a process of idealization—after becoming frightened in this way—should determine his reasons, rather than thinking that the relevant reason is grounded in his basic desires as they currently are.

Here again, I think that we are in greater agreement than it might initially appear. As clarified earlier, my view does not presuppose that all reasons are determined by engaging in actual or counterfactual reasoning processes. I am explicit about this in the book when discussing what it means to say that facts are reasons if a person is disposed to react to them in a certain way under conditions of rationality (see Jurjako 2024: 69–70). Furthermore, Smith's advice-model is compatible with the view that facts about the agent as they currently are determine what they should do. This model does not imply that the agent must counterfactually go through some explicit process of deliberation and information gathering that would fundamentally change the constellation of desires grounding their reasons.

Given this, do I claim that there is basically no difference between Reitsma's and my understanding of what ultimately grounds reasons? No, I would not say that. Reitsma's ultimate view seems to be that there are some basic, unmotivated desires that provide the most fundamental reasons. Insofar as there are such desires, they are simply the desires that generate reasons, and nothing more can be said about this. Moreover, there might be such desires whose reason-generating status should take precedence over what we might conclude after rationally deliberating about the world and our relation to it. I am not so sure that there are such desires, or why we should think of them as basic in this way. In fact, this is where I see the force of Sun's (2026: 42) objection, namely, why think that grounding reasons in

desires that might happen to be basic to an agent is any less problematic than grounding them in objective facts about value.

I think that idealized subjectivism, together with methodological naturalism, can provide an answer to this worry. If it really is the case that there are some basic desires that we would think of as sources of reasons, then this can be explained as the view we would come to accept because it fits with our naturalistic worldview. But if that is so, then we reach this conclusion through some form of reasoning about the issue. Furthermore, once this is allowed, it is not clear why we should stop ourselves from thinking that our reasons may be adjusted as we reflect on the situations we are in, and that in this process we might lose or revise some desires—even if we once regarded them as basic in some sense. I do not think that there is a necessarily settled fact about this issue prior to our reflecting on it. And I welcome this consequence of my view. What we find in the motivational set with which we are endowed will also include our ideals about how the world should look (Williams 1981). I see no reason why a subjectivist or naturalist about reasons should restrict themselves to the normative authority of some subset of desires that they might, at some moment, regard as basic.

6. Between the Scylla of Reductionism and the Charybdis of Anti-Reductionist Naturalism

6.1. The Incoherence Argument

Although my aim in the book is not to offer a full-blooded naturalist reduction—whether of the analytic or the non-reductive substantive kind—Tanyi (2026: Section III) correctly observes that, in my discussion of Parfit's incoherence argument, the distinction between these forms of naturalism appears to generate theoretical tension concerning the nature of normative reasons.

Parfit's incoherence argument involves three premises that mutually support each other but a proponent of idealized subjectivism about normative reasons cannot accept all of them on pain of contradiction.

(M) what we have most reason to do is whatever would best fulfill, not our actual present telic desires or aims, but the desires or aims that we would now have, or would want ourselves to have, if we knew and had rationally considered all of the relevant facts.

(N) when we are making important decisions, we ought if we can to try to learn more about the different possible outcomes of our acts, so that we can come to have better informed telic desires or aims, and can then try to fulfil these desires or aims.

(O) these possible outcomes may have intrinsic features that would give us object-given reasons to want either to produce or to prevent these outcomes, if we can.

(Parfit 2011: vol. 1, 93)

(M) states a version of idealized subjectivism of which (RD) is an instance. (N) introduces an epistemic or methodological principle according to which having better-informed desires or aims requires being informed about the likely consequences of acting on them. Before moving on, it seems appropriate in this context to address questions raised by Milevski (2026: 55–56) regarding epistemic standards and procedures for justifying judgments about normative reasons.

6.2. Epistemic Justification of Normative Judgments

In Chapter 6 of the book, I discuss and provide a naturalistic-evolutionary story aimed at accounting for how basic reason-relations can become stabilized and how and why we come to consider or experience certain facts or considerations as intrinsically counting in favor of something. I see this process as analogous, if not identical, to the establishment of basic semantic relationships that have both descriptive and imperative aspects, such as the vervet monkey's warning calls indicating the presence of predators, or the honeybee's waggle dance indicating the location of food (see, e.g., Harms 2004; Millikan 1995). Milevski (2026: 55) rightly interprets this project as genealogical. The aim is not to directly justify normative judgments as true, but rather to provide a story about why we come to regard them as true, and even as possessing a kind of mind-independent authority. However, he argues that this leaves open the question of how we are actually to justify our normative judgments, if not by appeal to an evolutionary story about their adaptability.

In the book, the focus is not primarily on this epistemic issue. Nevertheless, if I were to develop an account consistent with the book's framework, it would likely proceed along the lines of (N). It captures the commonsensical idea that, if one wants to determine one's reasons for acting, one should reflect on what one is doing, how valuable it is to oneself, and what the likely consequences of one's actions will be. That would be the rough outline of such an account.

Developing a more substantive account goes beyond the scope of the book and the present discussion. This is because, I believe, it would require establishing standards of justification grounded in first-order normative theories whose role is to regulate individual and social behavior. Given that we are, for the most part, social beings, the determination of these standards will depend on our interactions, relationships, shared history, and, importantly, shared ideals. Accordingly, such standards will be established, negotiated, and renegotiated in light of historical circumstances and the concrete demands of everyday life (see Gaus 2011, 2013).

The further issue, which is of more direct concern to the book, is what kind of account of normative reasons can provide the grounding for such a view. This brings us to the discussion of proposition (O).

6.3. Analytic Naturalism as a Response to the Incoherence Argument

According to Parfit (2011), (O) provides an objectivist account of what gives us reasons to want to perform certain actions or pursue certain aims. Crucially, he claims, neither (M) nor (N) can be true unless (O) is true. However, the subjectivist cannot accept (O), since it appears to ground reasons in mind-independent normative facts. The result, Parfit argues, is that subjectivism is incoherent. Since I am sympathetic to subjectivist theories encapsulated within (M) and that compared to objectivist theories, naturalism favors something along the lines of (M), I feel compelled to provide a response to Parfit.

In the book I respond by arguing that the subjectivist endorsing (M) does not need (O), rather it is sufficient to justify (M) and (N) by accepting something like (O’):

(O’) possible outcomes may have intrinsic features that would give us subject-given reasons to want either to produce or to prevent some outcome. (Jurjako 2024: 59)

In the book, I concede that, if (O’) is read as presupposing that reasons are grounded in what our idealized selves would be motivated to do, this would render (M) analytically true.

6.4. The Open Question Argument and the Triviality Objection

However, in that case, my defense becomes vulnerable to several objections. Milevski (2026: 54) suggests that, if I adopt analytic naturalism, my position becomes vulnerable to the Moorean open question argument. Whatever natural property is proposed as identical with the property of being a reason, we can always meaningfully ask whether that property really provides a reason for something. Moreover, as Tanyi (2026: 27–28) notes, my defense also becomes vulnerable to Parfit’s triviality objections against analytic forms of naturalism. The core of the worry is that normativity would then be reduced to psychological facts about the agent—namely, what they would predict themselves to want under certain conditions—thereby threatening to drain the account of genuine normative force.

At that point of the discussion in the book, I do not address the open question argument nor the triviality objection because my focus is on discussing options for deflecting the incoherence argument and adopting the interpretation that makes (M) analytic might be one way to go. However, the

two concerns raised by Tanyi and Milevski seem to be two different aspects of the same general type of objection (Lutz and Lenman 2024). And in the book, I find the triviality concern as more directly pertinent to my discussion. This is for two reasons.

First, despite the intuitive appeal of the traditional form of open question argument, it seems to presuppose that analytic reductions should be obvious to competent concept users. However, in many interesting cases this will not be true. As Smith (1994: 37–9) plausibly argued, conceptual analysis and conceptual reductions are not required to be obvious. Consider an analogy with analytic reductions in other domains that are not obvious. We can have a long logical or mathematical derivation in which one equation is derived from another through a series of transformations, such that the end result is the output of an analysis of the initial formula. Even though the final result may look very different from the original equation—and for some people may invite the question whether the first equation is equivalent to the second—that question can still be answered by providing a justification of how each step of the analysis was performed. Similarly, if we have a justification for analyzing the concept of normative reasons in other terms, the fact that the analysis is not immediately obvious and may initially give rise to a sense of openness about whether it concerns the same thing does not, by itself, invalidate the analysis.

Second, the traditional form of the open question argument has the most bite against accounts that purport to reduce normative properties to properties described in a completely non-normative language. However, even though some commentators read me as endorsing such a reductionist view, such as Milevski (2026: 54) when entertaining the idea that I aim to reduce truths about normative reasons to what is adaptively useful, this is not my goal in the book.

The triviality objection seems to be the flip side of the open question argument. One aspect of the triviality objection is the claim that, if you attempt to analyze the concept of a normative reason in terms of some purely descriptive natural property, you risk either eliminating normativity altogether or rendering claims about reasons trivial (see, e.g., Lutz and Lenman 2024: Sec. 2.2). In particular, what is relevant for my discussion is the form of the triviality objection according to which, by reducing reasons to purely descriptive natural facts, one loses the normativity associated with them. For instance, if normative reasons are understood as those considerations that an individual would take into account when rationally considering what to do, this can seem to transform normative claims into mere predictions about what that individual would do under certain conditions.

However, as Tanyi (2026: 27–28) notes, I have a response to the triviality objection, since I address a variant of this objection in Section 2.4 of the book, where I discuss Williams’s (1981) account of internal reasons. There I adopted the line according to which Williams’s claims about internal reasons do not become pure predictions about what an agent would decide, since the concept of an internal reason is explicated in terms of rationality, and I argue that (instrumental) rationality is normative insofar it provides criteria of conduct that can be fulfilled or not fulfilled (see also Roberts 2005). I am not sure whether this would count, by Tanyi’s lights, as an endorsement of analytic naturalism. Nonetheless, within the methodological naturalism I employ in the book, I take myself to be permitted to adopt such a view of normativity.

Moreover, taking this perspective should help clarify why I find Milevski’s (2026: 55) worry unfounded, namely that adopting a reductionist form of naturalism is insufficient for providing a criterion for critically assessing social and other practices. Milevski suggests that, by reducing normative reasons to facts about what is functional or adaptive, we undermine the possibility of questioning whether what is functional or adaptive ought to be endorsed as justified in our lives. It should be clear that adopting a Williams-style internal reasons interpretation, together with my refinement in the form of (RD), does not lead to such a view. Thinking about normative reasons from the perspective of what facts a rational agent would consider relevant in their decision-making is meant to account for the “counting in favor of” relation, not to impose fixed limits on the substantive reasons people may have for performing different actions or for determining criteria of moral and social evaluation.

When it comes to the specific reasons people might have for action, I do think, as Milevski (2026: 54–55) nicely points out, that we can provide a naturalistic story, often in terms of adaptive and pragmatic value of taking or experiencing some facts as inherently counting in favor of something. However, the overall view espoused in the book is that, ultimately, which facts will be regarded as normative reasons will be negotiated among rational agents against the background of their concerns. Of course, those concerns will have been shaped by evolutionary, cultural, and ontogenetic factors. At the same time, my naturalism leaves it open which of these concerns will be endorsed, revised, or eventually abandoned through various social and biological processes, including rational deliberation and negotiation (for further discussion, see Jurjako 2024: Ch. 4).

6.5. Adopting a Non-Reductionist Understanding of Normative Reasons

In addition, when discussing internal reasons and Parfit's incoherence argument, I make a concession to Parfit and argue that even if we give up the claim that the notion of a normative reason can be explicated in terms of a relatively thin notion of rationality and accept that the concept of a normative reason is irreducible and indefinable, still it does not follow that we should adopt a non-subjectivist ontology of reasons (see Jurjako 2024: 60–62). In alignment with Street (2017), I argue that even if the concept of a normative reason does not admit of reductive analysis, this does not entail that, at the ontological level, it refers to some irreducibly mind-independent normative property.

Here, Tanyi (2026: 29–30) raises the question whether this move is consistent with my earlier response to the incoherence argument, where (M) was interpreted as analytic. The concern is that if I now adopt the view that the concept of a normative reason is primitive and not definable in other terms—and that (M) is therefore not analytic—then I appear to face the very problem I initially attempted to avoid. In that case, I would have to abandon the analytic subjectivist strategy and could no longer rely on (O') in defense of (M), since, as previously noted, (O') renders (M) analytic.

I do not think there is a substantial tension here. In responding to the incoherence argument, I suggest that the subjectivist need not rely on (O) as formulated by Parfit, but only on something like (O'), which is intended to capture an explanation why (M) would hold, that is consistent with a subjectivist view of reasons. Nothing in the general formulation of (O') implies that it necessarily makes (M) analytically true. Admittedly, in the book I used the unfortunate wording which suggested the contrary, that (O') necessarily makes (M) analytic. However, as I explain later in the book (pp. 61–62), if we agree with Parfit in treating the concept of a normative reason as primitive, then (M) can be understood as specifying truth-conditions for statements about reasons rather than as providing a semantic analysis of the concept itself. In that case, any specification of (O') intended to explain or support (M) should not be read as offering an analysis of the notion of reason as articulated in (M).

The question, then, is what kind of understanding of (O'), or what more general considerations, could justify treating (M) as providing truth-conditions for statements about reasons. I want to proceed cautiously here, since I allow that (M) might ultimately not turn out to be the best naturalist or subject-based account of reasons. Still, for the sake of argument, let us assume that it is. Recall that Parfit treats (O) as supplying reasons for believing that (M) and (N) hold. In the present context, the issue concerns what

subjective reasons we have for accepting (M). An answer consistent with the book and the proposed reading of (O'), but going beyond what is explicitly stated there, would be the following. If we were to reflect carefully and rationally on the nature of normative reasons—taking into account the relevant information and thinking clearly—we would come to regard certain facts as reasons for action. On this picture, reasons are those considerations that would survive such informed and rational scrutiny.

At first glance, this might look circular. It can seem as if I am saying that rational reflection tells us that reasons are whatever rational reflection would endorse. That may appear trivial or even as triggering a regress, as it seems to invite a further round of rational reflection to justify why rational reflection has this authority, and so on without end.

However, the reflection at issue is not happening in a vacuum. It is constrained by a broader framework, namely, the methodological naturalism adopted in the book. Crucially, within that framework, we are comparing two general approaches to normative reasons—objective and subjective theories—and asking which fits better with this naturalistic picture. My claim has been that subjectivist theories fit that picture more plausibly than objectivist ones. Thus, given those constraints, a rational agent reflecting on the nature of reasons would be led to favor something like (M). For that reason, I would also argue that there is no vicious regress. The process of reflection is anchored by methodological naturalism, and it is within those constraints that I claim (M) would emerge as the most attractive account of the truth-conditions of reason statements.

This type of argumentation might initially be dismissed as a sleight of hand, as if it simply relocates the problem rather than resolving it. However, I do not think it is any more *ad hoc* than the objectivist position that Parfit appears to endorse. On his view, if (M) is true, that truth rests on the existence of irreducibly normative and necessary facts that make (M) true—and there is little further story such a view can offer about why we should accept it. In contrast, my naturalist approach does not retreat to a quietist metaphysics. It instead seeks to embed the reason-relation within a broader, naturalistically respectable framework that remains continuous with what we know about human agents, their psychology, and their social practices, while at the same time respecting our normative intuitions about which considerations count as reasons and how we come to recognize them.

Of course, any naturalistic account must stop somewhere, because philosophical naturalism is not grounded in self-evident, *a priori* foundations. But unlike Parfit's quietist view, the naturalist can offer an explanation of why we are driven to seek a grounding for normative facts and why this project culminates where it does—e.g., in reflective deliberation constrained

by naturalistic conditions. In this sense, the naturalist strategy provides a kind of closure that non-naturalist theories typically do not.

7. Evolutionary Debunking: Against Mind-Independence of Normative Reasons

7.1. Naturalism and Antirealism about Normativity

In Chapter 5 of the book, under a significant influence of Street's (2006) seminal discussion, I use a version of the evolutionary debunking argument to challenge the claim that truths involving normative reasons are mind-, stance-, or response-independent. My central claim is that, from a naturalistic perspective, especially once we take into account the explanatory power of evolutionary considerations regarding why we tend to endorse, or take ourselves to have reason to endorse, claims about avoiding pain, caring for our children, cooperating with others, and punishing rule-breakers, the best way to construe the truth of such claims (if we take them to be true at all), is to understand them as reflecting something about us as rational creatures with a particular evolutionary history and articulated ideals.

Now, Tanyi (2026: Sec. V) and Sun (2026: Sec 2.2) raise several questions about my use of evolutionary debunking arguments in this context. The first issue, similarly raised by Tanyi and Sun, is that, by employing evolutionary debunking against what I call robust normative realism, I appear to endorse a form of anti-realism about reasons. At the same time, I seem to adopt a naturalist realist view according to which truths about normative reasons are reduced to truths about our responses under idealized conditions. This creates an apparent tension, since Street's (2006) debunking argument is directed at both naturalist and non-naturalist forms of realism. The question, then, is why my view does not succumb to the same evolutionary debunking considerations.

I do not think my view is under threat here, precisely because I understand my response-dependence view as anti-realist in Street's sense. On this view, truths about reasons do not transcend truths about what we, as rational agents, would agree upon regarding which reasons there are.

Another worry might arise at the meta-level: what entitles me to ground normative reasons in the idealized responses of rational agents rather than in something else? How do I know that normative reasons are in fact grounded in such idealized responses? As before, my answer would appeal to the methodological naturalism that shapes the project. The general strategy is to ask which view of normative reasons would be favored from a naturalistic perspective. Among objectivist and subjectivist accounts, I argue that naturalism coheres more readily with subjectivist views. The response-

dependence account I defend in the book represents one way of developing a subjectivist position. Although I do not explore alternatives explicitly, I leave it open that other accounts within the subjectivist framework might, in some respects, prove superior. Nevertheless, whichever such account one adopts, I do not think it would be vulnerable to evolutionary debunking, insofar as it operates within the anti-realist constraints that debunking arguments are intended to establish.

7.2. Evolutionary Debunking with Ontological Upshots

The second question raised by Tanyi (2026: 34) concerns my presentation of the evolutionary debunking argument as having an ontological upshot. In contemporary debates, such arguments are typically construed as epistemological. Specifically, Street (2006) formulates an epistemic dilemma for normative realists: either evaluative realists must accept skepticism about the truth of even our most basic evaluative judgments, or they must abandon the claim that those judgments track objective, mind-independent truths. In contrast, I suggest that evolutionary considerations not only pose a dilemma for evaluative realists but also support the endorsement of an anti-realist-friendly subjectivist account of normative reasons. If we grant that there are truths about normative reasons concerning us as human agents, then, from a naturalistic point of view, those truths are best understood as grounded in facts about us as agents.

Tanyi (2026: 34) finds this orientation toward an ontological construal of evolutionary debunking arguments puzzling for two reasons. First, he claims that if Street is right that, under a realist construal, we can never know whether our normative beliefs correspond to an external, objective reality, then it might seem pointless to go on to argue that such a reality does not exist. After all, on that view, we could never know anything about it.

However, for me the ontological issue is more central, because the project is partly ontological: it concerns the nature of normative reasons. Simply agreeing that we could not know whether realism is true does not, for my purposes, settle the question of what normative reasons are. Moreover, I find the ontological turn congenial to what Street is ultimately doing with the debunking argument. Her aim is not merely to generate epistemological skepticism, but to motivate some form of anti-realist view about evaluative truth. That is also what I find most compelling when thinking about normativity in evolutionary terms. It is not just that evolutionary considerations leave us uncertain about whether we can know the nature of normative truths; rather, they seem to favor one understanding of normative truth over another.

The second reason Tanyi (*ibid.*) finds it puzzling to construe evolutionary debunking arguments in ontological terms is that such an understanding may appear to have little impact on certain non-naturalists who are among my primary targets. In particular, Derek Parfit (2011, 2017) and Tim Scanlon (2014) seem to adopt a relaxed ontological view of normative reality. On their accounts, there is no special realm of normative facts to which true normative claims correspond; rather, normative facts are simply understood as true normative propositions (see also Jack and Khuramy 2025). Moreover, Parfit (2017) appears to regard his position as compatible with the view that normative facts supervene on descriptive facts. Tanyi therefore suggests that authors who accept such a view may find ontological versions of debunking arguments less threatening than their more familiar epistemological counterparts.

I am not sure how deflating ontology in this way makes Parfitian normative realism more palatable. It seems instead to introduce an additional air of mystery concerning what exactly such realism is committed to. In any case, the aspect of realism that I take both Street's epistemic version and my ontological uptake of the evolutionary debunking argument to target is the realist commitment to the mind-, stance-, or response-independence of normative truths. As mentioned earlier, I believe that once evolutionary and other naturalistic considerations are taken seriously, subjectivist views that reject this independence condition will appear more plausible as accounts that make sense of the nature of normative reasons.

Sun (2026: 42–43) advances a similar objection. He argues that the ontological version of the evolutionary debunking argument does not work against all forms of normative realism. My argument is that if we can provide a plausible evolutionary explanation of why we take many intuitive judgments about normative reasons to be true—such as those pertaining to the value of our lives, our children's lives, and engaging in collaborative endeavors with others—then supposing that what makes them true is a mind-independent realm of normative facts would be explanatorily and metaphysically superfluous (for the actual discussion, see Jurjako 2024: 111–6). Sun claims that this type of argument might work against normative non-naturalism, but not against all mind-independent forms of normative realism. The idea is that the charge of being metaphysically superfluous applies only if normative facts are construed as irreducible to natural facts. However, if normative facts can be reduced to natural facts, then they might still play an explanatory role in accounting for why we hold certain beliefs to be true.

I would agree with this point. However, I would also claim that if the evolutionary explanation of our normative beliefs is plausible, then normative facts should be identified with some of the facts that form part of the

evolutionary story of those beliefs. At that stage, it seems to me, we face two options: either the facts we identify with true normative propositions are mind-independent, or they are mind-dependent. My claim is that the most plausible account, consistent with the methodological naturalism expounded in the book, will be a mind-dependent one. This is because, to reiterate some of the earlier reasoning, the evolutionary story will involve accounts of how natural and historical processes have shaped our psychological, emotional, and cognitive dispositions, which eventually became associated with our deeply held normative beliefs. In effect, this is a story about ourselves and our relations with other people. Moreover, this account will include a story about the cognitive evolution of human beings, who possess unprecedented cognitive capacities and behavioral flexibility that enable them to reflect on their place in the world and to make sense of themselves from a naturalistic perspective.

These themes are discussed in greater detail in Chapters 4 and 6 of the book. Importantly, one of the questions they will face—and this is the central theme in the book—is what this “counting in favor of” relation amounts to and, if it is ever true that some fact counts in favor of something, what makes that claim true. As mentioned before, my claim is that when confronted with this question and presented with the option of choosing between objectivist and subjectivist theories (whether of reductive or non-reductive varieties), a committed methodological naturalist will find some version of mind-dependent subjectivism more plausible.

7.3. Extending Evolutionary Debunking Arguments Beyond the Moral Domain

The third, more general question Tanyi (2026: 34) raises concerns why I extend Street’s (2006) original argument to the discussion of the nature of normative reasons rather than confining it to morality. The worry is that evolutionary debunking arguments are typically understood as having their primary force in the moral domain, and perhaps less so in other normative domains. Tanyi thus asks why considerations that are usually taken to debunk narrowly construed moral beliefs can also be taken to undermine the realist interpretation of beliefs about normative reasons more generally.

I would agree that evolutionary debunking arguments in the contemporary literature—following Street and others—are typically discussed in the context of moral judgments. In the book, however, I argue that it is not straightforward to delineate what counts as the moral domain or to identify what would make moral judgments special in a way that is relevant to evolutionary explanation (for further discussion, see Sackris and Rosenberg

Larsen 2023). For that reason, it becomes difficult to isolate which specific adaptations would be tied to morality as such, rather than to other socially and personally significant practical domains. Instead, I focus on more general judgments about what we might regard as normative reasons; for example, judgments about what we care about and those supporting claims about how we ought to behave toward one another. Tanyi does not find this move fully convincing and asks for further justification of extending the debunking argument in this way.

Here it should be noted that, despite appearances, Street's (2006) original argument is not directed solely at moral judgments (see also Street 2009). Rather, her discussion is more broadly about evaluative judgments and she argues that any realist understanding of evaluative judgments that are significantly influenced by evolution faces skeptical pressure. Moreover, when she offers examples of paradigmatic evaluative judgments, she frequently appeals to judgments about reasons, understood as considerations that count in favor of something (in particular, see Street 2006: ft. 2). In this respect, her discussion already concerns normative reasons that go beyond the strictly moral domain. My discussion follows her lead here.

Furthermore, my focus in the book is on normative practical reasons. Paradigmatic moral reasons, on most standard accounts, would form a subset of these. But normative practical reasons also include prudential reasons, reasons concerning our well-being, reasons relating to what we care about, and reasons governing how we ought to act toward our children, our conspecifics, and others more generally. These are precisely the kinds of reasons that bear directly on what we are to do. As such, intuitions and judgments about them are the most plausible candidates for being shaped by evolutionary and ontogenetic forces. For that reason, I believe, extending the scrutiny to the nature of normative practical reasons more generally is supported both by Street's original discussion and by the fact that thoughts about practical reasons are exactly the sort of normative phenomena we would expect to be influenced by evolutionary processes.

7.4. What Is the Role of Affordances in Naturalistic Explanations of Normative Reasons?

Reitsma (2026: 88–89) raises several questions about the role of the notion of affordances in my overall discussion. In chapter 6 of the book, I use this notion as part of an attempt to retell the story of the phenomenology of seemingly objective, mind-independent normative authority in terms that are more compatible with a naturalistic view of human beings and their motivational and cognitive capacities. In its most general form, Reitsma's

question is whether the appeal to affordances is intended to provide a complete explanation of value: can all values that strike people as possessing an irreducible authoritative quality be explained by appealing to the idea of affordances?

The appeal to affordances in the book is not meant to explain all issues related to value or our experience of it. Its primary purpose is to explain the phenomenology of certain facts as having an authoritative feel and as inviting specific directions of action. Examples include perceiving a person who is hurt as requiring help, or perceiving one's child as calling for our protection. In this respect, the perception of affordances involves a motivational aspect, since we experience them as pulling us toward certain actions, but also a descriptive aspect, since affordances are attributed to objects and states of affairs. Importantly, both the motivational and descriptive aspects can be explained by the structure of our motivational and cognitive systems, and by how phylogenetic and ontogenetic processes have shaped them.

The notion of affordance seems suitable for this task because its primary explanatory role is to capture our perception of certain features of the world as providing opportunities for action without requiring more demanding deliberative cognitive processing (see, e.g., McClelland 2020). If some aspects of our intuitions about the nature of value are more detached from this perceptual perspective, then they will be less easily captured by the notion of an affordance.

In that respect, the more general argument against holding a realist, mind-independent view of normative reasons would come from broader naturalistic and evolutionary considerations of the kind I discussed in previous sections and in Chapter 5 of the book. That being said, I do think that the intuitive pull toward a normativist realist understanding of the authority of more abstract values—such as justice, equality, virtuous character, and so on—may be diminished if we can explain the origins of such values in terms of more basic dispositions for valuing and acting that do not require positing an extravagant form of normative authority attached to external states of affairs.

Reitsma (2026: 90) asks why this type of reasoning should not simply be dismissed as a version of the genetic fallacy. Why think that explaining the origin of an intuition is sufficient to undermine its credentials as tracking the mind-independent truth of the matter? My response is that it becomes difficult to see why more abstract notions—such as justice—should possess a special *sui generis* authority if the phenomenology of such authority can ultimately be traced back to more basic perceptions of affordances, such as perceiving a person in need and responding with reciprocal or cooperative behavior that has itself been shaped by our evolutionary past. In my view,

normative realists owe us an explanation of where this additional normative authority comes from and why we need to posit it in the first place. Even Thomas Nagel (2012: 105), whom Reitsma mentions in this context, seems to recognize that normative realism struggles to provide a satisfying answer to this question, and consequently feels forced to bite the bullet and reject Darwinian explanations of our cognitive and motivational capacities despite the broader scientific consensus.

8. Concluding remarks

The goal of this paper was to engage with the comments on my book *Normative Reasons from a Naturalistic Point of View*. The commentaries and my responses, for the most part, focused on issues concerning naturalism and on whether a coherent subject-based or response-dependence theory of normative practical reasons can be sustained in the face of objectivist views about the kinds of facts we intuitively take to characterize normative reasons.

I would like to take this opportunity to thank my interlocutors once again for their insightful and critical comments on the book. I realize that I could not address all of their remarks and criticisms in sufficient depth. Nevertheless, I hope that the responses offered here have at least clarified where our views diverge and converge, and that these exchanges provide a solid basis for further discussion of this fascinating topic.⁵

References

- Bermúdez, J. L. 2011. *Decision Theory and Rationality* (Oxford: Oxford University Press).
- Brink, D. O. 1989. *Moral Realism and the Foundations of Ethics* (Cambridge: Cambridge University Press).
- Dohrn, D. 2024. “Can quasi-cyclical preferences be rational?”, *Theoria*. <https://doi.org/10.1387/theoria.26006>
- Enoch, D. 2005. “Why idealize?”, *Ethics*, 115(4), 759–87.
- Enoch, D. 2011. *Taking Morality Seriously: A Defense of Robust Realism* (Oxford: Oxford University Press).
- Gan, N. 2024. “Arbiters of truth and existence”, *European Journal of Analytic Philosophy*, 20(1), 1–23.

⁵ I would like to thank Luca Malatesti for reading and commenting on a previous version of this paper. This work was funded by the European Union – NextGenerationEU – KOB – uniri-iz-25-275. The views and opinions expressed are those of the author only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

- Gaus, G. F. 2011. *The Order of Public Reason* (Cambridge: Cambridge University Press).
- Gaus, G. F. 2013. "On theorizing about public reason", *European Journal of Analytic Philosophy*, 9(1).
- Harman, G. H. 2012. "Naturalism in moral philosophy", in S. Nuccetelli and G. Seay (eds.), *Ethical Naturalism: Current Debates* (Cambridge; New York: Cambridge University Press), 8–23.
- Harms, W. F. 2004. *Information and Meaning in Evolutionary Processes* (New York: Cambridge University Press).
- Jack, E. and M. Khuramy. 2025. "In defense of Parfit's ontology", *Acta Analytica*, 40(4), 643–58.
- Jurjako, M. 2024. *Normative Reasons from a Naturalistic Point of View* (Rijeka: University of Rijeka, Faculty of Humanities and Social Sciences).
- Jurjako, M. 2025. "How general are Marr's levels of analysis? An assessment of Bermúdez's view", *Theoria*. <https://doi.org/10.1387/theoria.25896>
- Karimi, P. 2024. "Relaxed realism, robust realism, and the truthmaker challenge", *Ratio*, 37(2–3), 134–44.
- Kauppinen, A. 2007. "The rise and fall of experimental philosophy", *Philosophical Explorations*, 10(2), 95–118.
- Khuramy, M. and E. Schulz. 2026. "We can defend normative error theory", *Acta Analytica*, 41(1), 99–107.
- Korsgaard, C. M. 2011. "The activity of reason", in R. J. Wallace, R. Kumar, and S. Freeman (eds.), *Reasons and Recognition: Essays on the Philosophy of T. M. Scanlon* (New York: Oxford University Press), 3–22.
- Lutz, M. and J. Lenman. 2024. "Moral naturalism", in E. N. Zalta and U. Nodelman (eds.), *The Stanford Encyclopedia of Philosophy* (Stanford University). <https://plato.stanford.edu/archives/win2019/entries/naturalism-moral/>
- McClelland, T. 2020. "The mental affordance hypothesis", *Mind*, 129(514), 401–27.
- Milevski, V. 2026. "Naturalism and normativity: a critical examination of the subject-based model of reasons", *Prolegomena*, 25(1), 51–63.
- Millikan, R. G. 1995. "Pushmi-pullyu representations", *Philosophical Perspectives*, 9, 185–200.
- Nagel, T. 2012. *Mind and Cosmos: Why the Materialist Neo-Darwinian Conception of Nature Is Almost Certainly False* (New York: Oxford University Press).
- Okasha, S. 2018. *Agents and Goals in Evolution* (Oxford: Oxford University Press).
- Olson, J. 2014. *Moral Error Theory: History, Critique, Defence* (Oxford: Oxford University Press).
- Papineau, D. 2023. "Naturalism", in E. N. Zalta and U. Nodelman (eds.), *The Stanford Encyclopedia of Philosophy* (Metaphysics Research Lab, Stanford University). <https://plato.stanford.edu/archives/fall2023/entries/naturalism/>

- Parfit, D. 2011. *On What Matters*, vol. 1 and vol. 2 (Oxford: Oxford University Press).
- Parfit, D. 2017. *On What Matters*, vol. 3 (Oxford: Oxford University Press).
- Pözlner, T. 2023. *A Philosophical Perspective on Folk Moral Objectivism* (New York: Routledge).
- Pözlner, T. and J. Cole Wright. 2023. "An empirical argument against moral non-cognitivism", *Inquiry*, 66(6), 1141–69.
- Pözlner, T., L. Zijlstra, and J. Dijkstra. 2024. "Moral progress, knowledge and error: do people believe in moral objectivity?", *Philosophical Psychology*, 37(8), 2073–109.
- Rast, E. H. 2024. "Better-making properties and the objectivity of value disagreement", *European Journal of Analytic Philosophy*, 20(1), 155–79.
- Reitsma, R. 2026. "Proceduralist subjectivism and the phenomenology of practical life: a critique of Jurjako", *Prolegomena*, 25(1), 65–92.
- Roberts, D. 2005. "Does the explanatory constraint on practical reasons favour naturalism about practical reasons?", *South African Journal of Philosophy*, 24(2), 97–108.
- Sackris, D. and R. Rosenberg Larsen. 2023. "Are there "moral" judgments?", *European Journal of Analytic Philosophy*, 19(2), (S1)1–23.
- Scanlon, T. M. 2014. *Being Realistic about Reasons* (Oxford: Oxford University Press).
- Smith, M. 1994. *The Moral Problem* (Oxford: Blackwell).
- Smith, M. 1995. "Internal reasons", *Philosophy and Phenomenological Research*, 55(1), 109–31.
- Sterelny, K. and B. Fraser. 2017. "Evolution and moral realism", *The British Journal for the Philosophy of Science*, 68(4), 981–1006.
- Street, S. 2006. "A Darwinian dilemma for realist theories of value", *Philosophical Studies*, 127(1), 109–66.
- Street, S. 2009. "Evolution and the normativity of epistemic reasons", *Canadian Journal of Philosophy*, 39(S1), 213–48.
- Street, S. 2010. "What is constructivism in ethics and metaethics?", *Philosophy Compass*, 5(5), 363–84.
- Street, S. 2017. "Nothing 'really' matters, but that's not what matters", in P. Singer (ed.), *Does Anything Really Matter* (Oxford: Oxford University Press), 121–48.
- Sturgeon, N. (1985), "Moral explanations", in G. Sayre-McCord (ed.), *Essays on Moral Realism* (Cornell University Press), 229–255.
- Sun, Y. 2026. "Why endorse response-dependent theories?", *Prolegomena*, 25(1), 39–49.
- Taccolini, J. 2024. "Can we defend normative error theory?", *European Journal of Analytic Philosophy*, 20(1), 131–54.
- Tanyi, A. 2026. "Naturalism, normativity, and reasons", *Prolegomena*, 25(1), 15–38.
- Ventham, E. 2023. "Morality without categoricity", *European Journal of Analytic Philosophy*, 19(2), (A4)1–23.

Williams, B. 1981. "Internal and external reasons", in *Moral Luck* (Cambridge: Cambridge University Press), 101–13.