

ARTICLE

A dilemma for reasons additivity

Geoff Keeling 

Stanford University, 559 Abbot Way, Stanford, CA 94305, USA

Email: gkeeling@stanford.edu

(Received 17 February 2021; revised 17 September 2021; accepted 16 October 2021; first published online 27 December 2021)

Abstract

This paper presents a dilemma for the additive model of reasons. Either the model accommodates disjunctive cases in which one ought to perform some act ϕ just in case at least one of two factors obtains, or it accommodates conjunctive cases in which one ought to ϕ just in case both of two factors obtains. The dilemma also arises in a revised additive model that accommodates imprecisely weighted reasons. There exist disjunctive and conjunctive cases. Hence the additive model is extensionally inadequate. The upshot of the dilemma is that one of the most influential accounts of how reasons accrue to determine what we ought to do is flawed.

Keywords: Reasons; weighing; aggregation; additive; normativity

1. Introduction

The additive model of reasons:

Whether or not one ought to ϕ is determined by the weighted sum of pro tanto reasons for and against ϕ ing. The weights are given by the normative force or strength of each reason. One ought to ϕ if, and only if, the sum total weight of the reasons for ϕ ing exceeds the sum total weight of the reasons against.¹

This paper presents a dilemma for the additive model. Either the additive model accommodates *disjunctive* cases in which one ought to ϕ just in case at least one of two factors obtains, or it accommodates *conjunctive* cases in which one ought to ϕ just in case both of two factors obtain. There exist both conjunctive and disjunctive cases. Hence the additive model is extensionally inadequate. The upshot of the

¹Here ϕ is a variable that refers to an act or judgement-sensitive attitude such as believing, desiring, hoping, or fearing that something is the case (Scanlon 1998: 20). ϕ is treated as an act hereafter for brevity. A pro tanto reason is a reason which has genuine weight but which can in principle be outweighed by counter-vailing considerations (Kagan 1989: 17n). Here the term ‘reason’ is used to mean pro tanto reason throughout. Furthermore, ‘one ought to ϕ ’ is understood as having overall, all-things-considered, or most reason to ϕ . That one has overall reason to ϕ does not imply that one has an additional reason to ϕ (Dancy 2004: 16). The relevant sense of ought is assumed to be permission and not requirement throughout.

dilemma is that one of the most influential accounts of how reasons accrue to determine what we ought to do is flawed. Section 2 introduces the model. Section 3 presents the dilemma. Section 4 argues that there exists no plausible modification to the model that avoids the dilemma. Section 5 concludes.

2. The additive model

Comparativism is the view that the comparative weight of reasons for and against performing some act ϕ is the sole determinant of ϕ 's deontic status, i.e. whether or not one ought or has overall reason to ϕ (Chang 2016: 215; Tucker *Forthcoming*).² The idea is that ϕ 's deontic status is settled in a *competition* between conflicting reasons that count for and against ϕ ing to differing degrees (Broome 2004: 36–8; Berker 2007: 113–18; Lord and Maguire 2016: 4–9; Snedegar 2018). The standard metaphor is a pair of scales: reasons for ϕ are placed in one pan, reasons against are placed in the other, and one ought to ϕ if the balance of reasons favours ϕ ing.

The additive model gives a straightforward and literal analysis of the weighing metaphor.³ The weights of the reasons for and against ϕ ing add up to determine ϕ 's deontic status. One ought to ϕ , in the sense of being permitted as opposed to required to ϕ , just in case the sum total weight of the reasons for ϕ ing exceeds the sum total weight of the reasons against. The commitment here is that reasons are separable (Lord and Maguire 2016: 13–14). What it means for reasons to be separable is that the weight that each reason contributes to the combined weight of a set of reasons can be assessed independently of the weight that other reasons contribute. Reasons are separable if they do not interfere with one another. In particular, whether a feature of a situation grounds a reason for or against ϕ ing is not conditional on the obtaining or non-obtaining of other reasons that count for or against ϕ ing or on the weights of those reasons. Furthermore, the weight of any particular reason for or against ϕ ing does not depend on facts about the obtaining or non-obtaining of other reasons for or against ϕ ing or their weights.

The upshot of the additive model is twofold. On one hand, as Lord and Maguire (2016: 14) explain, 'additivity, or something like it, is attractive because it is intuitive, workable, and theoretically straightforward'. In particular, the additive model provides a plausible metaphysical basis for the phenomenology of deliberation in that it accounts for the sense in which competing considerations pull in different directions when we decide between alternatives. On the other hand, the model makes good on the use of contrast cases in normative theorizing. That is, the method of establishing the normative relevance of particular factors by permuting one factor across two cases that hold fixed all other factors and showing that there is a normative difference between the cases (Kagan 1988: 14–18). Indeed, such is the abductive standing of additivity that some claim, like

²Chang (2016: 214–15) distinguishes *reasons* and *value* comparativism. The latter holds that comparative facts about values determine deontic status. Both views originate in the 1930s, most notably with Ross (1930, 1939) but also Broad (1930: 276–85). For early discussions of comparativism see Bair (1958: Ch. 3), Nagel (1979: Ch. 7) and Scanlon (1998: 65–6).

³For some recent and closely related discussions of reasons additivity see Dietrich and List (2013), Titelbaum (2019), Nair (*Forthcoming*) and Wedgwood (*Forthcoming*). For an excellent discussion of reasons accrual that is not strictly additive see Sher (2019).

Roger Crisp, that the aim of normative theorising is that of locating the ‘reason-giving properties of which we can safely assume the “additive assumption” holds’ (Crisp 2000: 35f).

The additive model is formalized as follows. First, a situation or state of affairs can be decomposed into normatively relevant factors x_i for $i \in \{1, 2, 3, \dots, n\}$. Each factor either obtains or fails to obtain. Accordingly, each factor x_i can be treated as a Boolean such that $x_i = 1$ if factor i obtains, else $x_i = -1$. To illustrate: Suppose ϕ is the act of going to the dentist, and that the normatively relevant factors in this case are, first, whether or not one has a toothache (x_1), and second, whether or not one will have to pay for the appointment (x_2). A situation is an assignment of values to factors. For example, one situation is where one has a toothache and one does not have to pay for the appointment, i.e. $x_1 = 1$ and $x_2 = -1$. Another is where one does not have a toothache and one does not have to pay, i.e. $x_1 = -1$ and $x_2 = -1$. Accordingly, each situation can be described as a vector, or an ordered list of 1s and -1 s, which describes the normatively relevant features of the situation, i.e. $\mathbf{x} = (x_1, x_2, \dots, x_n) \in \{1, -1\}^n$. Here $\{1, -1\}^n = \{1, -1\} \times \dots \times \{1, -1\}$ is the factor space (the product set containing the 2^n possible situations). The factors are assumed to satisfy a logical independence condition such that for any x_i and x_j such that $i \neq j$, each value of x_i is compatible with each value of x_j .⁴

Each factor x_i has a real-valued weight w_i representing the polarity and weight of the reason for or against ϕ ing contributed by x_i . The sign of w_i represents the polarity of the reason. If w_i is positive, then $x_i = 1$ counts for ϕ ing, and $x_i = -1$ counts against.⁵ If w_i is negative, $x_i = 1$ counts against ϕ ing, and $x_i = -1$ counts for ϕ ing. The magnitude of w_i reflects the strength of the reason. The further w_i is from 0, the stronger the reason. Call the reasons in a case $\mathbf{w} = (w_1, w_2, \dots, w_n)$. Finally, there is some function, σ , sometimes called the combinatorial function or the governing function, that takes as its input the reasons for and against ϕ ing, and outputs whether or not one ought to ϕ (Kagan 1988: 14; Berker 2007: 121; see also Nagel 1979: 128). The additive model holds that the function is additive. The combinatorial function can be formulated with one of two thresholds.

Strict Threshold: One ought to ϕ if, and only if, the sum total weight of the reasons for ϕ ing strictly outweighs the sum total weight of the reasons against ϕ ing, i.e. $\sigma(\sum w_i x_i) = 1$ if $\sum w_i x_i > 0$, else $\sigma(\sum w_i x_i) = -1$.

⁴Factor independence is required for non-trivial separability of reasons (Bader 2016: 51–5). Factor x_1 is separable from factor x_2 just in case $(a, c) \geq (b, c) \leftrightarrow (a, d) \geq (b, d)$ for all a, b in the domain of x_1 and all c, d in the domain of x_2 . Here \geq is an ordering over the (sub-)vectors in $\{1, -1\}^n$ that ranks the (sub-)vectors in terms of the strength of the overall reason for ϕ ing given the assignment of values to factors. If dependencies exist between the factors such that certain factor values cannot be combined, the rectangular field assumption will not hold (cf. Broome 1991: 80–1). The ordering \geq may be strongly separable over factors and thus additively representable but the separability in this case will be trivial.

⁵Snedegar (2018) discusses the treatment of reasons *for* and *against* in the additive model. Balancing is zero-sum. A reason that counts to some degree for ϕ ing is taken to count against $\neg\phi$ ing to that same degree. The non-obtaining of a feature that counts for ϕ ing to some degree is taken to count as a reason in favour of $\neg\phi$ ing to that same degree. Section 4.4 shows that something close to the dilemma raised against the additive model also arises if the non-obtaining of factors is assumed to have zero as opposed to negative reason-giving forced. On this model, disjunctive and not conjunctive cases can be accommodated.

Weak Threshold: One ought to ϕ if, and only if, the sum total weight of the reasons for ϕ ing is at least as great as the sum total weight of the reasons against ϕ ing, i.e. $\sigma(\sum w_i x_i) = 1$ if $\sum w_i x_i \geq 0$, else $\sigma(\sum w_i x_i) = -1$.

This section ends by flagging that the strict commitment to additivity can be relaxed in at least three respects. First, one might restrict the scope of the model such that reasons accrue additively in a strict subset of cases. One motivation for the restriction is that certain reasons may be incomparable or incommensurable. Hence one might restrict the scope of the additive model to cases in which reasons are commensurate and comparable (Snedegar 2018: 726). Second, one might drop the commitment to precisely weighted reasons. Broome (2004: 37) rejects a picture of reasons accrual that is ‘strictly analogous to mechanical weighing’ on grounds that ‘it often seems inappropriate to associate a reason with anything so precise as a number’. Nevertheless, Broome maintains that we can ‘retain the central elements of the [weighing] analogy’, on the assumption that ‘the weight [of reasons] need not be anything so precise as a number’ (cf. Berker 2007: 114). Third, one might maintain that the additive model applies to non-derivative reasons and not to derivative reasons, i.e. reasons that can be explained in terms of more basic reasons (Crisp 2000: 36–40; Gert 2004: 77–9; Nair 2016: 64–70). The idea here is to countenance non-additive cases of reasons accrual at the level of derivative reasons but nevertheless maintain that ultimate reasons for action are additive.

3. Extensional inadequacy

There are limits to what the additive model can represent.⁶ Extensional inadequacy charges against the additive model purport to show that there are certain features of the model’s target phenomenon, i.e. reasons accrual, that it fails to represent. The first half of this section explains some representational limits on the additive model and argues that these can straightforwardly be explained away. The second half offers a dilemma that presents a more serious extensional inadequacy problem.

3.1. ‘Non-additive’ reasons accrual

The additive model holds that reasons accrue additively. Counterexamples to the additive model aim to show that there exist certain cases involving non-additive reasons accrual. These cases are designed to transgress the representational limits of the additive model. In particular, the model lacks the resources to accommodate exclusive disjunction cases in which one ought to ϕ just in case exactly one of two factors obtains. Neither can it accommodate biconditional cases in

⁶Here the focus is on cases of reasons accrual that the additive model cannot represent. There are also certain purported normative phenomena that the additive model cannot represent. For example, moral dilemmas in which one ought both to ϕ and $\neg\phi$; and supererogation, i.e. cases in which the morally best option is not morally required. Tucker (Forthcoming) defends a ‘dual scale’ weighing model which addresses these shortcomings.

which one ought to ϕ just in case both or neither of two factors obtains (Horty 2012: 61; Nair 2016: 59–60). Understanding these representational limits of the additive model and why they arise will help to contextualize the dilemma and its significance. Consider,

(Exclusive Disjunction) A bomb will destroy all life on Earth if detonated. The bomb is controlled by two switches: one in London, one in Chicago. The bomb detonates if, and only if, the switches have opposing values (London UP and Chicago DOWN, or London DOWN and Chicago UP). Right now, the switches are both UP. Unbeknownst to each other, two villains plan to detonate the bomb at the stroke of midnight. One plans to use the London switch. The other plans to use the Chicago switch. Both switches are heavily guarded. We ought to panic if London villain makes it past the guards, or if Chicago villain makes it past the guards. We ought not panic if both or neither makes it past the guards.

Here there are two factors: London villain makes it past the guards (x_1) and Chicago villain makes it past the guards (x_2). The additive model cannot explain how we ought to panic if, and only if, $x_1 = 1$ and $x_2 = -1$, or $x_1 = -1$ and $x_2 = 1$. There exist no weights w_1 and w_2 that can be applied to the factors x_1 and x_2 such that $w_1x_1 + w_2x_2 \geq 0$ if, and only if, $x_1 \neq x_2$. For suppose that there exist w_1, w_2 such that: (i) $w_1 + w_2 < 0$; (ii) $w_1 - w_2 \geq 0$; (iii) $-w_1 + w_2 \geq 0$; (iv) $-w_1 - w_2 < 0$. Add (i) and (iv): $w_1 + w_2 - w_1 - w_2 < 0$. Hence $0 < 0$. ■

Because this proof strategy features heavily in the paper, it is worth taking the time to explain how the above proof works in plain English. The additive model accommodates exclusive disjunction cases only if there exists a set of weights that gives a weighted sum of factors that comes out positive just in case exactly one of the two factors obtains. Conditions (i) through (iv) are the conditions that must be satisfied by the two weights in order for it to be true that $w_1x_1 + w_2x_2 \geq 0$ if, and only if, $x_1 = 1$ and $x_2 = -1$, or $x_1 = -1$ and $x_2 = 1$. To illustrate: Condition (i) concerns the situation in which both factors obtain, i.e. $x_1 = 1$ and $x_2 = 1$. Here the weighted sum needs to come out negative because in an exclusive disjunction case one ought not ϕ if both factors obtain. In formal terms, then, we need it to be the case that $(1 \times w_1) + (1 \times w_2) < 0$. This simplifies to $w_1 + w_1 < 0$. The same idea holds for the other conditions. For example, (ii) concerns the situation where $x_1 = 1$ and $x_2 = -1$. Here the sum needs to come out positive, as one ought to ϕ in an exclusive disjunction case if exactly one factor obtains. Hence $(1 \times w_1) + (-1 \times w_2) \geq 0$. This simplifies to $w_1 - w_2 \geq 0$. The proof strategy above is to show that the conditions that the weights must satisfy in order to represent an exclusive disjunction case are inconsistent. The conditions, if satisfied, imply a contradiction. Hence no weights exist which satisfy the four conditions, and the additive model is unable to accommodate exclusive disjunction cases.

The same problem arises for biconditionals. Take John Horty's example:

(Biconditional) Suppose I am deliberating about an afternoon run, and that both heat and rain, taken individually, function as reasons to not run; still,




Figure 1. Biconditional (LHS) and Exclusive Disjunction (RHS).

the combination of heat and rain together might function as a weaker reason to not run (say, because the heat is less onerous when there is rain). (Horty 2012: 61; see also Prakken 2005: 87–8; Nair 2016: 59–60)

There are two factors: It is hot (x_1) and it is raining (x_2). The additive model cannot explain how one ought to run if, and only if, $x_1 = 1$ and $x_2 = 1$, or $x_1 = -1$ and $x_2 = -1$. There exist no weights w_1 and w_2 that can be applied to x_1 and x_2 such that $w_1x_1 + w_2x_2 \geq 0$ if, and only if, $x_1 = x_2$. For suppose there exist w_1, w_2 such that: (i) $w_1 + w_2 \geq 0$; (ii) $w_1 - w_2 < 0$; (iii) $-w_1 + w_2 < 0$; (iv) $-w_1 - w_2 \geq 0$. Add (ii) and (iii): $w_1 - w_2 - w_1 + w_2 < 0$. Hence $0 < 0$. ■

The additive model cannot represent these cases because the Boolean functions XOR and XNOR, which correspond to the exclusive disjunction and biconditional, are not linearly separable.⁷ Figure 1 illustrates. The white dots in each diagram represent the situations in which the agent ought to ϕ . The black dots represent the situations in which agent ought not ϕ . The additive model works by separating the ‘ought’ and ‘ought not’ situations with a straight line. The equation for the line is $w_1x_1 + w_2x_2 = 0$.⁸ The ‘ought’ situations are such that $w_1x_1 + w_2x_2 \geq 0$, and the ‘ought not’ situations are such that $w_1x_1 + w_2x_2 < 0$. The separating line runs through the origin and is orthogonal to the weight vector, which is shown as $w = (1, 1)$. The problem is that XOR and XNOR are such that no straight line can separate the ‘ought’ situations and the ‘ought not’ situations. Accordingly, there are no weights $w = (w_1, w_2)$ that deliver a separating line such that one ought to ϕ just in case both or neither factor obtains, or just in case one or the other obtains.

The point can be put differently: What is being presented in the diagrams is a geometric representation of the factor space. The x -axis denotes the value of x_1

⁷This is a well-known problem in machine learning (Minsky and Papert 1969; Hertz et al. 1991: 94–7). A Boolean function f is linearly separable if, and only if, $\exists w_1, \dots, w_n, t \in \mathbb{R} : f(x_1, \dots, x_n) = 1 \Leftrightarrow \sum w_i x_i \geq t$ (Crama and Hammer 2011: 404).

⁸In cases involving n factors, the 2^n possible situations form the vertices of an n -dimensional hypercube around the origin. The hypercube is separated into ‘ought’ and ‘ought not’ situations by the hyperplane $\sum w_i x_i = 0$. The separating hyperplane is orthogonal to the weight vector because $\sum w_i x_i = w \cdot x = \|w\| \|x\| \cos(\theta)$. Here $w \cdot x$ is the dot product of $w = (w_1, w_2)$ and $x = (x_1, x_2)$, θ is the angle between the two vectors, and w and x are the lengths of the vectors. Hence $\theta = 90^\circ$ when $\sum w_i x_i = 0$ because $\arccos(0) = 90^\circ$.

and the y -axis denotes the value of x_2 . Because factors take either the values 1 or -1 , we are interested in the corners of the square, i.e. the four logically possible factor combinations, $(1, 1)$, $(1, -1)$, $(-1, 1)$ and $(-1, -1)$. What we are looking for is a pair of weights w_1 and w_2 that give a line $w_1x_1 + w_2x_2 = 0$, such that the cases in which one ought to ϕ fall on one side of the line, and the cases in which one ought not ϕ fall on the other side of the line. Whatever our choice of weights, the line will be orthogonal to the weight vector $\mathbf{w} = (w_1, w_2)$, and will intersect the midpoint $(0, 0)$. In exclusive disjunction and biconditional cases, there is no line which can be drawn through the factor space, such that the cases in which one ought to ϕ fall on one side of the line, and the cases in which one ought not ϕ fall on the other side of the line. So, trivially, there are no weights w_1 and w_2 that allow for additive representation of these cases. Hence the additive model of reasons is unable to accommodate exclusive disjunction and biconditional cases.

These cases are not fatal to the additive model. Biconditional and exclusive disjunction cases can in principle be explained away either on grounds that the reasons in the relevant cases are misconstrued, or by appeal to the distinction between ultimate and derivative reasons (Crisp 2000: 36–40; Gert 2004: 77–9; Nair 2016: 64–70). For example, additivists can argue that each villain making it past the guards is not that in virtue of which one has reason to panic. Rather, the ground of the reason for panicking is the fact that the bomb will detonate. Hence the stipulated reasons are incorrect and these examples are illegitimate cases of non-additive reasons accrual. The other option is to concede that each villain making it past the guards grounds a reason for panicking, but to argue that these reasons are merely derivative reasons. The ultimate reason to panic is grounded in the fact that we are all going to die. Additivists can then concede that whilst the exclusive disjunction case is an instance of non-additive reasons accrual, the reasons that accrue are non-ultimate, and the additive model applies only to ultimate reasons. Neither move is obviously ad hoc. Thus the representational limits of the additive model so characterized are not obviously problematic.

3.2. The dilemma

This section articulates a more serious extensional inadequacy charge against the additive model. The exclusive disjunction and biconditional cases purport to show that there exist certain non-additive cases of reasons accrual. Additivists can say in response that the reasons in the relevant cases are misidentified, or that the reasons in the cases are derivative and the additive model applies only to ultimate reasons. These responses are inadmissible against the extensional inadequacy charge raised here. What is shown is a formal constraint on the additive model's representational capabilities and not a series of cases. Furthermore, the problem is not that there are apparent cases of non-additive reasons accrual that the additive model is unable to represent. Rather, the problem is that there exist certain *additive* cases of reasons accrual that the additive model is unable to represent. Consider,

- P1 There exist disjunctive cases, i.e. cases in which one ought to ϕ just in case at least one of two factors obtains; and there exist conjunctive cases, i.e. cases in which one ought to ϕ just in case both of two factors obtain.⁹
- P2 The additive model accommodates disjunctive cases if, and only if, it fails to accommodate conjunctive cases.
- C Therefore, there exist cases that the additive model fails to accommodate.

Take each premise in turn. First, P1. There exist conjunctive and disjunctive cases of reasons accrual. On one hand, conjunctive cases are cases in which two reasons are individually insufficient for ϕ ing but where the reasons together are sufficient for ϕ ing. These cases form a non-trivial component of what the additive model is intended to explain. Shyam Nair (2016: 57–8), for example, claims that theories of how reasons accrue purport to explain how '[two] reasons together [...] can have a strength that is an increasing function of the strengths of the individual reasons'. What constitutes conjunctive cases are two reasons that have a combined weight that is sufficient for ϕ , but whose individual weights are insufficient for ϕ . On the other hand, disjunctive cases are such that two reasons are individually and jointly sufficient for ϕ . These are cases in which reasons accrue but only trivially so. Each reason alone is sufficient for ϕ ing. Prima facie, both kinds of cases form non-trivial components of the explanandum of the additive model of reasons.

There is more to be said in support of P1. In particular, we can distinguish two kinds of argument for P1. The first is to appeal to cases. If there exist plausible examples of conjunctive and disjunctive cases, then there is good reason to believe that there exist conjunctive and disjunctive cases. Obviously, appeal to cases has limited dialectical force. There is always room to quibble. For example, it might be argued that the reasons in a particular case are misdescribed, or that the reasons described are derivative and not ultimate. Hence whilst cases may provide prima facie support for P1, the case for P1 should not rest entirely on cases. On the other hand, a stronger kind of argument for P1 seeks to establish a non-trivial dialectical cost to giving up either conjunctive or disjunctive cases. The idea is to show that denying the existence of conjunctive or disjunctive cases is in some significant respect unattractive for additivists. For example, by showing that the denial of either kind of case incurs a difficult explanatory

⁹It is important to distinguish three readings of conjunctive/disjunctive cases. On the strict reading, a conjunctive/disjunctive case involves *exactly* two factors, and one ought to ϕ just in case both/at least one of those factors obtains. On a moderately permissive reading, a conjunctive/disjunctive case can involve more than two factors, which are split between *background* factors x_1, \dots, x_{n-1} , and *foreground* factors, x_{n-1} and x_n . On this reading, a case is conjunctive/disjunctive provided (i) all else is equal between the background factors, i.e. the background considerations are exactly balanced; and (ii) holding fixed (i), one ought to ϕ just in case both/at least one of x_{n-1} and x_n obtains. Last, on a maximally permissive reading of conjunctive/disjunctive cases, same as before, except it does not matter whether the background considerations are equally balanced, i.e. there exist two factors x_{n-1} and x_n such that whatever the balance of x_1 through x_{n-2} , holding fixed the background factors, one ought to ϕ just in case both/at least one of x_{n-1} and x_n obtains. P2 holds for both the strict reading and the moderately permissive reading of conjunctive and disjunctive cases. It is not intended to hold for the maximally permissive reading. The salient upshot of this clarification is that it does not matter if background considerations are implicit in the examples provided to motivate P1, so long as all else is equal between the background considerations. Thanks to an anonymous reviewer for pressing for clarity on this point.

burden. What is at issue here is pitching P1 against the negation of P1, and showing that the negation of P1 is independently implausible or at least not straightforward to defend. It is clear that neither kind of argument *decisively* establishes P1. Obviously, however, insisting that a premise is decisively established is to impose an evidential standard that exceeds what might reasonably be expected in meta-normative theorizing. So, the arguments for P1 that follow are intended to provide good reason to believe P1, but there is no expectation that the arguments will prove universally compelling.

First, some cases. For the moment bracket the question of whether the reasons in these cases are ultimate or derivative. What is minimally clear is that conjunctive and disjunctive cases are ubiquitous if we are unconcerned about the distinction between ultimate and derivative reasons. Furthermore, unlike biconditional and exclusive disjunction cases, these cases are straightforward to construct and are not in the relevant sense *contrived*. For example, here are two epistemic cases:

(Conjunction) You are meeting a secret agent for lunch in a pub. You do not know what she looks like, but you have been informed that she will be wearing a grey trenchcoat and have a Campaign for Nuclear Disarmament badge affixed to her lapel. You arrive early. There is a lady in the corner. You ought to believe that she is the agent just in case she has a grey trenchcoat and a CND badge on her lapel. Neither of the reasons is sufficiently weighty for you to believe that the lady is the agent. But together both reasons are sufficiently weighty.

(Disjunction) Your teenage son is pretty cool. One day he comes home from a party. You ought to believe that he smoked marijuana at the party if he smells strongly of marijuana, or if he is visibly high, or both. These two reasons are individually and jointly sufficient for believing that he smoked marijuana.

Epistemic normativity offers particularly clear examples of conjunctive and disjunctive cases. Conjunctive and disjunctive cases are also apparently ubiquitous in practical rationality. Here are two practical examples. Consider,

(Conjunction) You want to buy a plant that requires infrequent watering and that flowers in the summer. Your friend suggests a particular plant. You ought to get the plant just in case it requires infrequent watering and it flowers in the summer. These reasons are individually insufficient but together sufficient.

(Disjunction) You have bought every item on your shopping list besides baked beans and a purple scarf. You walk past a shop. You ought to go into the shop if they sell baked beans, purple scarves, or both. Either reason is sufficient for you to go into the shop; and both reasons together are also sufficient.

There are, at least, two respects in which these cases can be challenged. First, it can be argued that these cases are not *genuine* conjunctive and disjunctive cases because there are certain background reasons at play that have not been described. However, P2 holds both for cases involving *exactly* two factors, and

cases in which there are background considerations at play but those considerations are balanced. Accordingly, so long as all else is equal amongst the background considerations, the argument goes through; and here there is no obvious reason to suppose that the background considerations tilt the scales in favour of or against ϕ ing. Second, it can be argued that the reasons at issue in these cases are not ultimate reasons, and that the dialectical burden is to provide examples of conjunctive and disjunctive cases at the level of ultimate reasons. This move is impossible to block absent a list of ultimate reasons. Without such a list, it is always an option for additivists to deny that the reasons described are ultimate. What can be done, however, is to provide examples of conjunctive and disjunctive cases that involve reasons that are typically taken to be ultimate reasons. This move is not decisive, as the reasons at issue may turn out not to be ultimate reasons. But it is the best that can be done in terms of providing examples. Consider the following pair of moral cases:

(Conjunction) You are a witness in court. You are asked to affirm or deny that you saw the Defendant plotting against the state authorities. The fact that you saw the Defendant plotting, i.e. the affirmation would be honest, is an insufficient reason for you to make the affirmation. For it might be the case that the state is illegitimate and the Defendant had good cause for plotting against the state, i.e. making the affirmation would be unjust. Likewise, the fact that the affirmation would be just, i.e. the state is legitimate, is insufficient to make the affirmation. For it might be that you did not see the Defendant plotting against the state, in which case it would be dishonest to make the affirmation. Honesty and justice are individually insufficient but jointly sufficient reasons to affirm.

(Disjunction) You promised Grandma that you would give money to the homeless should you have ample means. You have ample means, and you pass a homeless person in need of money. Here there are two reasons to give money to the homeless person: You promised Grandma and it is the kind thing to do. Either reason is sufficient to give money, and both are together sufficient.¹⁰

These cases are unlikely to convince everyone. But in lieu of a list of ultimate reasons, there is not much more that can be done in the way of providing examples. Furthermore, given the general lack of agreement about what is and what is not an ultimate reason, it is futile to pretend that cases will receive unanimous acceptance. For that reason these cases are intended only to provide *prima facie* support to P1.

The second argument for P1 is to pitch P1 against its negation, and argue that the negation of P1 is a theoretically costly commitment for additivists. Consider the costs of rejecting disjunctive and conjunctive cases in turn. First, disjunctive cases. Recall that disjunctive cases are cases in which reasons accrue but only

¹⁰This case is loosely based on Johnson King's (2020: 190–1) promise keeping case, although Johnson King's case concerns motivating reasons.

trivially. Two reasons are individually and jointly sufficient for ϕ . Denying that *these* cases exist implies that no two reasons which are individually sufficient for ϕ can co-occur for any act ϕ . Minimally, this commitment requires an error theory for apparent cases of trivial reasons accrual. But what is also required is a plausible meta-physical account of the mechanism which prohibits the co-occurrence. For example, it might be argued that no two reasons which are individually sufficient for ϕ can co-occur *because* there is a strict priority ordering over factors such that the obtaining of a dominant factor disables dominated factors from grounding reasons to ϕ . This claim, if true, would explain why no cases exist in which two reasons which are sufficient for ϕ ing can co-occur. Obviously, however, this explanation and relevantly similar explanations are unavailable to additivists as they imply that reasons are non-separable, i.e. explanations of this sort require the grounds of some reasons to act as disabling conditions on other reasons. So, it is at best unclear what mechanism is available to additivists which, first, guarantees that no two individually sufficient reasons can co-occur, whilst at the same time preserves the commitment to the claim that reasons are separable, where this latter claim is a precondition on additive representation of reasons accrual. Accordingly, giving up on disjunctive cases incurs a significant theoretical cost. Hence rejecting P1 on grounds that disjunctive cases do not exist is *prima facie* unattractive.

On the other hand, giving up conjunctive cases neuters the abductive support for the additive model. Recall that the abductive case for the additive model depends in large part on its simplicity. Now, conjunctive cases are such that two reasons are individually insufficient but together sufficient for ϕ ing. If there are no such cases, then, *prima facie*, reasons accrual consists only in trivial cases of accrual, i.e. cases in which two reasons are individually and together sufficient for ϕ ing. But if this is all that the target phenomenon consists in, then the additive model is one of several models that adequately captures the target phenomenon. An equally plausible model is one on which one ought to ϕ if there exists a sufficient reason for ϕ ing. Insisting that the weights of reasons must be combined to determine the deontic status of acts, as the additive model does, introduces needless complexity into the picture if the only cases of reasons accrual are trivial cases. The target phenomenon is adequately captured by simpler models which do not require combining weights. Hence in denying conjunctive cases, i.e. cases of non-trivial reasons accrual, additivists are restricting the phenomenon to be explained to such an extent that their model is no longer the simplest explanation of the target phenomenon. Thus, the risk in denying the existence of conjunctive cases is that additivists undermine the abductive support for their model over rival models of how reasons accrue.

That concludes the case for P1. This case is not decisive. But what should be clear, minimally, is that there is *prima facie* evidence in the form of cases to accept P1, and that the dialectical costs of denying P1 are significant. This provides good reason to take P1 seriously, even if the case for its acceptance is not decisive.

P2. This premise is a theorem. Recall that the additive model can be formulated either as having a strict or a weak threshold. The strict threshold is such that one ought to ϕ just in case the reasons for ϕ ing *strictly outweigh* the reasons against. The weak threshold is such that one ought to ϕ just in case the reasons for ϕ ing are *at least as weighty* as the reasons against. The theorem holds that the additive model accommodates conjunctive cases only if the strict threshold is used; and it

accommodates disjunctive cases only if the weak threshold is used. The additive model accommodates conjunctive or disjunctive cases but not both kinds of case.

Proof: First, disjunctive cases. The model can represent disjunctive cases on the weak threshold. Provided $w_1 = w_2 > 0$, the model gives the correct result. For example, let $w_1 = w_2 = 1$. Then: (i) $1 + 1 \geq 0$; (ii) $1 - 1 \geq 0$; (iii) $-1 + 1 \geq 0$; and (iv) $-1 - 1 \not\geq 0$, as required. The model fails to represent disjunctive cases on the strict threshold. There exists no pair of weights w_1 and w_2 such that $w_1x_1 + w_2x_2 > 0$ if, and only if, $\neg(x_1 = x_2 = -1)$. For suppose there exist w_1 and w_2 such that: (i) $w_1 + w_2 > 0$; (ii) $w_1 - w_2 > 0$; (iii) $-w_1 + w_2 > 0$; (iv) $-w_1 - w_2 \leq 0$. Sum (ii) and (iii): $w_1 - w_2 - w_1 + w_2 > 0$. Hence $0 > 0$. Contradiction. Second, conjunctive cases. The model can represent conjunctive cases on the strict threshold provided $w_1 = w_2 > 0$, e.g. let $w_1 = w_2 = 1$. Then: (i) $1 + 1 \not\geq 0$; (ii) $1 - 1 \not\geq 0$; (iii) $-1 + 1 \not\geq 0$; and (iv) $-1 - 1 \not\geq 0$, as required. The model fails to represent conjunctive cases on the weak threshold. There exist no weights w_1 and w_2 such that $w_1x_1 + w_2x_2 \geq 0$ if, and only if, $x_1 = x_2 = 1$. For suppose there exist w_1 and w_2 such that: (i) $w_1 + w_2 \geq 0$; (ii) $w_1 - w_2 < 0$; (iii) $-w_1 + w_2 < 0$; (iv) $-w_1 - w_2 < 0$. Sum (ii) and (iii): $w_1 - w_2 - w_1 + w_2 < 0$. Hence $0 < 0$. ■

That concludes the defence of the premises. The argument is valid. Hence the conclusion follows. Either the additive model fails to accommodate conjunctive cases or it fails to accommodate disjunctive cases. Both kinds of cases exist. So, there are certain cases that the additive model fails to accommodate.

The dilemma arises because the reasons in conjunctive and disjunctive cases are identical. Two factors have positive and equal weight. Accordingly, in situations such that one factor obtains and the other does not, two equally weighted reasons pull in opposing directions. What breaks the tie is the threshold. On the weak threshold, one ought to ϕ . On the strict threshold, one ought not ϕ . In formal terms, because the two factors in conjunctive and disjunctive cases have positive equal weight, the separating line which divides the 'ought' and the 'ought not' situations is the same. This line runs through both situations such that $x_1 \neq x_2$. If the threshold is weak, situations $(-1, 1)$ and $(1, -1)$ fall on the 'ought' side of the line, i.e. the case is treated as disjunctive. If the threshold is strict, the situations $(-1, 1)$ and $(1, -1)$ fall on the 'ought not' side of the line, i.e. the case is treated as conjunctive. Figure 2 illustrates. What this implies is that the additive model can either accommodate cases in which one ought to ϕ just in case at least one of two factors obtains; or the model can accommodate cases in which one ought to ϕ just in case both of two factors obtains. Because both sorts of cases are legitimate, the additive model is extensionally inadequate. There are cases it cannot accommodate.

4. Modifications

This section examines whether the additive model can be modified to avoid the dilemma. Four modifications are considered. (1) Maintain that the strict threshold applies in conjunctive cases and the weak threshold applies in disjunctive cases. (2) Introduce a bias term into the summation of reasons that stacks the additive model in favour of or against an act's performance. (3) Drop




Figure 2. Conjunction (LHS) and Disjunction (RHS).

the commitment to reasons having precise weights. (4) Maintain that a factor's non-obtaining has reason-giving weight of zero as opposed to a weight equal to the negative weight of the reason grounded in the factor when it obtains. It is argued that each option fails.

4.1. Variable threshold

The first option is to maintain that the combinatorial function, though additive in all cases, employs a weak threshold for disjunctive cases and a strict threshold for conjunctive cases. In disjunctive cases one ought to ϕ just in case the reasons for ϕ ing are at least as weighty as the reasons against; and in conjunctive cases one ought to ϕ just in case the reasons for ϕ ing are strictly weightier than the reasons against. This move in principle provides a straightforward solution to the dilemma.

The variable threshold is costly. First, introducing a variable threshold renders the additive model inconsistent with comparativism, i.e. the thesis that the deontic status of acts is solely determined by the comparative strength of the reasons for and against those acts (Chang 2016: 215; see also Tucker [Forthcoming](#)). This is for two reasons. First, whatever determines the appropriate threshold in a particular case is a partial determinant of deontic status. Consider an act ϕ , the deontic status of which is determined by two factors that have positive and equal weight. If one factor obtains and the other does not, then one ought not ϕ on the strict threshold and one ought to ϕ on the weak threshold. Hence what determines the correct threshold is a partial determinant of deontic status. Second, whatever it is that determines the appropriate threshold is not the weights of the reasons. This is because the comparative weight of the reasons in conjunctive and disjunctive cases is the same, i.e. the two factors have positive equal weight. The determinant of the threshold must therefore be something other than comparative weight of reasons.

Rejecting comparativism is not itself a theoretical cost. The cost is incurred via the explanatory burden involved in making a non-comparativist theory of additive reasons accrual that overcomes the extensional inadequacy charge. Extensional adequacy requires that the strict threshold applies in conjunctive cases and the weak threshold applies in disjunctive cases. What is needed is an explanation for why the strict threshold applies in conjunctive cases and the weak threshold applies in disjunctive cases. What determines whether a case is conjunctive or disjunctive is facts about the deontic status of the relevant act in different

situations. Hence one cannot claim that the strict threshold applies in a given case *because* that case is a conjunctive case. For then facts about an act's deontic status would determine the appropriate threshold, and the threshold would in turn determine facts about the act's deontic status. What is unclear is how we can distinguish between conjunctive and disjunctive cases as to ensure that the weak threshold applies in the latter cases and the strict threshold applies in the former cases. This cannot be done by appeal to the status of the cases as conjunctive or disjunctive on pain of circularity.

There are two options. The first is to suppose that there is a priority ordering over reasons grounded in different features. When two equally weighted reasons pull in opposing directions, the priority ordering determines whether one ought to ϕ or $\neg\phi$. The priority ordering is a tie-breaker. If the reason for ϕ ing ranks higher than the reason for $\neg\phi$ ing, then the weak threshold applies and the case is treated as a disjunctive case. In contrast, if the reason for $\neg\phi$ ing ranks higher than the reason for ϕ ing, then the strict threshold applies and the case is treated as a conjunctive case. This option allows additivists to distinguish conjunctive and disjunctive cases without appeal to the cases *qua* conjunctive or disjunctive cases, and consequently avoids the circularity charge outlined above. The explanatory burden is manageable. What is needed is a specification of the priority ordering.

The variable threshold model cannot obviously be made to work by appeal to a priority ordering over reason-giving features. The problem is that there exist pairs of cases that involve identical reason-giving features, but where one of the cases is conjunctive and the other is disjunctive. Consider the following examples:

Case 1: You are a witness in court. You are asked to affirm or deny that you saw Defendants 1 and 2 enter the bank prior to the robbery. You ought to affirm just in case you saw Defendant 1 and Defendant 2 enter the bank.

Case 2: You are a witness in court. You are asked to affirm or deny that you saw either Defendant 1 or 2 enter the bank prior to the robbery. You ought to affirm just in case you saw at least one of Defendants 1 and 2 enter the bank.

These cases involve two reason-giving features: You saw Defendant 1 enter the bank and you saw Defendant 2 enter the bank. In one case, you ought to ϕ just in case both factors obtain. In the other case, you ought to ϕ just in case at least one of the factors obtains. Accordingly, there exist both conjunctive and disjunctive cases that involve the same reason-giving features. Thus it cannot be the case that a priority ordering over the reason-giving features determines whether the weak or the strict threshold applies. For the threshold in one of these cases is weak and in the other it is strict even though the two reason-giving features are identical.

It can be objected that these cases concern *derivative* reasons and not *ultimate* reasons, and that the scope of the additive model is restricted to ultimate reasons (Nair 2016: 64–70; Crisp 2000: 36–40). Perhaps the ultimate reason that bears on making or not making the truthful affirmation is that doing so is honest. This move has some purchase, at least in the moral case. Furthermore, it is hard to block the move absent a list of reason-giving properties that are said to be ultimate and their priority ranking. We can, however, speculate about what might be ultimate reasons.

Crisp (2000: 37–8) claims that honesty and justice are ultimate reasons. Assume he is correct. Recall the following conjunctive case from Section 3.2:

Case 3: You are a witness in court. You are asked to affirm or deny that you saw the Defendant plotting against the state authorities. The fact that you saw the Defendant plotting, i.e. the affirmation would be honest, is an insufficient reason for you to make the affirmation. For it might be the case that the state is illegitimate and the Defendant had good cause for plotting against the state, i.e. making the affirmation would be unjust. Likewise, the fact that the affirmation would be just, i.e. the state is legitimate, is insufficient to make the affirmation. For it might be that you did not see the Defendant plotting against the state, in which case it would be dishonest to make the affirmation. Honesty and justice are individually insufficient but jointly sufficient reasons to affirm.

Here is a disjunctive case involving the same morally relevant factors:

Case 4: You and your friend have the same name. There is a mix-up at the Examinations Office. You receive your friend's grade and your friend receives your grade. Your friend's true mark is very high, and your true mark is very low. Your friend is bitterly disappointed with the low mark, whereas you are delighted with the high mark. The next day, by chance, you sit next to the Chief Examiner on the bus. You get talking, and he, being a chatterbox, discloses that there has been a mix-up involving two students with the same name. You and your friend are the only students in the class with the same name, so you know that the mix-up concerns you and your friend. The Chief Examiner tells you that he has no intention of fixing the mix-up because he fears that he will be fired as Chief Examiner if he admits the mistake. That evening, you are at your friend's house. They are in tears: 'How could this possibly have happened?', they sob, 'I revised so much for that exam.' You can either inform your friend that you know about the mix-up or you can say nothing. Presumably, the reasons to inform your friend are overdetermined. It would be the just thing to do, and the honest thing to do. Both reasons are individually and jointly sufficient.

Cases like these are only so convincing. There is always room to quibble. But the dialectical burden here is on additivists to rule-out such cases in order to make a plausible case for a priority-ranking over reason giving features. What is needed is a general explanation for why there exist no two reason-giving features such that in some cases one ought to ϕ just in case both features obtain, and in other cases one ought to ϕ just in case at least one of the two features obtains. Furthermore, a general error theory is required for why apparent cases in which the same two reasons appear to feature in conjunctive and disjunctive cases are merely illusory. Explaining away counterexamples on an ad hoc basis is unsatisfactory. For the burden of proof is on additivists to show that the priority ordering renders their position extensionally adequate. The burden is not on opponents of additivity to establish that the modification is extensionally inadequate. Accordingly, whilst

the counterexamples provided for both derivative and ultimate reasons do not rule out the priority ranking strategy *decisively*, what has minimally been shown is that the use of priority ranking to underpin a variable threshold is theoretically costly.

There is a second option for making the variable threshold work. This is to say that what determines the appropriate threshold in a given situation is judgement. What role judgement is required to play here must be carefully theorized. For additivists cannot follow W.D. Ross (1930) in assigning an ineliminable role for judgement in the balancing of reasons on grounds that there exist no general principles that explain how reasons trade-off to determine deontic status (cf. Audi 1996: 106; McNaughton 1996: 445–7; Berker 2007: 110). The Rossian position is comparativist. Reasons determine deontic status. Judgement is required only to *discern* which act is favoured on the balance of reasons absent general principles. Disjunctive and conjunctive cases involve two factors with positive equal weight. In situations where one factor obtains and the other does not, two reasons with equal weight pull in opposing directions. Judgement cannot be used to discern the stronger of the two reasons because the reasons are exactly balanced. Instead, what is required is that judgement *resolves* conflicts between conflicting and equally weighted reasons (cf. Nagel 1979; Gaut 1993: 36). Here the role for judgement is as a *determinant* of deontic status, not as a *discriminant*. Judgement is a tie-breaker for determining deontic status when reasons are exactly balanced.

Treating judgement as a determinant as opposed to a discriminant of deontic status is theoretically costly. This view of judgement admits a Euthyphro problem. Suppose we are in a situation in which two factors with positive equal weight are such that one factor obtains and the other does not. Then whether we ought to ϕ depends on judgement. In particular, the case can either be treated as a disjunctive case such that we ought to ϕ ; or as a conjunctive case such that we ought not ϕ . Either our judgement about the deontic status of ϕ is principled or unprincipled. If principled, then there exist principles by which the judgement can be evaluated. What is then unclear is why judgement is required to determine deontic status. For it seems that the additive model can be combined with the judgement evaluation principle to yield a non-comparativist account of how reasons accrue that provides a determinate verdict in all cases. This renders judgement eliminable, and an account is required as to what the relevant principle is. On the other hand, if the judgement is unprincipled, then the deontic status of ϕ is arbitrary. Admitting arbitrariness in deontic status seriously diminishes the theoretical standing of the additive model. This cost is unacceptably high absent an error theory for the intuition that in no situation is the deontic status of an act arbitrary.

4.2. Biased model

The second option is to introduce a bias into the model. A bias is an extra term in the summation that stacks the model in favour of a particular deontic status. Call this term b . Then one ought to ϕ if, and only if, $\sum w_i x_i + b > 0$ (strict threshold); or one ought to ϕ if, and only if, $\sum w_i x_i + b \geq 0$ (weak threshold). A negative bias makes it the case that one ought to ϕ just in case the weighted sum of reasons is at least $|b|$ (weak) or strictly greater than $|b|$ (strict). The effect is to bias the model

against the performance of ϕ . A positive bias makes it the case that one ought to ϕ just in case the weighted sum of reasons is at least equal to $-b$ (weak) or strictly greater than $-b$ (strict). The effect is to bias the model in favour of ϕ ing. In both cases, the bias shifts the separating line away from the origin. Figure 3 illustrates. This ensures that the separating line does not intersect the problematic situations in which $x_1 \neq x_2$.

Introducing a bias term is not free. One abductive virtue of the additive model is that it gives a plausible metaphysical basis for the phenomenology of deliberation. In particular, it provides a rationale for why it appears to us that what we ought to do depends on the balance of reasons that pull in conflicting directions. Whereas the reasons for and against ϕ ing postulated by the additive model track features of the phenomenology of deliberation, there is no obvious feature of deliberation that the bias term refers to. Hence introducing the bias term burdens the additive model with a one reason too many problem. The bias is, in effect, a reason for or against ϕ ing over and above the ordinary reasons for or against ϕ ing. Introducing a bias term thus diminishes the abductive standing of the additive model.¹¹ What is needed is an account of what the bias term is and an explanation for how the bias term is a partial determinant of deontic status despite not obviously featuring in deliberation.

We can nevertheless grant the plausibility of the bias term for argument's sake. The biased model can represent conjunctive cases with either the weak or the strict threshold. However, it requires a negative bias to do so. In contrast, the biased model can represent disjunctive cases with either the weak or the strict threshold. But it requires a positive bias to do so. Hence there is no unique bias b that enables the model to accommodate conjunctive and disjunctive cases simultaneously. A different bias will have to apply in conjunctive and disjunctive cases. Accordingly, similar problems emerge for the bias term and the variable threshold model. First, additivists cannot on pain of circularity claim that the status of cases as conjunctive or disjunctive determines the appropriate bias. For what makes a case conjunctive or disjunctive is facts about deontic status in different situations. It cannot be the case that facts about deontic status determine facts about what kind of case is at issue which in turn determine the appropriate threshold which in turn determines deontic status. That is circular. Hence additivists will have to appeal to either a priority ordering over the features or individual judgement to determine the appropriate bias in particular cases, and the same problems arise as before.

4.3. Imprecise reasons

The third modification is to drop the additive model's commitment to the idea that reasons have numerically representable weights. What motivates this move is the following objection: the dilemma concerns situations in which two precisely comparable reasons are exactly balanced in favour of two alternatives. However, additivists are free to deny that the additive model's use of precisely weighted

¹¹See Tucker (Forthcoming) for a discussion of a default bias towards permissibility in the context of weighing models of reasons such as the additive model. Note that 'default' in this sense differs from Horty's (2012) account of reasons as defaults.

reasons is a substantive metaphysical commitment (Broome 2004: 37; Berker 2007: 114). This feature of the model can be treated as an idealization. Perhaps, then, the dilemma reveals a limitation of the formal model, and not a problem in the metaphysical commitments that undergird the model. If so, additivists can avoid the dilemma by dropping the commitment to precisely weighted reasons.

Admitting imprecise reasons into the additive model fails to avoid the dilemma. The dilemma arises because the factors in conjunctive and disjunctive cases are identically weighted. The two factors have positive equal weight. The problem is that in situations in which one factor or the other obtains two conflicting and equally weighty reasons pull in opposing directions. Accordingly, the reasons are balanced and thus contain insufficient information to favour ϕ ing or $\neg\phi$ ing. Imprecise reasons cannot solve the dilemma because imprecise reasons contain less information than precise reasons. Because an informational deficit is the source of the dilemma, removing information from the model will not solve the problem. This point can be strengthened by showing that the dilemma resurfaces in a revised form of the additive model that accommodates imprecisely weighted reasons.

Here is the imprecise model: rather than representing the reason grounded in factor x_i as a sharp numerical value w_i , instead represent the reason as an interval, $[a_i, b_i]$. The idea is that the weight of the reason is imprecise, but we can still give lower and upper bounds on the imprecision.¹² This representational technique facilitates both slight imprecision and severe imprecision depending on the width of the intervals used. In cases involving two factors, the factors x_1 and x_2 ground imprecise reasons $[a_1, b_1]$ and $[a_2, b_2]$ respectively. These imprecisely weighted reasons yield a set of weight vectors given by each precisification of the weights:

$$\Omega = \{(\omega_1, \omega_2) \mid \omega_1 \in [a_1, b_1] \wedge \omega_2 \in [a_2, b_2]\}$$

Each $\omega \in \Omega$ defines a separating line \mathcal{L}_ω such that \mathcal{L}_ω is orthogonal to ω . Here \mathcal{L}_ω is the line $\omega_1 x_1 + \omega_2 x_2 = 0$, where $\omega = (\omega_1, \omega_2)$. Accordingly, in place of the single separating line that we find in the standard additive model, there is a set of separating lines, $\mathcal{L}_\Omega = \{\mathcal{L}_\omega \mid \omega \in \Omega\}$. The weak and strict thresholds do not work in the imprecise model. There is no determinate point at which the reasons for ϕ ing outweigh the reasons against. There are two plausible imprecise decision rules. Each is compatible with both a weak and a strict threshold. Consider,

Restrictive Rule: One ought to ϕ if on any precisification of the reasons, the reasons for ϕ ing outweigh the reasons against.

Permissive Rule: One ought to ϕ if there exists at least one precisification of the reasons such that the reasons for ϕ ing outweigh the reasons against.

The dilemma arises for the imprecise additive model because the restrictive rule accommodates conjunctive and not disjunctive cases, whereas the permissive rule accommodates disjunctive and not conjunctive cases. Hence the imprecise model cannot accommodate both conjunctive and disjunctive cases. Below is the proof.

¹²For discussion of interval representation see Chang (2005).

- (1) **Permissive rule accommodates disjunctive cases.** An example suffices. Let x_1, x_2 ground identical imprecise reasons $[0.5, 1]$. On the permissive rule, one ought to ϕ if, and only if, $\neg(x_1 = -1 \wedge x_2 = -1)$. *Case 1:* Suppose $x_1 = x_2 = 1$. Then $\exists \omega \in \Omega : \omega_1 x_1 + \omega_2 x_2 > 0$, e.g. $\omega_1 = \omega_2 = 1$. Here $1 + 1 = 2 > 0$, so one ought to ϕ . *Case 2:* Suppose $x_1 = 1 \wedge x_2 = -1$. Then $\exists \omega \in \Omega : \omega_1 x_1 + \omega_2 x_2 > 0$, e.g. $\omega_1 = 1, \omega_2 = 0.5$. Here $1 - 0.5 = 0.5 > 0$, so one ought to ϕ . *Case 3:* Same as *Case 2*. *Case 4:* Suppose $x_1 = x_2 = -1$. Then $\forall \omega \in \Omega : \omega_1 x_1 + \omega_2 x_2 \not> 0$. The upper bound on $-\omega_1 - \omega_2 = -0.5 - 0.5 = -1 < 0$, so one ought not ϕ . ■
- (2) **Permissive rule cannot accommodate conjunctive cases.** Let x_1 ground imprecise reason $[a, b]$ and x_2 ground $[c, d]$. We assume a non-degenerate case of imprecision such that $b > a$ and $d > c$. Suppose for purposes of contradiction that the permissive rule can accommodate conjunctive cases. Then $\forall \omega \in \Omega : (\omega_1 x_2 - \omega_2 x_2 \not> 0) \wedge (-\omega_1 x_1 + \omega_2 x_2 \not> 0)$. It follows from $b > a$ and $d > c$ that $\exists \omega \in \Omega : (\omega_1 > \omega_2) \vee (\omega_2 > \omega_1)$. First, suppose $\omega_1 > \omega_2$. If $\omega_1 - \omega_2 \not> 0$, then $\omega_2 \geq \omega_1$. Contradiction. Second, suppose $\omega_2 > \omega_1$. If $-\omega_1 + \omega_2 \not> 0$, then $\omega_2 \geq \omega_1$. Contradiction. Hence $\exists \omega \in \Omega : \neg((\omega_1 x_2 - \omega_2 x_2 \not> 0) \wedge (-\omega_1 x_1 + \omega_2 x_2 \not> 0))$. Hence $\neg \forall \omega \in \Omega : (\omega_1 x_2 - \omega_2 x_2 \not> 0) \wedge (-\omega_1 x_1 + \omega_2 x_2 \not> 0)$. Contradiction. ■
- (3) **Restrictive rule accommodates conjunctive cases.** An example suffices. Let x_1, x_2 ground identical imprecise reasons $[0.5, 1]$. On the restrictive rule, one ought to ϕ if, and only if, $x_1 = x_2 = 1$. *Case 1:* Suppose that $x_1 = x_2 = 1$. Here $\forall \omega \in \Omega : \omega_1 + \omega_2 \geq 1 > 0$. This follows from the lower bound 0.5. So, one ought to ϕ . *Case 2:* Suppose that $x_1 = 1 \wedge x_2 = -1$. Then $\exists \omega \in \Omega : \omega_1 x_1 + \omega_2 x_2 \leq 0$, e.g. $\omega_1 = 0.5, \omega_2 = 1$. Here $0.5 - 1 = -0.5 \leq 0$, so one ought not ϕ . *Case 3:* Same as *Case 2*. *Case 4:* Suppose that $x_1 = x_2 = -1$. Then $\exists \omega \in \Omega : \omega_1 x_1 + \omega_2 x_2 \leq 0$, e.g. $\omega_1 = 0.5, \omega_2 = 1$. Here $0.5 - 1 = -0.5 \leq 0$, so one ought not ϕ . ■
- (4) **Restrictive rule cannot accommodate disjunctive cases.** Let x_1 ground imprecise reason $[a, b]$ and x_2 ground $[c, d]$. We assume a non-degenerate case of imprecision such that $b > a$ and $d > c$. Suppose for purposes of contradiction that the restrictive rule can accommodate disjunctive cases. Then $\forall \omega \in \Omega : (\omega_1 x_1 - \omega_2 x_2 > 0) \wedge (-\omega_1 x_1 + \omega_2 x_2 > 0)$. It follows from $b > a$ and $d > c$ that $\exists \omega \in \Omega : (\omega_1 > \omega_2) \vee (\omega_2 > \omega_1)$. First, suppose $\omega_1 > \omega_2$. If $-\omega_1 + \omega_2 > 0$, then $\omega_2 > \omega_1$. Contradiction. Second, suppose $\omega_2 > \omega_1$. If $\omega_1 - \omega_2 > 0$, then $\omega_1 > \omega_2$. Contradiction. Hence $\exists \omega \in \Omega : \neg((\omega_1 - \omega_2 > 0) \wedge (\omega_2 - \omega_1 > 0))$. It follows that $\neg \forall \omega \in \Omega : (\omega_1 x_1 - \omega_2 x_2 > 0) \wedge (-\omega_1 x_1 + \omega_2 x_2 > 0)$. Contradiction. ■

4.4. No negative weight for non-obtaining factors

The final option is to suggest that the dilemma arises because the additive model, as characterized, assigns each factor x_i a value 1 if that factor obtains, and -1

otherwise.¹³ What is being assumed here is that a factor's non-obtaining grounds an equal and opposite reason to the reason grounded when that factor obtains. Additivists can deny this assumption, and maintain that the non-obtaining of a factor has no reason-giving weight for or against the act. If so, then factors x_i ought to be assigned a value 1 if the factor obtains, and 0 otherwise. Perhaps, then, the dilemma can be avoided if the 0, 1 formalism offers a more plausible representation of the additive model and the dilemma does not manifest in the 0, 1 model.

This proposal is *prima facie* plausible. Consider again,

(Disjunction) You promised Grandma that you would give money to the homeless should you have ample means. You have ample means, and you pass a homeless person in need of money. Here there are two reasons to give money to the homeless person: You promised Grandma and it is the kind thing to do. Either reason is sufficient to give money, and both are together sufficient.

Let an act's being kind and its fulfilling a promise each have positive 10 weight in favour of the act's performance. If either or both factors apply, the total reason-giving weight in favour of the act is either 10 or 20. Given a strict threshold, one ought to perform the act just in case at least one of the two features obtain (as in either case there is strictly greater than 0 weight in favour). This shows that the additive model can accommodate disjunctive cases on the strict threshold under the assumption that a factor's non-obtaining counts 0 weight rather than -1 weight.

In response: This move is plausible to the extent that the 0, 1 formalism has greater representational power than the $-1, 1$ formalism. However, this is not so. The problem is that, whilst the 0, 1 formalism allows disjunctive cases on the strict threshold, it does not allow conjunctive cases on the strict threshold. What is more, on the weak threshold, neither disjunctive nor conjunctive cases are accounted for. Hence the 0, 1 formalism faces a relevantly analogous extensional inadequacy concern. The additive model, so construed, accommodates only disjunctive cases.

Proof: (1) The strict threshold cannot represent conjunctive cases. There exists no pair of weights w_1 and w_2 such that: (i) $w_1 + w_2 > 0$; (ii) $w_1 \leq 0$; (iii) $w_2 \leq 0$; and (iv) $0 + 0 \leq 0$. Sum (ii) and (iii) to get $w_1 + w_2 \leq 0$. This contradicts (i). (2) The weak threshold cannot accommodate disjunctive cases. There exists no pair of weights w_1 and w_2 such that: (i) $w_1 + w_2 \geq 0$; (ii) $w_1 \geq 0$; (iii) $w_2 \geq 0$; (iv) $0 + 0 < 0$. (iv) is a contradiction. (3) The weak threshold cannot accommodate conjunctive cases. (iv) is the same for disjunctive cases. Contradiction. ■

To explain: First, consider the weak threshold. For any w_1, w_2 , the separating line intersects the point $(0, 0)$. This is because $(0 \times w_1) + (0 \times w_2) = 0$. The weak threshold cannot represent conjunctive or disjunctive cases because the situation in which neither factor obtains, i.e. $(0, 0)$, is on the separating line. That is, the weak threshold classifies all situations on or above the line as 'ought situations', $(0, 0)$ is on the line in the 0, 1 model, and $(0, 0)$ is an 'ought not situation' for conjunctive

¹³I am grateful to an anonymous referee for pushing me to address this objection.




Figure 3. Disjunctive Case (LHS) and Conjunctive Case (RHS).




Figure 4. The weak threshold classifies $(0, 0)$ as an ‘ought’ situation given any separating line, and thus cannot represent conjunctive or disjunctive cases; the strict threshold leaves no available separating line for conjunctive cases.

and disjunctive cases. Hence neither can be represented on the weak threshold. Second, the strict threshold cannot represent disjunctive cases because there is no line that intersects the point $(0, 0)$ such that $(1, 1)$ is strictly above the line and the other points are on or below the line. Figure 4 gives a geometric illustration.

To summarise, the dilemma can be avoided by adopting the $0, 1$ formalism as opposed to the $-1, 1$ formalism. This allows for disjunctive cases to be represented using the strict threshold. However, the price for accommodating disjunctive cases on the strict threshold is that conjunctive cases cannot be accommodated. What is more, on the weak threshold, neither conjunctive nor disjunctive cases can be accommodated. Hence adopting the $0, 1$ formalism does not solve the extensional inadequacy problem. The $0, 1$ model cannot accommodate conjunctive cases irrespective of the threshold used. Accordingly, adopting the $0, 1$ model amounts to accepting a particular horn of the dilemma as opposed to solving the dilemma.¹⁴

5. Conclusion

This paper presented a dilemma for the additive model of reasons. The model can accommodate disjunctive cases in which one ought to ϕ just in case at least one of two factors obtains if, and only if, it cannot accommodate conjunctive cases such that one ought to ϕ just in case both of two factors obtains. Hence the model can accommodate conjunctive or disjunctive cases but not both. Furthermore, the same dilemma shows up in a revised additive model that allows for imprecisely

¹⁴Note that the problem in the $0, 1$ model can be avoided by using biases. However, the problems with this approach are made clear in section 4.2.

weighted reasons. Because conjunctive and disjunctive cases are ubiquitous, the dilemma implies that the additive model is severely extensionally inadequate. Whilst there are certain modifications available to the additive model that avoid the dilemma, these modifications are inconsistent with comparativism about reasons and incur substantial theoretical costs. Hence the additive model ought to be rejected.

Acknowledgements. I am extremely grateful to Farbod Akhlaghi, Nick Axten, Yuna Blajer de la Garza, Kathleen Creel, Nick Fortescue, Leia Hopf, Miikka Jaarte, Abby Jacques, Todd Karhu, Tim Keeling, Nannerl O. Keohane, Benjamin Lange, Chong-Ming Lim, Anne Newman, Niall Paterson, Rob Reich, Jack Ryan, and Alicia Steinmetz for comments on earlier drafts. I am also grateful to audiences at Stanford University, the University of Bristol, and Google.

References

- Audi R.** 1996. Intuitionism, pluralism, and the foundations of ethics. In *Moral Knowledge?*, eds. W. Sinnott-Armstrong and M. Timmons, 101–136. New York, NY: Oxford University Press.
- Bader R.** 2016. Conditions, modifiers, and holism. In *Weighing Reasons*, eds. E. Lord and B. Maguire, 27–55. Oxford: Oxford University Press.
- Bair K.** 1958. *The Moral Point of View: A Rational Basis for Ethics*. Ithaca, NY: Cornell University Press.
- Berker S.** 2007. Particular reasons. *Ethics* **118**, 109–139.
- Broad C.D.** 1930. *Five Types of Ethical Theory*. London: Routledge and Kegan Paul.
- Broome J.** 1991. *Weighing Goods*. Oxford: Blackwell.
- Broome J.** 2004. Reasons. In *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*, eds. R. Jay Wallace, P. Pettit, S. Scheffler and M. Smith, 28–55. Oxford: Clarendon Press.
- Chang R.** 2005. Parity, interval value, and choice. *Ethics* **115**, 331–350.
- Chang R.** 2016. Comparativism: the grounds of rational choice. In *Weighing Reasons*, eds. E. Lord and B. Maguire, 213–40. Oxford: Oxford University Press.
- Crama Y. and P.L. Hammer** 2011. *Boolean Functions: Theory, Algorithms, and Applications*. Cambridge: Cambridge University Press.
- Crisp R.** 2000. Particularizing particularism. In *Moral Particularism*, eds. B. Hooker and M.O. Little, 23–47. Oxford: Oxford University Press.
- Dancy J.** 2004. *Ethics Without Principles*. New York, NY: Oxford University Press.
- Dietrich F. and C. List** A reason-based theory of rational choice. *Noûs* **47**, 104–134.
- Gaut B.** 1993. Moral pluralism. *Philosophical Papers* **22**, 17–40.
- Gert J.** 2004. *Brute Rationality: Normativity and Human Action*. Oxford: Oxford University Press.
- Hertz J., A. Krogh and R.G. Palmer** 1991. *Introduction to the Theory of Neural Computation*. Redwood City, CA: Addison-Wesley.
- Horty J.** 2012. *Reasons as Defaults*. Oxford: Oxford University Press.
- Johnson King Z.** 2020. Accidentally doing the right thing. *Philosophy and Phenomenological Research* **100**, 186–206.
- Kagan S.** 1988. The additive fallacy. *Ethics* **99**, 5–31.
- Kagan S.** 1989. *The Limits of Morality*. Oxford: Clarendon Press.
- Lord E. and B. Maguire**, eds. 2016. An opinionated guide to the weight of reasons. In *Weighing Reasons*, 3–24. Oxford: Oxford University Press.
- McNaughton D.** 1996. An unconnected heap of duties? *Philosophical Quarterly* **46**, 433–447.
- Minsky M.L. and S.A. Papert** 1969. *Perceptrons*. Cambridge, MA: MIT Press.
- Nagel T.**, ed. 1979. Fragmentation of value. In *Mortal Questions*, 128–41. Cambridge: Cambridge University Press.
- Nair S.** 2016. How do reasons accrue? In *Weighing Reasons*, eds. E. Lord and B. Maguire, 56–73. Oxford: Oxford University Press.
- Nair S.** Forthcoming. “Adding up” reasons: lessons for reductive and non-reductive approaches. *Ethics*.

- Prakken H.** 2005. A study of accrual of arguments. *Proceedings of the Tenth International Conference on Artificial Intelligence and Law*, 85–94. New York, NY: ACM Press.
- Ross W.D.** 1930. *The Right and the Good*. Oxford: Clarendon Press.
- Ross W.D.** 1939. *The Foundations of Ethics*. Oxford: Clarendon Press.
- Scanlon T.** 1998. *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Sher I.** 2019. Comparative value and the weight of reasons. *Economics and Philosophy* **35**, 103–158.
- Snedegar J.** 2018. Reasons for and reasons against. *Philosophical Studies* **173**, 725–743.
- Titelbaum M.** 2019. Reasons without reasons for. *Oxford Studies in Metaethics* **14**, 189–215.
- Tucker C.** Forthcoming. The dual scale model of weighing reasons. *Notus*.
- Wedgwood R.** Forthcoming. The reasons aggregation theorem. In *Oxford Studies in Normative Ethics*. Vol. 11, ed. M.C. Timmons. Oxford: Oxford University Press.

Geoff Keeling is a Bioethicist at Google and an Affiliate Fellow at the Institute for Human-Centered Artificial Intelligence at Stanford University. From 2020–2021, he was an Interdisciplinary Ethics Fellow at Stanford’s McCoy Family Center for Ethics in Society and Institute for Human-Centered Artificial Intelligence, during which time this work was completed. His research interests include decision theory, moral philosophy, and the ethics of artificial intelligence. URL: <https://hai.stanford.edu/people/geoff-keeling/>