Int J Semiot Law (2024) 37:139–163
https://doi.org/10.1007/s11196-023-10048-5

Subtracting Reasons in Normative Domains

Federico L. G. Faroldi[1,2,3]

Accepted: 26 August 2023 / Published online: 21 October 2023

Abstract
Practical reasons can be aggregated to decide what one ought to do. This paper explores an operation that undoes aggregation: subtraction. I consider several distinctions concerning subtraction: subtracting content and subtracting strength; and subtracting one reason from one other reason or from a set of reasons. I put forward a precise understanding of subtracting the content of one reason from another, based on an operation of difference on a state-like, structured notion of content. Finally, I apply my approach to subtracting reasons to normative domains, and in particular to belief revision and norm change.

Keywords Practical reasons · Subtraction · Theory of content · Truthmakers

1 Introduction

Suppose that attempting to murder John is a reason why you ought to be punished. Suppose that murdering John is also a reason why you ought to be punished.

How many reasons are there? And how strong are they? There are two plausible principles of aggregation in this type of situation:

> DOUBLE COUNTING: The attempted murder and the murder together form a stronger reason to be punished than either of the considerations taken individually.

> SINGLE COUNTING: The attempted murder and the murder do not together form a stronger reason to get punished than either of the two considerations taken individually.

✉ Federico L. G. Faroldi
federico.faroldi@ugent.be

1 Department of Law, Università di Pavia, Pavia, Italy

2 Swiss National Science Foundation, Institute of Computer Science, University of Bern, Bern, Switzerland

3 Research Foundation-Flanders (FWO), Centre for Logic and Philosophy of Science Ghent University, Ghent, Belgium

Intuitively, DOUBLE COUNTING suggests, roughly, that, notwithstanding (some) overlap in content, the attempted murder and the murder provide two distinct reasons, at least as far as their strength is concerned; whereas SINGLE COUNTING suggests the opposite. Most people (and the criminal law, in this case) think that DOUBLE COUNTING is a mistake, at least when the content of one reason is included in the content of the other,[1]

Only some preliminary and largely informal research (collected mostly in [22]) has addressed the aggregation of practical reasons. But what happens if two reasons that are to be aggregated are partly concerned with the same subject matter? This question is fundamental, and it can be tackled before we have a theory of the weight of reasons or their strength, which will not be presupposed in what follows. In fact, it has been argued that one of the most important applications of the theory of reason, i.e. to legal theory by [26], suffers from the double counting problem (cf. [7]).

Let's now introduce two related principles concerning the subtraction of reasons. Suppose that attempting to murder John is a reason why you ought to get punished, because there is a law that punishes attempted crimes, and murder is a crime. Suppose that murdering John is also a reason why you ought to get punished, because there is a law that punishes crimes like murders.

LEFTOVER: The law punishing attempts is cancelled. What's left, i.e. the law punishing crimes, is still a reason to get punished.

NO LEFTOVER: The law punishing crimes like murders is cancelled. What's left, i.e. the law punishing attempts (if it is left!) is not a reason to get punished (under the reading that if not punished, then murder is not a crime, and therefore attempting to murder John is not an attempt to commit a crime).

Even when the content of one reason is included in the content of the other, it is not clear in which cases what is left (if anything) supports the same things.

But scarce, if any, attention has been devoted to subtracting reasons, to what we may call the problem of reason subtraction. The question has to do with the general task of subtracting one or more reasons that do not necessarily overlap. This is rather surprising, given that its ”implicit” counterparts, like forgetting, (in epistemology), desuetude and derogation (in the law), and perhaps even forgiving, are reasonably well-studied: this kind of task plays an important role in our lives.

Now, it could be that subtraction can be understood in terms of aggregation.

But subtraction seems to be much subtler than aggregation. For simplicity, let keep context fixed and just consider the issue of overlap: if reason 1 and reason 2 have some content in common, one can just take their overlap for granted and ignore it when considering the resulting reason's strength or weight. But if reason 1 and reason 2 have some common content, and one subtracts reason 2 from reason 1, and "takes away" their overlapping (which would happen if we are to give subtraction its

[1] For some recent discussions, see [18, p. 179, 13, 8, Ch. 7, 23, 23] distinguish between load-bearing reasons and derivative reasons, with only the former having separate weight. Thus, only derivative reasons overlap with each other and with load-bearing reasons.

natural set-theoretic interpretation), then there is no guarantee that what left of reason 1 is still a reason for the same thing, even before considering what the resulting weight is.

On the other hand, if reason 1 for A and reason 2 for A are disjoint, and one subtracts reason 2 from reason 1 (or vice versa), A is still supported by reason 1 (or 2), because subtracting non-overlapping sets from each other has no effect.

If reason 1 is instead completely contained within reason 2, and one subtracts reason 2 from reason 1, then presumably no reason at all should be left, unless perhaps one adopts strict dditiveiews, where the result could be a negative reason, a reason against A.

The question is whether one is first aggregating reason 1 and reason 2, obtaining a reason (1+2), and then subtracting, say, reason 2 from reason (1+2), or simply and directly subtracting reason 2 from reason 1.

We, therefore, need an account of the structure of reasons, to account for aggregation, double-counting, and subtraction. Such an account will help us decide whether subtraction can indeed be understood in terms of aggregation or not.

Together with the theoretical issues just mentioned, reasons subtraction seems particularly relevant when considering the epistemic and practical domains since epistemic and practical reasons feature prominently in belief, knowledge, and normative contexts, for example when it comes to issues like forgetting, belief change, norm change, derogation.

In this paper, I discuss how to account for reason subtraction. In §1., I distinguish different types of subtraction (conceptual vs ontological, content vs strength, individual vs set), in order to focus on a conceptual, content, and individual kind of subtraction, which is the most philosophically significant and has not previously been recognized. In §2 I discuss two theories of reason subtraction, the flat theory and the structured theory. Each has at least two flavors when it comes to support: a universal version and an existential version. Thus, in §3 and 4, I spell out some consequences of each view with regard to the logic of reasons that they generate. I show that this kind of subtraction cannot be captured by straightforward adaptations of existing techniques of belief revision and norm change on one hand, and of content subtraction on the other in §5, where I also provide some brief consideration of applying reason subtraction to forgetting and derogation.

2 Types of Subtraction

A sizeable part of contemporary normative theory considers normative reasons as a key element in the explanation or grounding of other normative concepts or even of reasons themselves as the fundamental normative notion, from which all others (such as deontic modals and evaluatives) would be derived.[2]

[2] Cf. e.g. [26, 25, 28, 32].

There are many takes on practical reasons in the literature: considerations as to why some action or state of affairs is obligatory,[3] or considerations in favor of (or against) something. What follows does not assume a particular theory of reasons.[4]

2.1 A Case for Subtracting Reasons

Why care about reason subtraction in the first place, even independently from a definite theory of reasons? There's empirical evidence that people constantly overlook subtraction and subtractive change focusing instead on addition, even when it is simpler and more beneficial to subtract [1, 19]; this seems confirmed in the literature on reasons, which focus on reason aggregation. Even though this work focuses on the modeling aspect, let me put forward some substantive, if preliminary, arguments to investigate reason subtraction.

First, there is an argument from symmetry, focusing on the strength and the polarity of a reason. While we might conjecture that strength is expressed by the absolute value of a number, polarity can be positive (a reason *for*), or negative (a reason *against*). Now, if one admits of reason aggregation, one way to understand reason subtraction is by adding a reason with the opposite polarity: one could think of $5 - 2$ as $5 + (-2)$. If t and s are reasons, $t - s$ in this informal characterization could be understood as $t + \bar{s}$, where $\bar{s}$ indicates s with the opposite polarity.

Second, there's an argument that depends more heavily on one's metaphysical account of reasons. If reasons are entities that can, themselves, be subtracted, such as facts, or propositions, then reasons can also be subtracted, *qua* facts, or propositions, etc. But why not use a theory of propositional subtraction, then? Because reasons also have a role of supporting, and are essentially related to obligations. A theory of reason subtraction has to explain how subtraction and support interact.

Third, one sees subtraction in the dynamics of our normative lives. Suppose that you have an obligation to work from Monday to Friday. Later, you are granted permission to take Tuesdays off. Your obligations shrank, for, in a sense, Tuesdays have been subtracted from your obligatory work days.[5] Thus, suppose that fulfilling your employment contract gives you a reason to work on Monday, to work on Tuesday, etc. Now your boss's amendment to your contract can be seen as subtracting the reason you had to work on Tuesday, rather than adding a reason not to work on Tuesday, and thus forcing a weighing, or balancing act to decide what you have most reason to do on Tuesday.[6]

There's a final consideration I want to gesture at. Dancy has done a lot to enrich our understanding of reasons and surrounding concepts. Two such concepts are

[3] Cf. for one prominent example [4]

[4] The literature on (practical) reasons is vast and the terminology is often inconsistent, so some amount of stipulation is necessary. [5, 6] contributed decisively to the topic. Useful recent takes include [21, 34] and [29]. The discussion will be limited to justificatory reasons (rather than motivational, for instance).

[5] For a recent take on permissive updates in the context of truthmaker semantics, see [35].

[6] This is not to say, obviously, that the usual balancing of reasons has to be excluded if one admits of reason subtraction.

those of disablers and attenuators (cf. [5, 6]). Consider that someone is in need: this is a reason for me to help them. If I am the only person around, this is not a reason to help them *per se*, but strengthens my reason to do so (that they are in need). If I learn that they are in need because they were trying to hurt someone, this might be a consideration that reduces the strength of my reason to help them. In Dancy's theory, disablers and attenuators are not reasons, but considerations that modify reasons, reducing their strength (attenuators) until they potentially the strength of the reason in question is reduced to zero (disabler). Not everyone agree with seeing disablers as a limit case of attenuators, and not everyone considers disablers and attenuators not to be reasons (cf on both points [17, pp. 141ff.]). I want to suggest that reason subtraction can reasonably contribute to this debate. As it will emerge later on in the paper, I deal with subtraction at the level of content. What is subtracted from a reason does not have to be a reason in itself (although it might, of course). *Prima facie*, this seems a good set-up to interpret disablers and attenuators, to the extent that they are not necessarily reasons. There is the question of whether what is left is still a reason, and if so, what it supports, if the same things supported by the original reasons, or less. *Prima facie*, this also seems like a good-set up to model at least part of what attenuators and disablers do. Now, for a *prima facie* disanalogy: disablers and attenuators are thought to act on the strength (Dancy) or ordering (Horty) of reasons, not on their content, like the kind of reason subtraction I explore in this paper. This is not an insurmountable problem, for at least two reasons. First, before not so long ago, not many theorists considered the content of reasons as a worthy topic of inquiry, so it might very well be that the theory of disablers and attenuators can be adapted to a theory that includes talk of content; second, it is reasonable that the strength of a reason is a (direct) function of its content: if its content gets diminished, it could be that the strength decreases. If this were the case, then reason subtraction could at least indirectly capture or model attenuators (which would e.g. subtract part of the content) and disablers (which would subtract e.g. all of the content). I do not want to suggest that a theory of reason subtraction will solve the issue of modeling disablers and attenuators, but just that, once we have a fairly well-developed theory of reason subtraction, we have novel theoretical tools to revisit that debate.

2.2 Types of Subtraction

Let's now focus in some detail on the kind of reason subtraction that we will deal with in the rest of the paper.

Ontological vs Conceptual We can think about reason subtraction ontologically or conceptually. The former understands subtracting reasons as a real change, which creates a shift in context. The latter, in contrast, understands subtraction as similar to abstracting from a certain context.[7] This paper will be concerned with conceptual subtraction. Among other considerations, it would be difficult to study reason subtraction, understood ontologically, without taking a substantive stance on whether reasons are holistic.[8],[9]

Content vs Strength Two kinds of subtraction need to be distinguished, both in order to avoid confusion and to explore their (possible) interaction. The *first* subtraction is at the level of content[10]; the *second* subtraction is at the level of strength. Naively, one might think that one could apply familiar set-theoretic techniques to content subtraction, more or less in the vicinity of set-theoretic difference, and familiar algebraic techniques to strength subtraction, more or less in the vicinity of an operation that inverts addition or aggregation. Note that subtraction at the level of strength, but not at the level of content, seems to require reasons to be reasons for the same things.

There are at least two issues with this naive picture. The first is a familiar one and stems from the well-known questions about the applicability of quantitative techniques to reason aggregation (cf. e.g. [8, Ch. 7]); I will not go in to it here. Second, even assuming that a well-defined, quantitative operation for reason aggregation (and, potentially, subtraction) is available, it is not clear how this would interact with the content of reasons.[11] In fact, in most accounts, reasons are essentially regarded as atomic. There is then a function that associates a weight with each reason, and operations can then be performed.[12] If reasons aggregation at the level of strength is possible at all, has not been subjected to sustained philosophical investigation, so we will set quantitative strength aside in this paper (or, alternatively, adopt a qualitative approach, assigning to all reasons the same strength) and focus instead on subtraction at the level of content.

[7] Abstracting a reason from its context, however, should not be confused with performing an isolation test, which has to do with considering the reason in question in a specific context, i.e., the context where it is alone. See [3] for similar ideas.

[8] Roughly, *holism* is the view that what counts as a reason (for something) in a certain context, does not count as a reason (for the same thing) in a different context. *Atomism* is the opposite view. Cf. [3].

[9] On the other hand, one might argue that one cannot consider reason subtraction only conceptually, without taking into account contextual considerations, because ignoring contextual considerations would amount to treating subtraction as context-invariant, and therefore to effectively endorsing an atomistic stance. The obvious response is that, in order to proceed with the conceptual understanding, one must abstract away from context for all reasons.

[10] In the case of belief revision, as it will be discussed in Sect. 5, one subtracts from a belief set, i.e., removes one or more sentences from a set of sentences, which has certain closure conditions. Reason subtraction works in a finer-grained way, as it is usually conceived of in terms of subtracting an individual reason from another individual reason, without the requirement that these reasons be atomic. Moreover, even when considering a set of reasons, this set presumably will not be closed under the very demanding closure conditions of a belief set.

[11] Indeed, even in one of the most extensive discussions of overlap in the literature (cf. [23]), overlap is defined not in terms of weight, not content: two reasons overlap if they do not provide separate weight.

[12] This is the case, for instance, in one of the most recent and fully developed formal account of reasons: [31].

Individual vs Set A further distinction needs to be made: whether we are subtracting a reason from a set of reasons or an individual reason from another individual reason. Moreover, it could be that a reason no longer supports one proposition, but can still support all the others in the set.[13]

The focus in this paper is on the individual case, on subtracting an individual reason from another individual reason, with the idea of subsequently generalizing this operation. Note that 'individual' is not the same as 'atomic'; as will emerge later on, an individual reason can have a structure.

3 Adding and Subtracting Content

This section offers an account of the content of reasons allowing for their subtraction, understood *conceptually*, and at the level of the *content* of *individual* reasons, in line with the distinctions discussed in the previous section.

Can we account for it by adapting existing accounts of, for example, (i) logical subtraction, or (ii) belief revision and norm change? In final section, I argue that neither is adequate: with regard to (i), the literature on logical subtraction, while (sometimes) offering a theory of content subtraction, cannot account for how what is left can continue to support normative consequences. With regard to (ii), while the (vast) literature on belief revision and norm change somewhat caters to this remaining support (but in a weak sense, to be specified), it cannot capture content subtraction, because beliefs are treated as sets of propositions closed under certain properties.

In this section, I introduce and discuss my own proposal to account for research subtraction.

I will assume that we have at our disposal a set of states equipped with a parthood relation (thus, the states are partially ordered, where a partial order is a reflexive, anti-symmetric, and transitive binary relation). I will take two theories into account: which I call the flat theory and the structured theory. According to the flat theory, a reason is a set of states. According to the structured theory, a reason is a set of states closed under parthood. Reasons are reasons *for* something. We, therefore, need to explain how sets of states support propositions. We have at least two choices: a universal and an existential interpretation. We will examine how universal and existential accounts of support interact with the flat and the structured theories.

3.1 Flat vs Structured Support

When we consider the content of reasons, for instance to avoid double counting, or in order to have a theory of subtraction, two issues immediately present themselves. First, what kind of entities reasons are. Second, what theory of content we should adopt. The two issues seem to be related. One might think that if reasons are facts,

[13] Lehmann and Thomas [20] formalizes this kind of contraction in the context of justification logic.

then what their content looks like would be different from what it would be if reasons were propositions, or mental states, and so on.

Regardless of what reasons really are, a theory of their content needs to fulfill two desiderata. First, it must have *structure*, in one way or another: it needs to be able to accommodate intuitions about reasons overlapping and having parts. Second, it needs to be sufficiently *fine-grained* to take into account the sensitivity of reasons to hyperintensional issues,[14] among other considerations.

Since the focus of this paper is on subtraction, I do not take a metaphysical stance, and only adopt an instrumentalist perspective.

Specifically, I will take a reason to be a set of states from a state space which is endowed with a parthood structure, in this case a partial order with some completeness requirements (i.e., that it is a semi-lattice, where any two states have a fusion).[15] This gets us the first requirement (structure) for free, and will also cater to the second requirement (fine-grainedness), as these structures have many hyperintensional applications.[16] It is advantageous to use states due to their generality: states are compatible with a multitude of metaphysical theses on the nature of reasons. Moreover, and still in an instrumentalist spirit, we can take states to verify (or falsify) sentences in a manner reminiscent of the old possible-world framework, but without the usual consistency and maximality assumptions. Moreover, there is at least one quite well-developed theory of content based on states ([11, 12]).[17]

More precisely, given a language with propositional letters $p \in Prop$ and reasons terms $r \in Tm$, we can define a state structure:

> STATE STRUCTURE A state structure is a tuple $(S, \sqsubseteq, V, *)$ such that S is the set of states, $\sqsubseteq$ is the parthood relation, $V : Prop \to 2^S$ is a valuation function and $* \colon Tm \times S \to 2^S$ is an assignment to reasons terms at each state, closed under different conditions depending on whether the flat or structured theory is adopted.[18]

[14] Once a relation between states and propositions, such as truthmaking, has been properly defined: cf. [9] for a hyperintensional theory of practical reasons. In brief, practical reasons are argued to be hyperintensional (i) because one consideration may be a reason for something, but not for anything that is intensionally equivalent to it; (ii) a consideration may be a reason for something, but an intensionally equivalent consideration may not be a reason for the same thing.

[15] Furthermore, a semi-lattice can be complete, where every subset has a join. This immediately makes available an operation of fission and a maximum and minimum element, whose definitions are standard from order-theory.

[16] For an overview, cf. [10].

[17] However, many of these operations go through if one uses sets of worlds, Yablo-style: cf. e.g. [36].

[18] We abbreviate $*(t, s)$ as t^{*_s}, although reference to a state in the interpretation of a reason term will be assumed to be constant and thus omitted. Both propositions and reason terms are sets of states. Two questions can be raised at this point. First, does this setup make reasons propositional, thus forcing us to abjure our neutrality regarding the ontology of reasons? And if so, second, since we now have two categories of propositional entities, is it possible to understand the notion of support (in our notation, ':') in terms of a propositional connective, such as some kind of implication? The answer to the first question is no: states need not be conceived as states of affairs, and moreover propositions would require further conditions on sets of states. The answer to the second question is a qualified yes, cf. e.g. [14] for similar thoughts.

With the basic background now in place, let's move on to explore two theories of the structure of reasons: the flat theory and the structured theory.

FLAT THEORY On the *flat theory*, a reason is taken to be a set of states without closure conditions (i.e., under parthood or under fusion).

If one imagines states to be (possible) pieces of reality, then the underlying intuition is that one's background normative theory tells us which parts of reality are normatively relevant from the perspective of a single state. This entails that, from the perspective of different states, different parts of reality are normatively relevant. Moreover, if you isolate a part of a piece of reality that is normatively relevant, it does not follow that this part, taken in isolation, is also normatively relevant (from the perspective of the same state).

STRUCTURED THEORY On the *structured theory*, a reason is taken to be a set of states closed under parthood.

Here too, one's background normative theory tells one which parts of reality are normatively relevant, from the perspective of a single state. In the structured theory, however, contrary to the flat theory, if you isolate a part of a piece of reality that is normatively relevant it does follow that such a smaller part, taken by itself, is also normatively relevant (from the perspective of the same state).

Let me advance a couple of considerations. The flat theory is flatter than the structured theory, but still more structured than a theory that just takes reasons to be sets of worlds (cf. e.g. [20]): in fact, while reasons themselves are not structured, their content (i.e., states) is structured. However, while being less structured than the structured theory, the flat theory allows for more freedom: in fact there is a higher degree of choice (or arbitrariness) in the states that constitute a reason than is allowed by closing off the reason itself under parthood. In the flat theory one can have "jumps", as it were, so that, given $s_1 \sqsubseteq s_2 \sqsubseteq s_3$, we could have a reason r with the following content: $r = \{s_1, s_3\}$. On the contrary, on the structured theory, given $s_1 \sqsubseteq s_2 \sqsubseteq s_3$, if $s_3 \in r$, then $r = \{s_1, s_2, s_3\}$.

Claims of the form: "r is a reason for ϕ" (written 'r: ϕ') are then evaluated at a state.[19] States verify propositions. There is now a choice to be made. We can evaluate "r is a reason for ϕ" existentially. Or, we can evaluate it universally.

Here is the formulation for the *flat theory*.

FLAT SUPPORT The *existential* clause states that "r is a reason for ϕ" is true at a state s just in case there is at least one state in the content of the reason r that verifies ϕ. The *universal* clause states that "r is a reason for ϕ" is true at a state s just in case the content of the reason r is a subset of the states verifying ϕ.

[19] Can this step be skipped? Yes. One would however lose the different points of view that such an account provides.

Here is the formulation for the *structured theory*.

STRUCTURED SUPPORT The *existential clause* states that “r is a reason for ϕ” is true at a state s just in case there is at least one state which is a part of r, that verifies ϕ. The *universal clause* states that “r is a reason for ϕ” is true at a state s just in case all of the states that are part of r verify ϕ.

These clauses are summarized in Table 1 (with a slight, but deliberate, misuse of notation, I am using $s' \sqsubseteq r^{*_s}$ improperly, as a remainder that r^{*_s} has a different structure depending on the background theory). If there are no states in r^{*_s}, then we assume that $s \nvDash r : \phi$, for any ϕ.

We can introduce (and we will in a bit) an operation of subtraction in the state structures as we just defined them. However, we are not guaranteed that the result of such an operation is always defined and is unique. For an analogy, consider subtraction in the natural numbers: in some case it is not defined. One has to move to a structure, such as the integers, with enough of the right elements for subtraction to be always uniquely defined. We can do something similar and define a remaindered state structure:

REMAINDERED STATE STRUCTURE A state structure is remaindered iff it is distributive and for all states s, t, with $s \sqsubseteq t$, there is a remainder r of t given s. A remainder of r of t given $s \sqsubseteq t$ is such that $r \sqcup s = t$ and r and s are disjoint. In a remaindered structure, r is unique and we denote it $t - s$. In the general case, when $s \not\sqsubseteq t$, $t - s = t - (t \sqcap s)$.[20]

Now, let's suppose that we are in a remaindered space, and therefore a remainder always exists and is unique. If we use the universal support clause, then have the following result:

Theorem 1 *If* $s : \phi$ *(and* $a \sqsubset s$*),* $s - a : \phi$

Thus, any proper, as it were, subtraction one can make from the original reason will result in the same propositions being supported. This does not capture the intended meaning of our reason subtraction, for it must be possible that subtracting some part of a reason result in it not supporting the same propositions anymore.

We have a few choices at this point. The first is not to use a remaindered space for the structure of reasons.

[20] Cf. Fine and Jago, Sect. 5.

The second choice is to endorse an existential clause for support, while remaining in a remaindered space. This would ensure that sometimes the result of the subtraction will not support the same propositions as before.

There are, of course, other, less natural, possibilities, but for the rest of this chapter I won't consider remaindered structures, because they, in a certain sense, prejudge the issue.

3.2 Flat vs Structured Subtraction

The universal and existential clauses have different consequences for subtraction. We will first consider subtraction for the flat and structured theories, and then return to these clauses.

> FLAT SUBTRACTION. In the flat theory, subtraction is defined as set-theoretic difference: given that t and s are (terms for) reasons, and a reason is a set of states without closure conditions, $t - s$ results in the (possibly empty) set of states that are in t but not in s. More precisely: $(t - s)^* = t^* - s^*$.

The flat theory holds that aggregation and subtraction of the content of reasons are no more than the union and set-theoretic difference of the sets of states which constitute their content. Note that this set-theoretic setup guarantees that subtraction is always well defined (although it might be empty). Also note that the remainder, while it always exists, is not guaranteed to be a reason. In fact, in the flat theory it is not the case that any set of states has to be assigned to a reason.

Such a theory differs considerably from more truthmaker-theoretic approaches, which, instead of the set of states, would generally consider the fusion and difference of the sum of the states that compose the reasons we wish to add or subtract. Here's one approach.

> STRUCTURED SUBTRACTION. In the structured theory, given that t and s are (terms for) reasons, $t - s$ is the unique state d, if it exists, such that all $v \sqsubseteq d$ iff $v \sqsubseteq t$ and there is no state v' (different from the empty state) that overlaps with s (i.e., such that it is not the case that $v' \sqsubseteq v$ and $v' \sqsubseteq s$). In other words, $t - s$ is the greatest-$\sqsubseteq$ part of t that does not overlap with s (if it exists). More precisely: $(t - s)^* = \{d\}$, if d exists; undefined otherwise.

This operation of subtraction may be undefined if the required unique state does not exist.[21]

Let's see concretely how these two approaches fare in the simple toy model of the red and scarlet case, before moving on to a more general case.[22] Consider the

[21] Why not use set-theoretic difference with the structured theory of content? This position is theoretically possible, but would defeat the spirit of the structured account because everything will be flattened out.

[22] Note that these toy models may be partial because generally they do not represent structures where completeness assumptions have been respected.

following fact: the ball is scarlet (i.e., bright red for the purposes of this example). Now consider the following fact: the (same) ball is red.[23] Suppose that the ball's being red is a reason why you ought to get the ball. Suppose that the ball's being scarlet is (also) a reason why you ought to get the ball.

LEFTOVER: Subtract being scarlet from being red. What is left (suppose e.g. the ball's being red, perhaps a different shade of red) is a reason to get the ball.

NO LEFTOVER: Subtract being red from being scarlet. What is left (suppose e.g. the ball's being some other bright color than red or scarlet) is not a reason to get the ball.

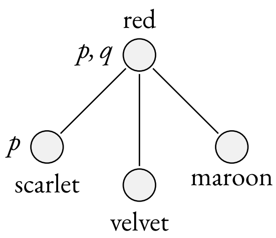

Model 1: Top-down: Red as the sum of its shades

Let's consider Model 1, where the state of the ball being red is the sum of its being scarlet, velvet, maroon. We assume existential clauses. The ball being red, *red* is a reason to get the ball: *red* : *p*. The ball being scarlet, *scarlet*, is also a reason to get the ball: *scarlet* : *p*. We have fixed a point of evaluation for simplicity. A state makes true the letter(s) next to it. Lines represent the ordering (so the state of being scarlet is part of the state of being red, and so on).

In Model 1, double counting is avoided: *red* + *scarlet* would just be *red*, so (regardless of whether we adopt the universal or existential clause), *red* is a reason for *p*, getting the ball, and that is the only reason there is. As a consequence, if we

[23] This example is inspired by [16], who, however, use it for aggregation.

Table 1 Interaction between the content of a reason and the clause for support

$s \vDash r : \phi$	$\forall$	$\exists$
Flat	$\forall s', s' \in r^{*_s}, s' \vDash \phi$	$\exists s', s' \in r^{*_s}, s' \vDash \phi$
Structured	$\forall s', s' \sqsubseteq r^{*_s}, s' \vDash \phi$	$\exists s', s' \sqsubseteq r^{*_s}, s' \vDash \phi$

also had a weight for each reason, the weight of *red* + *scarlet* would coincide with the weight of *red*.

Let's consider *red* − *scarlet*. In the *flat* case, red is taken to be just *red*, and "nothing happens": there is still a reason to get the ball, namely, *red*. So *red* : *p* = (*red* − *scarlet*) : *p*.

In the *structured* case, *red* is closed under parthood, so that it amounts to *red*, *scarlet*, *velvet*, *maroon*. What is the result of *red* − *scarlet*, and what does it support? It is still *red*, so in this case (*red* − *scarlet*) : *p* as well.

More interestingly, let's consider *scarlet* − *red*. In the flat theory, red is taken to be just *red*, so *scarlet* − *red* results in there still being a reason to get the ball, namely, *scarlet*. However, if what it is to be red is to be scarlet, velvet, and maroon, as in the structured theory, such that red amounts to *scarlet*, *velvet*, *and maroon*, then *scarlet* − *red* will result in there being no reasons to get the ball.

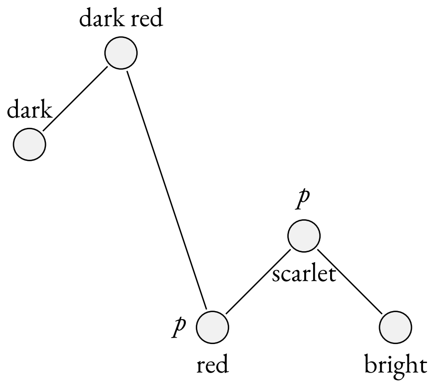

Model 2: Bottom-up

Let's now consider Model 2, where the ball being scarlet is the sum of it being red and it being bright. We shall stick with existential clauses.

In this case, double counting is also avoided: *red* + *scarlet* would consist just in *scarlet*, so *scarlet* is a reason for *p*, getting the ball.

Let's consider *red* − *scarlet*. In this case, if we adopt the flat theory and *scarlet* is taken to be just *scarlet*, then nothing happens, and there is still a reason to get the ball, namely, *red*. However, if we adopt the structured theory, *scarlet* is taken to be the fusion of *red* and *bright*, and there is no reason to get the ball.

More interestingly, let's consider *scarlet* − *red*. If *scarlet* is taken to be just *scarlet*, then there is still a reason to get the ball, namely, *scarlet*. However, if *scarlet* is taken to be the fusion of *red* and *bright*, then we are left with *bright* which, in this particular model, is not a reason to get the ball. However, were *bright* to make p true, or have a part that does, then it would be a reason to get the ball.

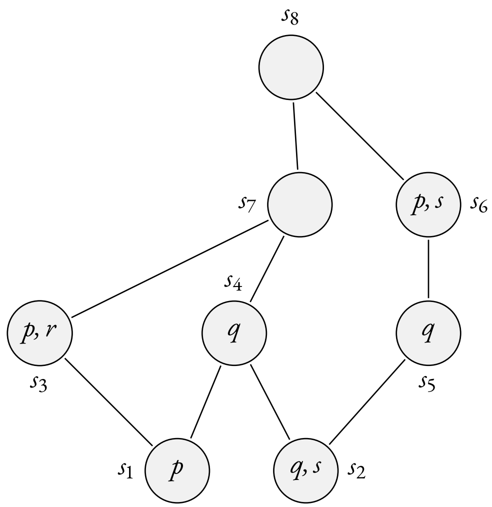

Model 3: Complex reasons

3.3 Complex Reasons

So far we have considered subtracting reasons when these reasons are, in a sense, simple: while they have a content (i.e., a set of states, in our model), they form a unity. What happens when we try to subtract complex reasons? If we go back to the logical subtraction literature, people puzzle, for instance, over the meaning of $A - (A \vee B)$. It is, of course, unclear what a connective like $\vee$ could mean in a reason-theoretic setting. In the literature on justification logic there is a operation $+$ between reason terms that is mostly interpreted disjunctively (i.e., as a union of the sets of interpreted terms, and as a disjunction of supported formulas), or, in at least one case (by [20]), conjunctively (i.e., as an intersection of the sets of interpreted terms, and as a conjunction of supported formulas).

Since our operations on contents are compositional, and as long as we have agreed on an interpretation, we can see what happens at the level of content. We can, for instance, interpret $+$ as the fusion (i.e., $\sqcup$) or fission (i.e., $\sqcap$) of states. Whether these exist, however, depends on the completeness assumptions we make about the state structure. If we assume that the fusion and fission of states always exist, then we always have a result for $+$, and subtracting complex reasons collapses into the case of subtracting simple reasons.

There is another problem, however. For the sake of concreteness, let's consider Model 3 and suppose that $A^* = \{s_1, s_2, s_3, s_4\}$ and $B^* = \{s_2, s_5, s_6\}$, where $*$ specifies the semantic assignment to reasons terms (we keep an s fixed and omit it as usual). With this example it is apparent that we cannot continue to abuse notation and casually identify a reason with its greatest state. This would still work for reason B, which has a greatest state (namely, s_6) but not for reason A, which does not have a greatest state.

This conundrum suggests the following correction to the structured theory: if the state structures are complete, we can identify the reason in question with the fusion of the states in its content, which we are guaranteed exists. In this particular case, $A = \bigsqcup A^*$ and $B = \bigsqcup B^*$. In Model 3, $A = \bigsqcup A^* = s_7$ and $B = \bigsqcup B^* = s_6$. However, while s_6 was already in B^*, s_7 was not in A^*.

This raises another question: whether we need to adapt our clauses for support. Recall the clauses in Table 1: they required that the state(s) verifying what the reason supports are in the assignment of that reason term. This creates a tension with the expanded definition we need for structured subtraction. We have several options at this point: ignore whether or not $\bigsqcup r*$ introduced new states, or take into account $\bigsqcup r*$ in our clauses for support, for instance in the following way:

(I) $r : \phi$ iff $\bigsqcup\{s | s \in r^*\} \vDash \phi$

where r^* is the interpretation of the reason (i.e., the states forming its content) and reference to points of evaluation is suppressed.

Such a clause causes the universal clauses to collapse into existential clauses, and makes for a coarser approach to the structure of reasons.

However, adjoining the fusion does not seem promising, because it introduces extraneous elements in the content of the reasons, i.e., states that are potentially irrelevant. What the "expanded" reason supports might also change, depending on what the added content does or does not support. In fact, there are no inheritance requirements, so we cannot be sure that what we add, even if it is a greater state, still supports what was supported by its parts.

For all of these reasons, we would do better accept that subtraction can sometimes be undefined.

3.4 Subtraction and Support

Now that we have a theory of content subtraction, we need to investigate how subtraction interacts with support. Let's assume that $\sigma \vDash t : \phi$. What happens if we subtract reason s from the originally supported ϕ? The answer obviously depends on whether we endorse the flat or the structured theory, and whether we adopt a universal or an existential clause for support.

We begin by formally introducing an operation on terms, $-$ defined at a state σ: $(t-s)^{*_\sigma} =_{df} t^{*_\sigma} - s^{*_\sigma}$, where $-$, in the flat theory, is simply understood as set-theoretic difference, i.e., more explicitly as $(t-s)^{*_\sigma} =_{df} \{\theta \in t^{*_\sigma} : \theta \notin s^{*_\sigma}\}$.

This simple-minded definition, for instance, implies that when t and s are disjoint, $t-s$ is just t.[24] But suppose that t and s are *not* disjoint. Then there are fewer states in t^{*_σ} left. Those that are left verify the same statements: $t : A \rightarrow (t-s) : A$, because of basic set-theoretic principles. This fact, however, is (in general) counterintuitive: if we subtract content from a reason, we don't, intuitively, have a guarantee that what is left supports the same things as before. Your making a promise is a reason for you to fulfill it, but your making a promise except that you did not actually make it yourself is not (by itself) still a reason for you to fulfill it. The structured theory gives a more nuanced account of reason subtraction.

Tables 2 and 3 summarize the findings for the flat and structured theories respectively.[25]

It now remains to consider whether subtraction is extensional or not. Since we have a theory (more than one in fact) of the structure of reasons, we can now say that two reasons have the same content, $s \equiv t$, if they are made of the same states. In our simplified theories here, this is enough to guarantee that they also support the same propositions, so that subtraction is extensional:

EXTENSIONAL-CONTENT SUBTRACTION: If $s \equiv t$, then $w - s \equiv w - t$.

However, in a more realistic version of the approach, the content of the reasons and what they support is specified separately. In this case, it could happen that reasons with the same content support different propositions, and that, at least in theory, reasons with different content support the same propositions. We denote that two reasons support the same propositions as $s \doteqdot t$. In this case, we can formulate a different version of extensionality:

EXTENSIONAL-SUPPORT SUBTRACTION: If $s \doteqdot t$, then $w - s \doteqdot w - t$.

The latter strikes me as unintuitive.

4 Discussion

We have been considering four theories of subtraction: flat-universal, flat-existential, structured-universal, and structured-existential. Is one of these clearly superior?

[24] Moreover, as a straightforward set-theoretic fact, since $A \cap B = (A - (A - B))$ (as a special case of $A - (B - C) = (A \cap C) \cup (A - B)$), we can show that an operation similar to [27]'s $\barwedge$ operation on terms ($(t \barwedge s)^{*_\sigma} \supseteq t^{*_\sigma} \cap s^{*_\sigma}$) is equivalent to this subtraction operation: $(t \barwedge s) = (t - (t - s))$, under minimal assumptions. (Moreover, [27] has a relevant conditional in the following axiom: $t : A \wedge s : B \rightarrow (t \barwedge s) : A \wedge B$).

[25] The general case of non-empty intersection, which does not degenerate in t^* and s^* coinciding or being properly included in the other, is not shown. However, it is $\sigma \vDash (t-s) : \phi$ in the universal case, and undefined in the existential case.

Flat theories of content have one advantage: they rely on set-theoretic difference, which is always defined (contrary to truthmaker difference); however an operation of subtraction has to interact with support: with the universal clause all cases are settled, while with the existential clause some cases are not.

Structured theories pay the price for the fact that, in general, truthmaker difference may not be defined.

Note that there is a difference between (1) set-theoretic difference resulting in the empty set, (2) truthmaker difference being undefined and (3) a claim like $\sigma \vDash (t - s) : \phi$ being settled.

For the first case, set-theoretic difference consisting in the empty set, this results in the state not making the subtraction true (recall that the clause for support requires that r^* be non-empty—among other conditions—for $r : \phi$ to be true at a state).

The second case, truthmaker difference being undefined, is differs to having the empty set as the result of the operation. One obviously could, if desired, set this case to the empty set in the clauses for support, but this would obscure a conceptual difference (between resulting in nothing, or at least the empty set of states, and not resulting at all).

The third case, that a claim like $\sigma \vDash (t - s) : \phi$ is settled, means that it is either true or false. In other terms $t - s$ is defined.

In the findings reported in the preceding section, we made an assumption about what t supports, namely $\sigma \vDash t : \phi$, in order to obtain certain results for $t - s$. However, we made no assumptions about s, i.e., the reason we subtracted from t. Such a move seems justifiable for at least two reasons.

First, we have been focusing on operations at the level of content, not at the level of support. What something supports depends on what its content is. When we aggregate two reasons, for instance, we do so for reasons independent of what they support. We look at what they support only after the aggregation. There is no reason why this should be different when it comes to subtraction.

Second, we deliberately left open the possibility that a reason could support both ϕ and $\neg\phi$, because this depends on the underlying truthmaking structure. This has nothing to do with reasons themselves: support is derived from truthmaking, in a sense. In the most general case, in fact, states are not required to be complete or consistent. We can of course force consistency (or completeness) on the underlying state spaces. However, this would be neither necessary nor sufficient: we should instead require that the assignment of reason terms to states be consistent (or complete).

Context In Section 2 we decided to focus on subtraction by abstracting from the role of context. Context, however, is quite important in philosophical debates on reasons. Atomism about reasons suggests (roughly) that a reason for A in a given context c remains both a reason and a reason for A in a different context c'. Holism about reasons, on the contrary, suggests (roughly) that a reason for A in a given context c may not be a reason for A in a different context c'.

There is an easy way to approximate these ideas available in the formalism we have already adopted. In fact, all assignments here, such as the * function, are relative to a state. What we formalized as r^* is more correctly written r^{*_σ}. Moreover,

Table 2 Interaction between the FLAT content of a reason and the clause for support

Ass: $\sigma \vDash t : \phi$	if $t^* = s^*$	if $t^* \cap s^* = \emptyset$	if $t^* \subset s^*$	if $s^* \subset t^*$
Flat-∀	$\sigma \nvDash (t-s) : \phi$	$\sigma \vDash (t-s) : \phi$	$\sigma \nvDash (t-s) : \phi$	$\sigma \vDash (t-s) : \phi$
Flat-∃	$\sigma \nvDash (t-s) : \phi$	$\sigma \vDash (t-s) : \phi$	$\sigma \nvDash (t-s) : \phi$	$\sigma \nvDash (t-s) : \phi \vee \neg(t-s) : \phi$

Table 3 Interaction between the STRUCTURED content of a reason and the clause for support

Ass: $\sigma \vDash t : \phi$	if $t = s$	if $t \sqcap s = \emptyset$	if $t \sqsubset s$	if $s \sqsubset t$
Structured-∀	$\sigma \nvDash (t-s) : \phi$	$\sigma \vDash (t-s) : \phi$	–	$\sigma \vDash (t-s) : \phi$
Structured-∃	$\sigma \nvDash (t-s) : \phi$	$\sigma \vDash (t-s) : \phi$	–	$\sigma \nvDash (t-s) : \phi \vee \neg(t-s) : \phi$

reasons, (i.e., claims of the form $t : \phi$) are true at states: $\sigma \vDash t : \phi$. We just put this aside in order to better focus on something else. However, if we identify a change of state as a change of context, there is a fairly straightforward way to differentiate between atomistic and holistic models, that is, by imposing (or not) further constraints on the interpretation function *:

A model is *atomistic* (or *context-invariant*) iff $\forall \sigma, \sigma' \in S, \forall r \in Tm, r^{*_\sigma} = r^{*_{\sigma'}}$.

There are multiple options to define holistic models. Here is one:

A model is *holistic* (or *context-dependent*) iff $\exists \sigma, \sigma' \in S, \exists r \in Tm, r^{*_\sigma} \neq r^{*_{\sigma'}}$.

It is interesting to note that we could, hypothetically, suppose that the way support works changes if a reason is aggregated or subtracted from another, due to (extreme) contextual considerations.[26]

Aggregation + Subtraction How do aggregation and subtraction interact? Does subtraction undo aggregation, as was our initial, intuitive suggestion? In other words, given $t : \phi$ and $(t+s) : \phi$, does $(t+s) - s : \phi$? The answer is complex, but not difficult. Subtraction undoes aggregation given certain conditions, summarized in Table 4.

Tables 4 and 5 show that assuming that $t : \phi$ makes a difference once we adopt an existential support clause, regardless of whether we adopt the flat or the structured theory.[27]

But there is another question to discuss, namely whether aggregation undoes subtraction (Table 6 and 7). If it did, we would have established that, in a certain precise sense, the aggregation and subtraction of reasons are each other's inverse operations. The question can be understood as follows: given that $t : \phi$ and $t - s : \phi$, will $(t-s) + s : \phi$ hold?

[26] I discuss these issues in a separate paper.

[27] This raises the interesting question of why subtraction seems to always be well-defined in this case.

Tables 4 and 5 show that these cases are not new, given that $t - s : \phi$. Since the subtraction is, by assumption, guaranteed to exist, whatever it is, we can treat it as a new simple term. The problem therefore reduces to the aggregation of s.

5 Arguments for a Flat Existential Theory of the Content of Reasons

We are now in a position to argue for the superiority of one theory over the others, at least on comparative grounds, if not for intrinsic reasons.

First, let's check which theory respects common intuitions about aggregation and subtraction as each other's inverse. If we assume that $t : \phi$ and $(t + s) : \phi$, $(t + s) - s : \phi$ holds in every case, but if we forgo the assumption that $t : \phi$, then $(t + s) - s : \phi$ holds only with the universal clause, regardless of whether the theory is flat or structured. We now have to check the other case, i.e. we assume $t : \phi$ and $(t - s) : \phi$, then this reduces to aggregation, so $(t - s) + s : \phi$ only with an existential support clause, regardless of whether the theory is flat or structured.

The maximally explanatory choice is therefore (1) assuming that $t : \phi$, which is a fairly neutral assumption, and (2) endorsing an existential clause of support. Of course at this point, even if we go existential, we still do not know whether to pick a flat or a structured theory.

Second, thus, let's check which theory allows for subtraction to be always defined. If we do not stipulate to use remaindered structures, which is an available option, then just because of simple set-theoretic facts we have that only flat theories allow for subtraction to be always defined (although degenerately so, i.e. as the empty set).

The clear winner, albeit a conditional one, is the flat existential theory. The conditions (or perhaps better consider them desiderata) are that we think that (1) subtraction has to undo aggregation, and that (2) subtraction should always be defined (more or less as we think of subtraction as an operation between numbers) and (3) we do not use remaindered structures.

Of course there might be philosophical motivations why the structured theory is to be preferred; this, however, would also mean to drop some of the above conditions.

In future work, it would be desirable to check whether this theoretical superiority, as it were, of the flat existential theory also extends to other phenomena not considered here, e.g. partial reasons.

Table 4 $(t + s) - s$

Ass: $t : \phi \wedge (t + s) : \phi \rightarrow$	$(t + s) - s : \phi$?
Flat-∀	Yes
Flat-∃	Yes
Structured-∀	Yes
Structured-∃	Yes

Table 5 $(t + s) - s$, fewer assumptions

Ass: $(t + s) : \phi \rightarrow$	$(t + s) - s : \phi$?
Flat-∀	Yes
Flat-∃	No
Structured-∀	Yes
Structured-∃	No

Table 6 $(t - s) + s$

Ass: $t : \phi \wedge (t - s) : \phi \rightarrow$	$(t - s) + s : \phi$?
Flat-∀	No
Flat-∃	Yes
Structured-∀	No
Structured-∃	Yes

Table 7 $(t - s) + s$, fewer assumptions

Ass: $(t - s) : \phi \rightarrow$	$(t - s) + s : \phi$?
Flat-∀	No
Flat-∃	Yes
Structured-∀	No
Structured-∃	Yes

6 Reason Subtraction, Epistemic and Practical Normativity

In this section, I consider some existing approaches to subtraction, with regard to belief subtraction, i.e., so-called belief contraction. This section provides arguments why none of these approaches is suited to subtracting reasons: either they are too flat, or they do not provide an account of how subtraction interacts with support.[28]

6.1 Can Belief Revision and Norm-change Explain Reason Subtraction?

6.1.1 Belief Revision and Contraction

Some theoretical insight might be gained by considering the field of belief revision, and in particular of contraction, the removal of a belief B from a knowledge base K (in short $K - B$) without unnecessarily removing other information. In the belief revision literature, and according to the so-called AGM model in particular, beliefs are represented by sentences in a formal language (although the overall approach is

[28] This is not to say, obviously, that reason subtraction can only be modeled with the approach developed in this work. Two thoughts are that, when suitably interpreted, Spohn's ranking theory and formal argumentation theory may prove to be good approaches to reason subtraction.

metalinguistic, cf. [2]). The set of beliefs held by an agent is assumed to be closed under logical consequence.[29]

Some considerations and comparisons are now in order. The belief revision approach to contraction does not seem to be able to be applied straightforwardly to reasons in light of some very general considerations, even before taking into account the specifics of the postulates; some ideas, however, are worth considering further,[30]

First, this approach to contraction is eminently propositional: beliefs are (represented via) sentences, closed under logical consequence, and treated extensionally (i.e., logically equivalent propositions are substitutable in a contraction context). On the contrary, practical reasons are not necessarily propositional or representable by sentences, are not closed under logical consequence, and have been argued to be hyperintensional [9, 8].

Second, this approach to contraction is eminently "flat": it does not really care about the content of single beliefs; rather it treats all beliefs as atomic, and the contraction operator is defined between a set of atomic beliefs and atomic beliefs. On the contrary, practical reasons seem to crucially require attention to their content when investigating aggregation and subtraction.[31] The entrenchment relation $\leq$ is quite promising: it could be linked to the importance or strength ordering that reasons display.[32]

[29] $K - B$, i.e., the result of removing B from K, is the largest subset of K that does not contain or imply B. Given a sentence p and a set A, the remainder set $A \perp p$ is the set of inclusion-maximal subsets of A that do not imply p. $A - p$ is then one of the remainders. Alchourrn, Grdenfors and Makinson have highlighted several problems with this construction, and have instead suggested "partial meet contraction", i.e., interpreting contraction as the intersection of some of the remainders. Concretely, let g be a selection function on K. Given that $K \perp p$ is not empty, $g(K \perp p)$ selects the "best" remainders, i.e., elements of $K \perp p$ (in case $K \perp p$ is empty, $g(K \perp p) = K$). In AGM then, $K - B = \bigcap g(K \perp p)$. Partial meet contraction satisfies Grdenfors's six basic postulates. Grdenfors also proposed that contraction of beliefs should be ruled by a binary relation, epistemic entrenchment, such that the beliefs with the lowest entrenchment should be most readily given up, because they have less epistemic or deliberative value. Given an entrenchment relation $\leq$, this gives rise to an operation of entrenchment-based contraction $\div$ according to the following definition: $q \in K \div p$ iff $q \in K$ and either $p < (p \vee q)$ or $p \in Cn(\emptyset)$, i.e., p is a theorem.

[30] Some or most of the following objections apply, *mutatis mutandis* to approaches broadly within the AGM framework, such as belief sets, which are not closed under logical consequence (cf. [15], and various dynamic logics (cf. for one [30]).

[31] The only attempt I know of to deal with belief revision with explicit reasons for belief, [24], which combines belief revision and default theory tools with justification logic, suffers from exactly these problems.

[32] There is another approach in Bayesian epistemology related to the debate about forgetting or memory loss. Traditionally, Bayesian accounts take into account (via conditionalization) the learning of new information. There is, however, an assumption (perhaps due to ideal rationality constraints) that information cannot be lost (for instance due to forgetting or memory loss), by passing from more to fewer certainties. [33] proposes a novel epistemological account based on the idea that when one loses certainties and then regains them, one ends up with the same beliefs. His approach is, however, reliant on propositional logic and flat set-theoretic techniques, and thus cannot be adapted to our present concerns.

6.1.2 Norm Change

How does subtracting reasons compare with norm change and, in particular, norm derogation? There has been some formal work on the derogation of norms, originally in the AGM framework, and later on norm change. The derogation of a norm in a normative system occurs when a norm (and, potentially, all norms implying it) is eliminated.[33] There has also been some work on forgetting norms in normative multiagent systems (Hollander and Wu 2011 and Mahmoud et al.), especially in the context of the evolution (or life cycle) of a normative system. These tools cannot be straightforwardly applied to subtracting reasons for similar reasons as the work on AGM belief revision. First, reasons are not norms, but rather they ground them (e.g. by grounding oughts), perhaps only partially and not conclusively; second, reasons are not necessarily atomic, it is possible to remove only part of their content.

6.2 Can Reason Subtraction Explain Belief Revision and Norm-change?

In the previous two sections, I maintained that (at least some) standard accounts of belief revision and norm change cannot provide an account of reason subtraction. In this section, I gesture very broadly in the opposite direction, namely, to what extent reason subtraction bears on the questions of belief revision and derogation or norm change.

6.2.1 Belief Revision and Contraction

The extent to which reason subtraction bears on the questions of belief revision depends on the theory of epistemic/doxastic reasons we adopt. Let's consider, for the argument's sake, a rough theory according to which you believe in A just in case you have sufficient reason to (e.g. warrant, evidence, etc.). Then, reason subtraction can be taken to model phenomena such as belief contraction, as the balance of reasons changes in such a way as to shrink, for example, because certain pieces of evidence are lost, discredited, forgotten, etc. The approach we have developed in the course of the paper provides a principled and precise way to identify the remainder once some reasons have been taken away, in order to know whether said remainder still epistemically supports what it used to support before the subtraction.

6.2.2 Norm Change, Derogation and Abrogation

As in the case of belief, we need to settle on a theory that links reasons and legal theory. The most prominent such theory is presumably that of [26], who, roughly,

[33] At this point there is a technical distinction to be made in the legal domain between abrogation and annulment, but this does not concern us here.

holds that rules are both first-order reasons for action, and negative second-order reasons not to act for (first-order) reasons against the former. Such negative second-order reasons are called exclusionary reasons, and they do not intervene in the balance of reasons but rather exclude reasons from the field altogether. There is a large literature on Raz's notion of exclusionary reason and its import for legal theory. One idea is that exclusionary reasons can be modeled with an operation of subtraction.

But just taking Raz's theoretical basis as a starting point, i.e. roughly that legal rules are reasons for action, we can see how reason subtraction can model both derogation and abrogation. Derogation is defined by Black's Law Dictionary as: ”the partial repeal or abolishing of a law, as by a subsequent act which limits its scope or impairs its utility and force (ad vocem)”, whereas abrogation is ”the entire repeal and annulment of a law (ib.)” With this in mind, we can now connect the dots, and model derogation as a subtraction operation whose remainder is still supporting some of the original actions; and model abrogation with a subtraction operation whose reminder is not sufficient to support what it supported before the subtraction.

This section presented some initial ideas on applying the theory of reason subtraction to the doxastic and legal domain. Obviously, much more rigorous work is needed to work out the details of such potential application.

7 Conclusion and Future Work

In this paper, I mounted an initial investigation of reason subtraction, conceived of in terms of content, understood conceptually, and focused on individual reasons. I argued that current accounts of belief revision and norm change are not adequate, because they are either propositional or atomic, or cannot express the relation of support typical of reasons.

This investigation builds on a theory of the structure of reasons and on a theory of support that can be specified in (at least) two ways: flat and structured with respect to content, and universal and existential with respect to support. Accounts of aggregation and subtraction that can solve problems like double-counting emerge naturally from this setup, and I studied their interaction. I showed that, in a specific sense, aggregation undoes subtraction and vice versa. This account allows for a formalization of the atomistic and holistic positions, at least when understood as context-relative. The aim of this paper was the introduction of a formal framework in order to understand reasons subtraction more precisely and to open the way to further investigations. Thus I have not argued for one strategy over another on philosophical grounds, although I did point out that, given certain intuitive desiderata, the flat existential theory is better than its rivals. This is only an intermediate conclusion, though, because once we have the ability to investigate the content of reasons, other phenomena (such as partial reasons or partial support) have to be studied before arriving at a final judgment.

Funding Open access funding provided by Università degli Studi di Pavia within the CRUI-CARE Agreement.

References

1. Adams, Gabrielle S., Benjamin A. Converse, Andrew H. Hales, and Leidy E. Klotz. 2021. People systematically overlook subtractive changes. *Nature* 592 (7853): 258–261.
2. Alchourrón, C.E., P. Gärdenfors, and D. Makinson. 1985. On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *Journal of Symbolic Logic* 50: 510–530.
3. Bader, Ralf. 2016. "Conditions, Modifiers, and Holism", in Lord et al. [2016b].
4. Broome, John. 2013. *Rationality Through Reasoning*. London: Wiley-Blackwell.
5. Dancy, Jonathan. 2004. *Ethics Without Principles*. Oxford: Oxford University Press.
6. Dancy, Jonathan. 2018. *Practical Shape: A Theory of Practical Reasoning*. Oxford: Oxford University Press.
7. Essert, Christopher. 2012. A Dilemma for Protected Reasons. *Law and Philosophy* 31 (1): 49–75.
8. Faroldi, Federico L. G. 2019. *Hyperintensionality and Normativity*. Dordrecht: Springer.
9. Faroldi, Federico L. G., and Tudor Protopopescu. 2019. Hyperintensional Logics of Reasons. *Logic Journal of the IGPL.* 27 (4): 411–433.
10. Faroldi, Federico L. G., and Frederik Van De Putte. 2019. *Forthcoming Kit Fine on Truthmakers*. Springer, Newyork: Relevance and Non-Classical Logic.
11. Fine, Kit. 2017. A theory of truthmaker content I: conjunction, disjunction and negation. *Journal of Philosophical Logic* 46 (6): 625–674.
12. Fine, Kit. 2017. A theory of truthmaker content II: subject-matter, common content, remainder and ground. *Journal of Philosophical Logic* 46 (6): 675–702.
13. Fogal, Daniel. 2016. "Reasons and Reason: Count and Mass", in Lord et al. [2016b].
14. Giordani, Alessandro. 2020. "Reason-based Deontic Logic", In Deontic Logic and Normative Systems 15th International Conference, DEON 2020/2021, ed. by Paul Portner Fenrong Liu Alessandra Marra and Frederik Van De Putte, College Publications.
15. Hansson, Sven Ove. 1999. *A Textbook of Belief Dynamics. Theory Change and Database Updating*. Dordrecht: Springer.
16. Hawthorne, John, and Ofra Magidor. 2018. Reflections on the Ideology of Reasons. In *The Oxford Handbook of Reasons and Normativity*, ed. Daniel Star. Oxford: Oxford University Press.
17. Horty, John F. 2012. *Reasons as Defaults*. Oxford: Oxford University Press.
18. Kearns, Stephen. 2016. "Bearing the Weight of Reasons", in Lord et al. [2016b].
19. Klotz, Leidy. 2021. *Subtract*. New York: Flatiron Books.
20. Lehmann, Eveline, and Studer Thomas. 2019. Subset models for justification logic. In *Logic, Language, Information, and Computation -WoLLIC 2019*, ed. R. Iemhoff, M. Moortgat, and R. de Queiroz, 433–449. Newyork: Springer.
21. Lord, Errol., and Barry Maguire. 2016a . "An Opinionated Guide to the Weight of Reasons", in Lord et al. [2016b].
22. Lord, Errol, and Barry Maguire, eds. 2016b. *Weighing Reasons*. Oxford: Oxford University Press.
23. Maguire, Barry, and Justin Snedegar. 2021. NormativeMetaphysics for Accountants. *Philosophical Studies*. https://doi.org/10.1007/s11098-020-01435-w.
24. Pandzic, Stipe. 2020. On the Dynamics of Structured Argumentation: Modeling Changes in Default Justification Logic. In *Foundations of Information and Knowledge Systems*, ed. Andreas Herzig and Juha Kontinen, 222–241. Cham: Springer.
25. Parfit, Derek. 2011. *On What Matters*. Oxford: Oxford University Press.
26. Raz, Joseph. 1975. *Practical Reason and Norms*. London: Hutchinson.

27. Savic, Nenad, and Thomas Studer. 2019. Relevant Justification Logic. *Journal of Applied Logic* 6 (2): 397–412.
28. Scanlon, Thomas M. 2014. *Being Realistic about Reasons*. Oxford: Oxford University Press.
29. Schroeder, Mark. 2021. *Reasons First*. Oxford: Oxford University Press.
30. Segerberg, K. 1995. Belief revision from the point of view of doxastic logic. *Logic Journal of the IGPL* 3: 535–553.
31. Sher, Itai. 2019. Comparative Value and the Weight of Reasons. *Economics and Philosophy* 35: 103–158.
32. Skorupski, John. 2010. *The Domain of Reasons*. Oxford: Oxford University Press.
33. Titelbaum, Michael G. 2012. *Quitting Certainties:ABayesian FrameworkModeling Degrees of Belief*. Oxford: Oxford University Press.
34. Weaver, Bryan R., and Scharp Kevin. 2019. *Semantics for Reasons*. Oxford: Oxford University Press.
35. Stephen, Yablo. 2018. Reply to Fine on Aboutness. *Philosophical Studies* 175 (6): 1495–1512.
36. Yablo, Stephen, and Daniel Rotschild. 2018. *Permissive Updates*. Dordrecht: Springer.