Weighing Reasons Against

Chris Tucker April 2021
William & Mary cstucker@wm.edu

Forthcoming 2022 in *Oxford Studies in Metaethics* 17

0. Introduction

The idea of weighing reasons brings to mind a single (double pan balance) scale. The reasons for φ ($R_φ$) go on one pan and the reasons for ~φ ($R_{\sim φ}$) go on the other. The relative weights, as indicated by the relative heights of the two sides of the scale, determine the deontic status of the act. This model, *Single Scale*, assumes that a reason has justifying weight for φ (pushes φ down toward permissibility) only to the extent that it has requiring weight for φ (pushes ~φ up toward impermissibility). Thus, Single Scale is committed to *Single Proportion*: all reasons have the same proportion of justifying and requiring weight. Values that always come in the same proportion are *dependent* variables: if you know the value of one, then you can calculate the value of the other.

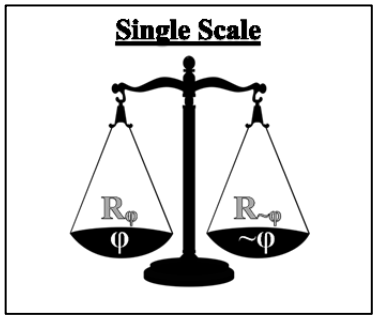

To reject Single Proportion is to deny that reasons always have the same proportion of justifying and requiring weight. For example, a standard account of supererogation holds the altruistic reasons have one proportion (equal justifying and requiring weight) whereas self-interested reasons have a different proportion (have justifying weight but no requiring weight). When values don't always come in the same proportion, they are *in*dependent variables: you must track them independently insofar as you can't calculate the value of one from the value of the other. If justifying and requiring weight are independent variables, then Single Proportion is false.

Many normative theories hold that justifying and requiring weight are *in*dependent variables. Single Scale cannot represent such theories because of its commitment to Single Proportion. My model, *Dual Scale*, can. Because it relies on *two* scales, we'll see that it has no problem representing justifying and requiring weight as independent (or dependent) variables.

Ethicists reject, for a variety of reasons, the scale as a useful metaphor for weighing reasons. Nonetheless, such ethicists generally retain the metaphor of a reason's *weight*. This combination is incoherent. The metaphor of weight is committed to this minimal idea: as φ is pushed toward permissibility, it is pushed away from impermissibility. I will show that this minimal idea entails Dual Scale. Fans of the weight (force, pressure) metaphor should be fans of Dual Scale.[1]

Justin Snedegar (2018) offers a specific challenge to scale-based models. He observes that they are committed to:

Equivalence: Insofar as R is a reason *against* φ, R is a reason *for* some alternative.

[1] An object's weight while on the φ-pan is just the downward force (pressure) it exerts on the φ-pan and the upward force (pressure) it exerts on the ~φ-pan. Consequently, Dual Scale can easily be recast in terms of force, pressure, vector sums, etc. if that is your preferred terminology.

Take Single Scale, for example. A reason *against* φ is represented as a reason *for* ~φ. Snedegar rejects Equivalence, and so he infers that "the balancing [scale] metaphor is not even approximately correct" (742). If he's right, then Dual Scale is false.

At first glance, his rejection of Equivalence seems right. There are two children in the burning building. Since you can easily carry both children at the same time, the cost to save the children is the same: you will be severely burned. You can save one child (Save1) or save both children (Save2). The severe burns involved in Save2 seem to be a reason *against* Save2. Yet they don't seem to be a reason *for* Save1. After all, you would get those very same burns if you Save1. As plausible as such examples seem to be, Equivalence is a *conceptual* truth.[2] We will see that every reason against φ, by definition, plays the same functional role in fixing deontic status as some reason for ~φ.

In §1, I present Dual Scale and explain how it resolves the problem with Single Scale. In §2, I clarify the reason for/against distinctions. Yes, you read that right. There are *two* reason for/against distinctions, because a reason can have justifying weight for/against, and it can have requiring weight for/against. In §3, I show how easily Dual Scale accommodates the various reasons for and against. In §4, I establish that Equivalence is a conceptual truth. In §5, I show that the metaphor of a reason's weight entails Dual Scale.

The version of Dual Scale I present in §1 makes the simplifying assumption that there are exactly two options, often represented by φ and ~φ. In §6, I consider the possibility that there might be more than one alternative to φ. This possibility raises:

> **The Which Alternatives Question**: if R is a reason against φ and every reason against φ is a reason for some alternative (as Equivalence assumes), then *which* alternatives is R a reason for?

Snedegar doesn't think that scale-based views can provide a satisfactory answer to this question. He's wrong. I will say a little about how to generalize Dual Scale to cases in which there are three or more options, just enough to provide a satisfactory answer to the Which Alternatives Question. This generalization will reveal an important respect in which the reasons for φ can vary as you vary the alternative.

By the end of the paper, then, I will have clarified the relationship between reasons for and against, as well as have provided three considerations in favor of Dual Scale: (1) it follows from the metaphor of a reason's weight; (2) it can represent normative theories that reject Single Proportion, such as the standard account of supererogation; and (3) it can weigh both reasons for and reasons against.

1. From Single Scale to Dual Scale

In §1.1, I remind you of Gert's distinction between justifying and requiring weight for, as well as forestall a potential misunderstanding of requiring weight. In §1.2, I explain why Single Scale can't represent the standard account of supererogation. In §1.3, I introduce Dual Scale and show that it can represent the standard account.

1.1. Justifying vs Requiring Weight for

A consideration has **justifying weight for** φ iff the consideration makes it permissible to φ in the absence of sufficiently weighty countervailing reasons (more precisely:

[2] White (2017) objects to the converse of Equivalence, i.e., if R is a reason for φ, then R is a reason against some alternative. I will establish the converse as a conceptual truth too, so White's arguments must go wrong somewhere. For where they might go wrong, see Kiesewetter (2018: §III).

sufficiently weighty requiring weight for ~φ). The fact that eating this cookie will make me happy has justifying weight insofar as it is a consideration that pushes *eating the cookie* toward permissibility. A consideration has **requiring weight for** φ iff it makes ~φ-ing impermissible in the absence of sufficiently weighty countervailing reasons (more precisely: sufficiently weighty justifying weight for ~φ). If stopping the car is the only way to avoid killing one of the children playing in the street, then you have a reason for stopping that has requiring weight, a reason that pushes not stopping toward *im*permissibility.

You'll notice that 'justifying weight' and 'requiring weight' are (inter)defined in typical functionalist fashion. To have justifying weight for φ is just to play a certain role in fixing deontic status. It is to push φ toward permissibility. To have requiring weight for ~φ just is to play an *opposing* role in fixing deontic status. It is to push φ toward *im*permissibility. Something's *amount* of justifying weight and/or requiring weight is a way of specifying which role wins the competition in a given case.[3] This functionalist approach to reasons and their relation to deontic status will allow us to state precisely the relationships between reasons for and reasons against. I will not defend this functionalist approach, except insofar as the cool things I do with it count in its favor.[4]

The difference between justifying and requiring weight is not a difference *in kind*. It is a difference of *direction*. There is just one kind of physical weight, but we can nonetheless distinguish between *which direction* physical weight pushes and *what option* it pushes in that direction. If you put my phone on the left pan, it will push the left pan down and the right pan up. It pushes the two pans in opposite directions. Likewise, even if there is only a single kind of normative weight, we can distinguish between *which deontic direction* it pushes (toward permissibility or toward impermissibility) and *what option it pushes* in that direction (φ or the alternative).

Following Portmore (2011: 88-9) and Dorsey (2016: 166), I use a thick notion of requirement: φ is a **requirement** iff both φ is permissible and ~φ is impermissible. I use a different term for Gert's thinner, purely negative notion of requirement: φ is a **commitment** iff ~φ is *im*permissible. Commitment to φ is what you add to permission to φ in order to make φ required. Nonetheless—and in contrast to Dorsey (2016: 166)—I follow Gert in using a thin sense of 'requiring weight' that does *not* entail justifying weight. So understood, an undefeated requiring reason to φ does not, by definition, entail a requirement to φ (both φ is permissible and ~φ is impermissible) but only a commitment to φ (~φ is impermissible). Were Gert's terminology not ingrained in the literature, I would have avoided this infelicity by replacing the term "requiring weight" with "committing weight".

1.2. The Problem for Single Scale Explained

[3] Gert sometimes claims that, if justifying and requiring weight were independent variables, then *amounts* of justifying and requiring weight would be incomparable (2007: 537). This is a mistake. Downward and upward forces are distinct but comparable: the downward force can be greater than, less than, or equal to upward force. Contrary to what Gert himself asserts (2007: 536), justifying and requiring weight are opposing weights/forces to be understood on analogy with downward and upward force. Hence, justifying weight and requiring weight are distinct but comparable: justifying weight for φ can be greater than, less than, or equal to requiring weight for ~φ.

[4] For further defense of this functionalist approach, consult Dorsey 2016: 8-12 and especially Gert 2004: ch 4.

Single Scale entails that justifying and requiring weight exist. Justifying weight for φ is *nothing more than* weight that pushes the φ pan down, making φ permissible (in the absence of sufficiently weighty countervailing reasons). So R_φ have justifying weight for φ. They also have requiring weight for φ. Requiring weight for φ is *nothing more than* weight that pushes the alternative up, making the alternative impermissible (in the absence of sufficiently weighty countervailing reasons).

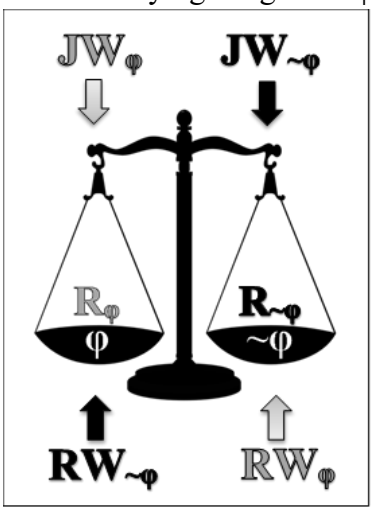

So far, so good. But Single Scale doesn't stop there. Since a reason pushes the φ-pan down (has justifying weight for φ) only insofar as it pushes the ~φ-pan up (has requiring weight for φ), Single Scale is committed to:

> **Single Proportion:** reasons always have the same proportion of justifying and requiring weight.[5]

Single Proportion is a substantive normative judgment about the *weights* that reasons can have. In my forthcoming, I argue that a model of *weighing* reasons should not presuppose Single Proportion even if it is true. Here I assume that Single Proportion is false, because it is incompatible with the standard account of supererogation.

Liv can save the lives of five soldiers by jumping on a grenade, thereby sacrificing her own life. Or she can remain in safety and allow the fellow soldiers to die. Liv's jumping on the grenade would be a paradigmatic instance of supererogation, of going beyond the call of duty (in a morally good way). The standard account provides the following explanation:

> The lives of the five soldiers have both justifying and requiring weight, because Liv would be morally required to save the five soldiers if she could do so at no cost to herself. Morality is other-centered, so Liv's self-interested reasons lack requiring weight. Yet she is permitted to remain in safety because her life has considerable justifying weight. Since both alternatives are permissible and it is morally better to save the others, she goes "beyond the call of duty" in jumping on the grenade.[6]

Single Scale can't represent the standard account of supererogation, because the latter depends on certain reasons (altruistic ones) having a different proportion of justifying and requiring weight than others (self-interested ones). The former have both justifying and requiring weight; the latter have only justifying weight.[7]

1.3. Dual Scale

[5] Single Scale actually has a bit more flexibility than this (my forthcoming, §3), but the details won't concern us.

[6] Portmore (2008; 2011: ch 5) is perhaps the clearest proponent of this account. His account is assumed by Massoud (2016) and Archer (2016). Similar accounts are given by Clark (1978), Curtis (1981), Muñoz (forthcoming), and Raz (1990: 94). Raz's response focuses on "exclusionary permissions" rather than merely justifying reasons. Yet "exclusionary permission" seems to be an awkward way of referring to an undefeated merely justifying reason (Gert 2004: 107-110).

[7] Other theories which entail that justifying and requiring weight are independent variables include: the existence of prerogatives (Scheffler 1982) and mere permissions (Hurka and Shubert 2012); solutions to Kamm's Intransitivity Paradox (Kamm 1996; Archer 2016; Muñoz forthcoming) and Horton's All or Nothing Problem (Muñoz forthcoming); and certain versions of satisficing theory (Tucker 2017).

My model, **Dual Scale**, uses two scales to determine the full deontic status of an action. **Permission Scale** determines whether φ is permissible by comparing the justifying weight for φ (JW_{φ}) with the requiring weight for ~φ ($RW_{\sim\varphi}$). **Commitment Scale** determines whether φ is a commitment by comparing the requiring weight for φ (RW_{φ}) with the justifying weight for ~φ ($JW_{\sim\varphi}$). The two scales working together determine whether φ is required (whether it is a permissible commitment).

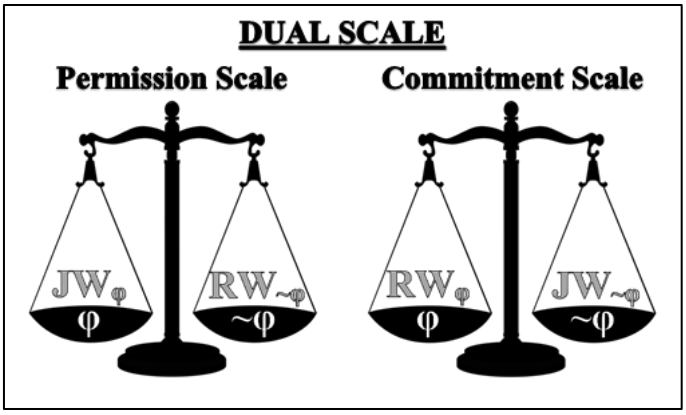

Compare Dual Scale with the illustration of how JW and RW map onto Single Scale. Permission Scale represents the same competition as the left side of Single Scale (JW_{φ} vs $RW_{\sim\varphi}$). Commitment Scale represents the same competition as the right side of Single Scale ($JW_{\sim\varphi}$ vs RW_{φ}). The only difference is that there is no bar between the Permission and Commitment Scale that forces JW and RW to be dependent variables. In other words, the sole difference between Single and Dual Scale is that only the former takes a substantive, controversial stand on the relation between justifying and requiring weight.

Any balance scale model must assign relative weights to deontic status. I work with:

Permission Assignment: φ is permissible iff the justifying weight of all reasons for φ is *at least as weighty* as the requiring weight of all reasons for ~φ, and

Commitment Assignment: φ is a commitment iff the requiring weight of all reasons for φ is *weightier than* the justifying weight of all reasons for ~φ.

By relying on two scales, Dual Scale can represent what Single Scale can't: normative views that allow justifying and requiring weight to be independent variables (some reasons have a different proportion than others). Liv must choose whether to sacrifice her life (**Sacrifice**) or remain in safety (**Safety**). The standard account of supererogation treats altruistic reasons, such as the lives of the five soldiers, as requiring reasons (say, 500 $JW_{Sacrifice}$ and 500 $RW_{Sacrifice}$). The account treats self-interested reasons, such as the value of her life, as very weighty merely justifying reasons (say, 1,000 JW_{Safety} and 0 RW_{Safety}). Permission Scale entails that Sacrifice is permissible, because the justifying weight of Liv's altruistic reason is weightier than the non-existent requiring weight of her self-interested reason (500 $JW_{Sacrifice}$ > 0 RW_{Safety}). Commitment Scale entails that Sacrifice isn't a commitment, because the requiring weight of the altruistic reason is outweighed by the justifying weight of her self-interested reason (500 $RW_{Sacrifice}$ < 1,000 JW_{Safety}).

(when assessing deontic status of Sacrifice)

Permission Scale		**Commitment Scale**	
500 $JW_{Sacrifice}$	0 RW_{Safety}	500 $RW_{Sacrifice}$	1,000 JW_{Safety}
Verdict: Sacrifice is permissible		Verdict: Sacrifice is not a commitment	
Combined verdict: Sacrifice is permissible but not required			

Since the altruistic act is permissible but not a commitment (and so not required), one goes beyond the call of duty in performing it. At least, one goes beyond the call on the plausible assumption that it is morally better to save the five soldiers than to save one's own life.

In short, Dual Scale holds that it takes two competitions to determine the full deontic status of φ: it takes a single use of Permission Scale to determine whether φ is permissible and a single use of Commitment Scale to determine whether φ is a commitment (whether ~φ is impermissible). Together the two competitions reveal whether φ is required.

The above version of Dual Scale assumes that there are exactly two options, usually represented by φ and ~φ. In §6, I gesture at how it should be generalized to handle cases in which there are three or more options. The above version is simplified in other respects too. For example, it may seem incompatible with the existence of conditions (enablers, disablers) and modifiers (intensifiers, attenuators). Dual Scale is easily refined to make room for such things, but I do not have the space to explore, clarify, and defend all its various features. I do much of that in my forthcoming and manuscript. Here I clarify the relationship between reasons for and against (§2, §4, §6), explain how Dual Scale handles such reasons (§3), and use these results to argue for Dual Scale (§5).

1.4. An Alternative to Dual Scale?

If Snedegar gives up his attack on balance scale models, he has the resources to construct a balance scale model of weighing reasons that may seem to be an alternative to Dual Scale and the justifying/requiring distinction. The model would be stated in terms of a single distinction between reasons for and against. φ is permissible iff the reasons for φ are at least as weighty as the reasons against φ. To make sense of the standard account of supererogation, Snedegar could deny that every reason for φ is a reason against the alternative: Liv's death would be a reason for Safety without being a reason against Sacrifice. Sounds pretty good so far, right? That's because, so far, this alleged alternative—call it *Dualplicitous Scale*—is just a terminological variant of Dual Scale. It uses 'reasons for φ' as code for justifying weight for φ (reasons that push φ toward permissibility); uses 'reasons against φ' as code for requiring weight against φ (reasons that push φ toward *im*permissibility); and denies Single Proportion, that justifying and requiring weight always come in the same proportion.[8]

Whether Dualplicitous Scale remains a terminological variant of Dual Scale depends on how it is developed to handle the issues in the remainder of the paper. Be warned, though, that its terminology is apt to obscure. It draws a *single* distinction between reasons for and against when there are really *two* (§2.1).[9] It is no surprise, then, that Snedegar misses the conceptual equivalences between reasons for and reasons against (§2.2, §4), equivalences which underwrite the argument from weight to Dual Scale (§5).

2. Reasons For vs Reasons Against

[8] In contrast to Single Scale, the alleged alternative makes reasons for and against *in*dependent variables; therefore, like Dual Scale, it needs a second scale to account for commitment and requirement. The first scale determines the permissibility of φ by comparing the reasons for and against *φ*. The second scale takes into account the reasons for and against *~φ*: φ is required only if the reasons against *~φ* outweigh the reasons for *~φ*.

[9] So far, Dualplicitous Scale has no terms that correspond to my justifying weight against or my requiring weight for. How big of a problem is this? That depends on how useful it is to have such terms. It seems pretty useful to me (see the end of §2.2; cf. the discussion of opportunity costs and benefits in §6.3).

2.1. Justifying and Requiring Weight Against

There are *two* for/against distinctions, one for justifying weight and one for requiring weight. In §1.1, I defined the functional roles of justifying and requiring weight *for*. These definitions imply the functional roles for justifying and requiring weight *against*. Justifying and requiring weight for/against are comprised by two general distinctions that cut across one another. The justifying/requiring distinction specifies *which direction* the weight pushes (toward permissibility or toward impermissibility). The for/against distinction specifies *what* is pushed in that direction (φ or the alternative).

Justifying weight always pushes something toward permissibility. Justifying weight for φ pushes φ toward permissibility (§1.1). Hence, **justifying weight *against*** φ pushes *the alternative* toward permissibility. You can take an outside job offer (Go) or stay at your current job (Stay). If you Go, you will live in a less desirable location. The worse location has justifying weight *against Go*. That is, it pushes *Stay* toward permissibility.

Requiring weight always pushes something toward *im*permissibility. Requiring weight for φ pushes the alternative toward impermissibility (§1.1). Hence, **requiring weight *against*** φ pushes *φ* toward impermissibility. If you Go, you will break your promise to your nearby sister, a single mother, to watch her kids once a week. The broken promise has requiring weight *against* Go insofar as it pushes Go toward *im*permissibility.

Reasons often have both justifying and requiring weight. The promise has requiring weight against Go (pushes Go toward impermissibility), but it also has justifying weight against go (pushes Stay toward permissibility). It wouldn't be shocking if most reasons had equal amounts of justifying and requiring weight; however, this isn't true of all reasons. To represent the standard account of supererogation, we assume that justifying and requiring weight can vary independently of one another.

2.2. Conceptually Necessary Functional Equivalences

Single and Dual Scale are different ways of capturing the idea that an act's deontic status depends on the reasons that bear on φ *and* the reasons that bear on the alternative, ~φ. Single Scale holds that there are four relevant categories of reasons: two categories that bear on φ (reasons for φ and reasons against φ) and two categories that bear on ~φ (reasons for ~φ and reasons against ~φ). Four categories of reasons may sound like a lot. Yet Single Scale is problematic precisely because these categories are too few to represent normative theories which allow justifying and requiring weight to be independent variables.

Once we consider the possibility that justifying and requiring weight might be independent variables, the deontic status of φ is revealed to be a function of *eight*(!) categories of weight. If we use '**»«**' to represent against, we get:

Eight Weight Slate

	For	**Against**
Justifying weight	justifying weight for φ (JW_{φ}) justifying weight for ~φ ($JW_{\sim\varphi}$)	justifying weight *against* φ ($JW_{\text{»«}\varphi}$) justifying weight *against* ~φ ($JW_{\text{»«}\sim\varphi}$)
Requiring weight	requiring weight for φ (RW_{φ}) requiring weight for ~φ ($RW_{\sim\varphi}$)	requiring weight *against* φ ($RW_{\text{»«}\varphi}$) requiring weight *against* ~φ ($RW_{\text{»«}\sim\varphi}$)

There is good news for those who are starting to get overwhelmed by the distinctions. Everything in the against-column is functionally equivalent to something in the for-column. Every functional equivalence fits this pattern: hold fixed whether it is justifying or requiring weight and then vary the for/against *and* the φ/~φ. For example, justifying weight *for φ* is functionally equivalent to justifying weight *against ~φ*. They both push φ toward permissibility. We can summarize the equivalences with:

Eight Weight Mates[10]

	for	**Against**
Justifying weight	JW_{φ} $JW_{\sim\varphi}$	$JW_{\gg\ll \varphi}$ $JW_{\gg\ll \sim\varphi}$
Requiring weight	RW_{φ} $RW_{\sim\varphi}$	$RW_{\gg\ll \varphi}$ $RW_{\gg\ll \sim\varphi}$

If an arrow connects two categories, then they are functionally equivalent (by definition). If there is *no* arrow between two categories, then they are *not* functionally equivalent, at least not by definition.[11]

When two categories of weight are functionally equivalent, any reason that fits one category fits the other. For example, any reason that plays *the justifying weight against ~φ* role in a competition with a given alternative also plays the *justifying weight for φ* role in that competition, and vice versa.

According to Eight Weight Mates, nothing in the justifying weight row is functionally equivalent with anything in the requiring weight row. This shouldn't be surprising. We are stressing the justifying/requiring weight distinction precisely because we want to represent normative views that allow justifying and requiring weight to be *in*dependent variables.

Recall the Liv case. If Liv jumps on the grenade, she sacrifices her life. Is (the value of) her life justifying weight *for* Safety or justifying weight *against* Sacrifice? The choice may seem arbitrary. That's because the two types of weight, by definition, are functionally equivalent. Insofar as Safety and Sacrifice are direct competitors, justifying weight *for Safety* just is justifying weight *against Sacrifice*, i.e., weight that pushes Safety toward permissibility.

You might wonder, "if *justifying weight against φ* and *justifying weight for the alternative* really were conceptually equivalent, then what's the point in having the distinction at all?" Consider an analogy. *Joe is taller than Jack* is conceptually equivalent

---

[10] In case the table isn't clear already, I will spell out the four conceptual equivalences that it identifies, starting with the top of the against column:

($JW_{\gg\ll \varphi}$): By definition, justifying weight against φ pushes ~φ toward permissibility, which justifying weight for ~φ also does by definition.

($JW_{\gg\ll \sim\varphi}$): By definition, justifying weight against ~φ pushes φ toward permissibility, which justifying weight for φ also does by definition.

($RW_{\gg\ll \varphi}$): By definition, requiring weight against φ pushes φ toward *im*permissibility, which requiring weight for ~φ also does by definition.

($RW_{\gg\ll \sim\varphi}$): By definition, requiring weight against ~φ pushes ~φ toward *im*permissibility, which requiring weight for φ also does by definition.

[11] If justifying for and requiring weight for are dependent variables after all, then those conceptually distinct types of weight might not be distinct. Analogously, 'Biden' is conceptually distinct from 'the President of the US' but Biden is not distinct from the President of the US.

to *Jack is shorter than Joe*. Nonetheless, we find it convenient to have both ways of talking about the same relation of relative height. The former emphasizes Joe's side of the relation. The latter emphasizes Jack's side.

A similar point holds for reasons. Liv's life counts as a reason. If we want to emphasize how her life bears on Sacrifice's deontic status (i.e., it pushes sacrificing her life away from being a commitment and so away from being required), we can say that Liv's life is a reason *against Sacrifice*. If we want to emphasize how Liv's life bears on Safety's deontic status (i.e., it pushes remaining in safety toward permissibility), we can say that her life is a reason *for Safety*. *Justifying weight for φ* and *justifying weight against the alternative* are conceptually equivalent ways to talk about the same reason. Nonetheless, their pragmatics are different. The former emphasizes the reason's relation to φ's deontic status. The latter emphasizes its relation to the alternative's deontic status.[12]

3. Dual Scale and Reasons For/Against

Eight Weight *Slate* tells us that there are four categories of weight for and four categories of weight against. Recall the original illustration of Dual Scale (pictured again). It tells us where to put the four categories of weight for. It doesn't tell us where to put any category of weight against. Thankfully, it is easy to fill in those details with the help of Eight Weight *Mates*.

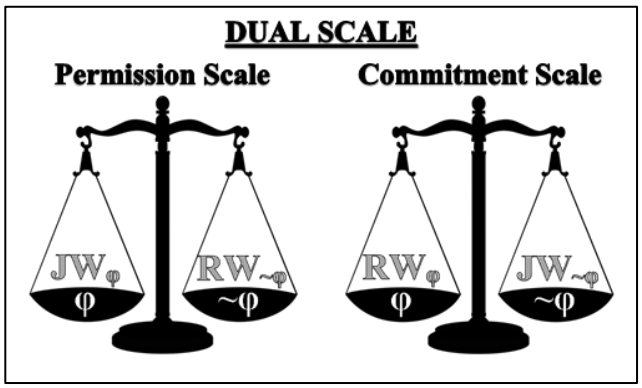

For X and Y to be on the same pan is just for X and Y to play the same functional role in fixing deontic status. If two categories of weight, by definition, play the same functional role in fixing an act's deontic status, then they go on the same pan. When we combine the illustration of Dual Scale with Eight Weight Mates, we find where to put every category of weight against:

Where to Put the Weight
(when assessing the deontic status of φ)

Permission Scale		**Commitment Scale**	
Left Pan of PS	**Right Pan of PS**	**Left Pan of CS**	**Right Pan of CS**
JW for φ	RW for ~φ	RW for φ	JW for ~φ
JW against ~φ	RW against φ	RW against ~φ	JW against φ

In short, Dual Scale easily weighs all four categories of weight for and all four categories of weight against.

[12] Kiesewetter defends the idea that a reason for φ always correlates to some reason against an alternative, but he does so under the label "Reasons Transmission" (2015: §IV, 2018: §III; cf. Snedegar 2016: §3). This is misleading. The correlation holds, but not because the reasons for φ transmit to something the property of *being a reason against the alternative*. 'Reason for φ' and 'reason against the alternative' are just two different ways to talk about the same reason.

4. Equivalence Established

Snedegar (2018: 738) rejects:

Equivalence: Insofar as R is a reason *against* φ, R is a reason *for* some alternative.

At first glance, he seems right. Recall our earlier example. We assumed that the burns you'll receive if you Save2 are a reason *against* Save2 in its competition with Save1. (We'll revisit this assumption shortly.) It may seem implausible that those burns are also a reason for Save1 when Save1 results in those very same burns (cf. 732). Yet he's wrong. Eight Weight Mates reminds us that, in Save2's competition with Save1, justifying weight *against Save2* just is justifying weight *for Save1*. Justifying weight against Save2 pushes its direct competitor toward permissibility, which is exactly what justifying weight for Save1 does.

The equivalence between justifying weight *against φ* and justifying weight *for ~φ* is non-negotiable. We did not arrive at Eight Weight Mates by reflection on the scale metaphor, much less idiosyncratic details of Dual Scale. We arrived at it by pointing out how the various functional roles relate to one another *by definition*. Eight Weight Mates reports *conceptual truths* about the connection between the eight categories of weight. Dual Scale just provides a helpful visualization of their conceptual equivalence. Recall the Where to Put the Weight table. Since justifying weight against φ and justifying weight for the alternative go on the same pan, we can "see" that they make the same functional contribution to an act's deontic status. [13]

Maybe we are missing Snedegar's point because we assumed that the burns were a reason against Save 2 in its competition with Save1. That assumption itself presumes **noncomparativism**: a reason for/against φ might consist in facts solely about φ, regardless of how φ compares to other alternatives. Noncomparativism holds, in other words, that reasons for/against φ don't need to be *differences* between φ and some alternative. In the competition between Save1 and Save2, the burns might be a reason against Save2 (and a reason for Save1) even though you get exactly the same burns in Save1.

Snedegar rejects noncomparativism[14] in favor of **comparativism**: all reasons are comparative, i.e., all reasons are *differences* between alternatives. On this view, since the burns are equally bad in Save1 and Save2, they aren't a reason for/against either option after all. No problem. Equivalence is trivially preserved. The burns' *lack* of weight against Save2 is matched by a *lack* of weight for Save1.

In contrast, since you get the burns in Save2 but not in Bystander (just stand around and do nothing), the comparativist will hold that the burns *are* a reason against Save2 when

---

[13] I take no stand on whether Equivalence is true for the third role of reasons, *(dis)commending*. That is, I take no stand on whether every reason that commends φ also discommends some alternative. Commending has a role to play in supererogation by making some acts better than alternatives *with the same deontic status*. Yet it is irrelevant to fixing deontic status. In my manuscript, I argue that commending and Dual Scale work together to explain supererogation and resolve the All or Nothing Problem.

[14] Snedegar finds it odd to say that the burns are a reason for Save1 when you get those very same burns in Save2 and Save2 is your only alternative. This oddness allegedly counts against noncomparativism because a theory of reasons should "explain the claims that we are willing to accept about reasons" (733). I'm more optimistic than Snedegar about the prospects for a pragmatic explanation of why it sounds odd to say that the burns are a reason for Save1 (cf. Snedegar 733, nt 11). When it is obvious to everyone that a reason for φ is cancelled out by the same reason for ~φ, then there is no point in talking about it in everyday decision-making. So when we are choosing between Save1 and Save2, we are going to focus on only those reasons that apply to one option but not the other. With that said, I am neutral between comparativism and noncomparativism.

it competes with Bystander. Yet this is no counterexample to Equivalence either. The burns push Bystander toward permissibility (and Save1 away from requirement), which is the defining functional role of both justifying weight *against Save2* and justifying weight *for Bystander*. Consequently, the burns are both a reason against Save1 and for Bystander. Equivalence is again preserved. Comparativism about reasons affects which considerations count as reasons, but it doesn't threaten the claim that all reasons obey Equivalence.

Given comparativism, the burns are a reason against Save2 when Save2 competes with Bystander, but not when it competes with Save1. This suggests that there may be some respect in which weight (and so reasons) are relative to the alternative. We will explore this relativity in §6.1.

5. From Weight to Dual Scale

Snedegar holds that we can keep the metaphor of a reason's *weight* (741) while rejecting the metaphor of weighing reasons on a *scale* (742).[15] Yet these metaphors are *not* separable. They appear to be separable in part because the scale metaphor tends to be corrupted by Single Proportion and the assumption that there is only a single scale. Once we drop these corrupting influences, we'll see that Dual Scale adds *nothing* to the metaphor of weight—except a useful visualization.

A complete defense of this claim would identify all features of Dual Scale and show that they are entailed by the features present in the weight metaphor. My defense here will be more limited, but it will get the job done well enough. I identify the *minimal idea* of a reason's weight. I then construct a scale—*Proto* Permission Scale—that captures only this minimal idea. I'll then show that this scale is conceptually equivalent to Permission Scale and that Commitment Scale is just a substitution instance of Permission Scale. Hence, if you endorse the metaphor of weight, you are committed to Dual Scale.

Suppose that a given bench press bar weighs 45 pounds. That 45 pounds of *weight* just is the *force* that pushes the bar down. In the real world, forces can push up, down, left, right, and everything in between. The metaphor of a reason's weight is just a convenient way to focus our attention on forces in two opposing directions: toward permissibility and toward impermissibility. If you push something toward anything at all, you thereby push it away from something else. Hence, the metaphor of weight (pressure, force) gives us this **minimal idea**: as φ is pushed toward permissibility, it is pushed away from impermissibility.

To construct a scale that captures only this minimal idea, start with a single balancing scale. Ordinarily, the scale metaphor lets one pan represent φ and the other represent ~φ. Here, though, let *both* pans of the scale represent φ.

> *Left Side of Proto Permission Scale:* Let down represent toward φ's permissibility and up represent toward φ's *im*permissibility. Hence, on the left pan we put justifying weight for φ (JW_{φ}), i.e., those reasons that push φ down toward permissibility.

As the left pan goes down, the right side goes up. Consequently:

> *Right Side of Proto Permission Scale: Up* is toward φ's permissibility and *down* is toward φ's impermissibility. Hence, on the right pan we put requiring weight against φ ($RW_{\text{»« }\varphi}$), i.e., those reasons that push φ down toward *im*permissibility.

[15] Dancy (2004: 105) agrees.

The resulting scale is **Proto Permission Scale**. See the below illustration. The minimal idea of weight is that, as φ is pushed toward permissibility, it is pushed away from impermissibility. Proto Permission Scale illustrates this minimal idea on both sides of the scale (in an unhelpful way) but otherwise adds nothing to it. To go from Proto Permission Scale to Dual Scale's Permission Scale, we need to modify the bottom, top, and the middle of Proto Permission Scale's *right* side. Each of these modifications results from replacing the relevant part of Proto Permission Scale with something conceptually equivalent.

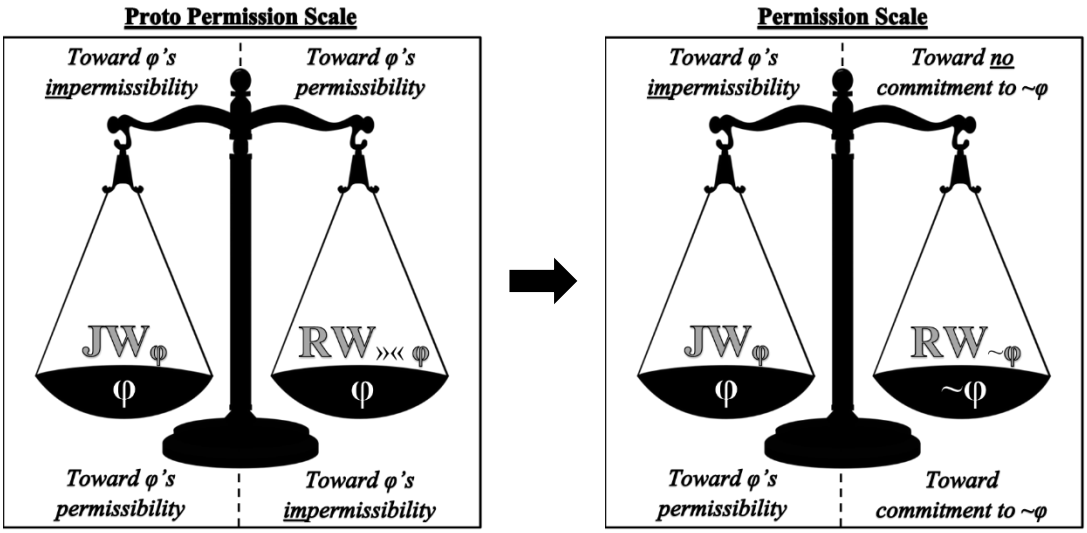

Bottom Right Modification to Proto Permission: For φ to be impermissible just is for ~φ to be a commitment, per the definition of 'commitment' to ~φ (§1.1). So pushing φ toward impermissibility is conceptually equivalent to pushing ~φ (the alternative) toward commitment. This conceptual equivalence vindicates the modification to the bottom right of Proto Permission Scale: replace "Toward φ's impermissibility" with "Toward Commitment to ~φ".

Top Right Modification to Proto Permission: By the definition of commitment, to the extent that a reason pushes φ toward permissibility, it pushes ~φ toward *not* being a commitment. This conceptual equivalence vindicates the modification to the top right of Proto Permission Scale: replace "Toward φ's permissibility" with "Toward no commitment to ~φ".

Middle Right Modifications to Proto Permission: Two changes must be made in the middle. The first is to the pan and the second is to what goes on the pan. We modified the bottom right so that it is now toward commitment to ~φ. We modified the top right so that it is now toward no commitment to ~φ. These changes vindicate changing the right φ pan into a ~φ pan. For, on the right side, it is ~φ that is pushed toward or away from commitment. The Eight Weight Mates table reminds us that $RW_{\gg\ll \varphi}$ is conceptually equivalent to $RW_{\sim\varphi}$. This conceptual equivalence vindicates the modification to what goes on the right pan: replace "$RW_{\gg\ll \varphi}$" with "$RW_{\sim\varphi}$".

Proto Permission Scale adds nothing to the minimal idea of the weight metaphor, except an illustration. This illustration is unhelpful because both sides of the scale illustrate the same simple idea: as φ is pushed toward permissibility, it is pushed away from *im*permissibility. Dual Scale's Permission Scale is conceptually equivalent to Proto

Permission Scale. Likewise, then, Dual Scale adds nothing to the minimal idea of the weight metaphor, except an illustration. In this case, however, the illustration is helpful. For it illustrates how certain reasons relevant to φ's deontic status affect the deontic status of φ's alternative: as φ is pushed toward permissibility, ~φ is pushed away from commitment (and so pushed away from requirement).

To get Dual Scale's Commitment Scale, start with its Permission Scale. Swap φ and ~φ and then flip the two sides of the scale (e.g., after the substitution, the *left* side of Permission Scale is the *right* side of Commitment Scale).[16] For your convenience, I again give you the original illustration of Dual Scale so that you can see for yourself that Commitment Scale is the mirror image of Permission Scale after swapping φ and ~φ. We shouldn't be surprised by the relation of the two scales. A requirement to φ is just a combined deontic verdict of permissible to φ (determined by Permission Scale) and *im*permissible to ~φ (determined by Commitment Scale).

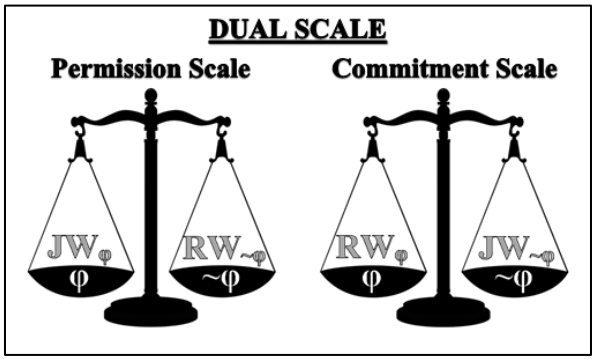

When we combine Permission and Commitment Scale, we get Dual Scale. Permission Scale is conceptually equivalent to the minimal idea of a reason's weight. Since Commitment Scale is conceptually equivalent to Permission Scale, it too is conceptually equivalent to the minimal idea of a reason's weight. Thus, if you are comfortable talking about a reason's weight, then you should be comfortable weighing that reason on Dual Scale.

6. The Which Alternative Question

6.1. Any Number of Options and the Relativity of Weight

Thus far I have made the simplifying assumption that there are exactly two options, usually represented by φ and ~φ. Snedegar will worry that this simplifying assumption ignores a challenging question for defenders of Equivalence. When there is more than one alternative to φ, Snedegar (2016: 728) raises:

> **The Which Alternatives Question (WhichAlt):** if R is a reason against φ and every reason against φ is a reason for some alternative (as Equivalence assumes), then *which* alternatives is R a reason for?

See my book manuscript for a complete explanation of how to generalize Dual Scale to cases with three or more options. Here I just explain the two things you need to know in order to answer WhichAlt: (i) what it is for an option to win a tournament, and (ii) that weight is contrastive, roughly, that weight can vary as you vary the alternative (i.e., the contrast).[17]

[16] Commitment Scale can also be constructed by going from the minimal idea of a reason's weight to a "Proto Commitment Scale" and then to the Commitment Scale itself. But going straight from Permission Scale to Commitment Scale saves us some time.

[17] Snedegar defends something similar to the tournament idea (2017: §3.4.3) and what I call 'contrastive weight' (2017: 7-8). Yet, in his 2018, he misses the way that balance scale models can be combined with these ideas to answer WhichAlt. Also, a word of caution: I *think* Snedegar builds more into contrastive

Recall that, given exactly one alternative, Dual Scale holds that *two* competitions determine the full deontic status of φ. A single use of Permission Scale determines whether φ is permissible. A single use of Commitment Scale determines whether φ is a commitment (whether ~φ is impermissible). The two scales working together reveal whether φ is required. Let's ignore the second scale and focus solely on how Dual Scale determines the permissibility of φ.

When the only options are φ and ~φ, Dual Scale holds that φ is made permissible by winning one competition with one alternative. It's also true, though, that φ is made permissible by winning a *tournament*, a pairwise competition with *every* alternative. This latter idea is the key to generalizing Permission Scale to contexts in which there are more than two options. The generalized Permission Assignment is this: φ is permissible iff, for every alternative A, $JW_{\varphi} \geq RW_{A}$. The metaphor of Permission Scale forces us to compare φ with one alternative at a time. We compare φ with each alternative, and if φ wins every comparison, then φ is permissible.

To see that weight is contrastive, consider the **Café or Kid Case**.[18] There is a single child in a burning building. I have three potential options. I can save the child and get severe burns (**Save1**). I can just stand around and watch the events unfold (**Bystander**). Or I can keep my promise to meet my friend at the Café (**Café**). If you had to choose between Bystander and Save1, then Bystander would be permissible and Save1 would be supererogatory. When you also have Café as an option, Save1 is supererogatory, Café is permissible, but now Bystander is *im*permissible. Somehow the addition of Café makes Bystander impermissible. This is puzzling. Consider the cost of breaking two fingers. This cost would justify not keeping your promise, but it wouldn't justify not saving the child. Intuitively, it takes a greater cost to justify not saving someone's life (100 RW_{Save1}) than to justify not keeping your promise (25 $RW_{Café}$). So how can $JW_{Bystander}$ beat 100 RW_{Save1} but not beat 25 $RW_{Café}$?

The answer is that there is more JW for Bystander when Bystander competes against Save1 than Café. This is easiest to see given comparativism about reasons, the view that a reason for φ is a way that φ is better than some alternative. If you vary the alternative, you can vary the ways and extent to which φ is better.[19] This is just what happens in the Café or Kid case.

> *Bystander vs Save1:* you avoid the burns in Bystander but not Save1. We are working with the standard account of supererogation which makes self-interested costs very weighty merely justifying reasons (§§1.2-1.3), so assume that this difference has 500 $JW_{Bystander}$. 500 $JW_{Bystander}$ > 100 RW_{Save1}.
>
> *Bystander vs Café:* the burns are irrelevant to the competition, because you don't get the burns in either alternative. Since there are no other ways Bystander is better than Café, there is 0 $JW_{Bystander}$. 0 $JW_{Bystander}$ < 25 $RW_{Café}$.

---

weight than I do. He seems to think it entails comparativism about reasons (2017: 8), but I don't (see nt 19 below).

[18] Earlier discussions of the Café or Kid Case include Kamm (1996: 313-5) and Muñoz (forthcoming). My explanation of this case is essentially a more developed version of Muñoz's (EV 9).

[19] If you endorse noncomparativism about reasons, then you need to rely on Eight Weight Mates (i.e., the conceptual equivalence of reasons for and against) to explain the Café or Kid Case. When Bystander competes with Save1, the burns have justifying weight *against* Save1 and so have justifying weight *for Bystander*. When Bystander competes with Café, there is no justifying weight against Café (the burns are irrelevant to that competition) and so no justifying weight for Bystander.

Bystander is impermissible, because it loses the tournament. It wins the individual competition with Save1, but it loses to Café.

The Café or Kid Case illustrates that weight is **contrastive**: varying the alternative to φ can vary the reasons (better: grounds) relevant to φ. The burns are a relevant reason in Bystander's competition with Save1, but not its competition with Café. Consequently, there is more $JW_{Bystander}$ when Bystander competes with the former than the latter.[20]

6.2. Equivalence's Answer to WhichAlt

When we generalize Dual Scale to any number of options, it says that φ is permissible if and only if it wins a tournament, a pairwise competition with each alternative ($JW_{\varphi} \geq RW_A$ for all alternatives A). Since weight is contrastive, there may not be a single value for JW_{φ} that applies to the whole tournament. There is more $JW_{Bystander}$ when Bystander competes with Save1 than when it competes with Café. Hence, any general answer to WhichAlt must be stated at the level of pairwise competitions, not at the level of the whole tournament. We can clarify this point by taking a closer look at the reasons involved in the Café or Kid Case.

There are three reasons at play: the life of the potentially saved child (**Child**), the potential severe burns (**Burns**), and the potentially kept promise (**Promise**). When comparing Save1 and Bystander, Burns and Child are relevant but Promise is not. When comparing Bystander and Café, Promise is relevant but Burns and Child are not. When comparing Save1 and Café, all three reasons are relevant. The following table sums up the reasons that are relevant to each tournament and individual competition:

Competitions	**Relevant Reasons**
Tournament that Determines Bystander's Deontic Status:	
Bystander vs Save1	Burns (for Bystander/against Save1) Child (against Bystander/for Save1)
Bystander vs Café	Promise (against Bystander/for Café)
Tournament that Determines Save1's Deontic Status:	
Save1 vs Bystander	Burns (for Bystander/against Save1) Child (against Bystander/for Save1)
Save1 vs Café	Burns (against Save1/for Café) Child (for Save1/against Café) Promise (against Save1/for Café)
Tournament that Determines Café's Deontic Status	
Café vs Bystander	Promise (for Bystander/against Café)
Café vs Save1	Burns (against Save1/for Café)

[20] Contrastive weight does *not* entail **holism**, roughly the idea that the same reason's weight can vary as you vary the context (or alternative). Contrastive weight involves *changes in which reasons are relevant*. Holism involves *changes in the weight of a relevant reason*. Suppose some comparativist endorses **atomism**, the idea that a reason's weight never varies, i.e., that holism is false. Now consider the difference in burns between Bystander and Save1. The atomist comparativist says *both* that this difference in burns is a relevant reason whenever the difference applies *and* that the reason when relevant has 500 JW no matter the context. The holist comparativist will agree that the difference in burns is a relevant reason whenever it applies, but she may insist that the relevant reason's weight can vary as you vary the context. Maybe its justifying weight is 500 in Café or Kid, as we assumed, but only 300 in some very different context. In short, contrastive weight is compatible with holism but doesn't entail it.

	Child (for Save1/against Café) Promise (against Café/for Save1)

This table confirms Equivalence. Whenever a reason applies to a given competition, it is both a reason for and a reason against. For example, when we take a micro perspective and focus on the individual competition between Bystander and Save1, it seems arbitrary whether we treat the burns as having justifying weight against Save1 or justifying weight for Bystander. Either way, the weight's function is to push Bystander toward permissibility.

The table also helps answer WhichAlt. When we zoom out, we see that the rational significance of the severe burns "follow" or "target" Save1, not Bystander.[21] The burns have justifying weight *against* Save1 in Save1's competition with Bystander *and* its competition with Café. But they have justifying weight for Bystander only in Bystander's competition with Save1. In contrast, some reasons against φ *don't* track or follow φ. Child is a reason against Café in its competition with Save1; but Child is irrelevant to Café's competition with Bystander. Child is a reason against Café *only insofar as* Café is an alternative to Save1.

These considerations reveal that the answer to WhichAlt is *not* as simple as: if R is a reason against φ, then it is a reason for *every* alternative. The answer is too simple, because it is stated at the level of the tournament. When R is a reason against φ in some pairwise competitions but not others, we should likewise expect that R is a reason for some alternatives but not others. Child is a reason against Café in its competition with Save1 but not its competition with Bystander. It is no surprise, then, that Child is a reason for Save 1 (in Café's competition with Save1) but not for Bystander (in Café's competition with Bystander).

It is easy to answer WhichAlt once we know that the answer should be stated at the level of individual competitions. Equivalence actually gives us a completely general, if somewhat unhelpful, answer to WhichAlt.

> **Equivalence's Answer to WhichAlt:** In the pairwise competition between φ and some alternative A1, R is a reason for A1 iff R is a reason against φ.

When Café competes with Save1, Child is a reason against Café, so it is a reason for Save1. When Café competes with Bystander, Child isn't a reason against Café and so it isn't a reason for Bystander either.

6.3. More Specific Answers to WhichAlt

Equivalence's answer to WhichAlt is so thin that it isn't very satisfying, and you may wonder whether more specific and satisfying answers can be given. They can, but they presuppose some substantive judgment about the *individuation* and/or *weights* of reasons. Hence, any more specific answer to WhichAlt goes beyond the strict purview of Dual Scale or any other model of *weighing* reasons.

Given comparativism about reasons, it is easy to give a more specific answer to WhichAlt. If R is a reason, it will be some difference between two alternatives. Consider Bystander and Save1. The difference in burns is a reason *for Bystander*, because it is a way that Bystander is better. It is a reason *against Save1*, because it is a way that Save1 is

[21] Snedegar emphasizes that "paradigmatic reasons against" target the option they are against (2018: 729; cf. 732), though I'm not sure we mean exactly the same thing by "target".

worse than Bystander. This same reason is relevant to any other pairwise competition in which the same difference applies. The difference in burns also applies to Save1's competition with Café, and so the burns are a reason *against Save1* and *for Café*. In contrast, recall the pairwise competition between Bystander and Café. Since you avoid the burns either way, the difference in burns doesn't apply and so aren't a reason for or against either option (§4). Generalize and we get this basic idea:

> **Comparativism's Answer to WhichAlt:** for any alternative to φ, A, R is a reason for A (and against φ) iff R is a way that A is better than φ.[22]

We need one more distinction to answer WhichAlt if we assume noncomparativism. Let's work with **Toy Noncomparativism**, which holds that all reasons boil down to *non-comparative* facts about (the consequences of) a particular option. For example, the burns are a reason against Save1 simply because Save1 results in those burns, *regardless of how severely one would be burned in any other option*. This reason against Save1 is local to Save 1. **Local** reasons for/against φ are located, or grounded, in φ (or φ's consequences).

Given Toy Noncomparativism, not all reasons are local in this sense. **Remote** reasons for/against φ are located, or grounded, in an alternative (or an alternative's consequences). Economists distinguish between local and remote reasons using the language of opportunity costs and benefits (cf. Snedegar 2018: 729). An outside job offer comes with a higher salary. This higher salary is a local reason to take the job offer. It is also a remote reason against, or an *opportunity cost* of, staying put. To stay put is to miss the opportunity for the higher salary provided by the outside offer.

In the Café or Kid Case, the burns you get if you Save1 are a remote reason *for Bystander*. They are a reason for Bystander that is grounded in an alternative, Save1. This reason for Bystander is still a noncomparative reason. The reason isn't *you get burned worse in Save1 than in Bystander*, as comparativism would have it. The reason is *you get severely burned in Save1*.

Every local reason is a remote reason, and vice versa. A local reason against φ is a remote reason for the alternative. Eight Weight Mates tells us that this relationship is a conceptual truth. For example, the burns' justifying weight *against Save1* is conceptually equivalent to the burns' justifying weight *for Bystander* (§2.2). Insofar as Save1 and Bystander directly compete, it is conceptually impossible for the burns to be a local reason against Save1 without also being a remote reason for Bystander.

Nonetheless, local and remote reasons have different relevance conditions. Since local reasons against φ are grounded in φ itself, those reasons are relevant no matter which alternative φ competes with. Burns are a local reason against Save1, so it is no surprise that they are relevant to Save1's competition with both Bystander and Café. When a reason against φ follows or targets φ—when it is a reason against φ in the competition with every alternative—then we have evidence that the reason is local.

In contrast, there is no presumption that remote reasons against φ will be a reason for every alternative, precisely because those reasons are grounded in some specific alternative(s). In principle, a remote reason against φ can be a reason for every alternative. When there are only two options, for example, every remote reason against φ will be a reason for every alternative. Yet remote reasons against φ tend not to be reasons for every alternative. Child is a reason against Bystander insofar as Bystander competes with Save1.

[22] Different versions of comparativism will cash out 'ways to be better than' in different ways. See, for example, Snedegar's (2018: 736-7) appeal to the promotion or respect of objectives.

Yet it is not a reason against Bystander in its competition with Café. If R is a remote reason against φ, then it is a reason for only certain alternatives: those alternatives for which R is a local reason.[23]

Let's return to WhichAlt. If R is a reason against φ, which alternatives is R a reason for? Our discussion of local and remote reasons gives us:

> **Toy Noncomparativism's Answer to WhichAlt:** If R is a *local* reason against φ, it is a reason for each alternative. If R is a *remote* reason against φ, it is a reason for only certain alternatives: those alternatives for which R is a local reason.

Conclusion

This paper concerned two topics: (i) the relationship between reasons for and reasons against and (ii) how reasons are to be weighed. Regarding the first topic, the paper has established that:

1. There are two categories of reasons for (justifying and requiring weight for) and two categories of reasons against (justifying and requiring weight against);
2. Every reason against is conceptually equivalent to some reason for;
3. Reasons are contrastive, i.e., which reasons are relevant may vary as you vary the alternative; and
4. A perfectly general answer to WhichAlt is available: in the pairwise competition between φ and some alternative A1, R is a reason for A1 iff R is a reason against φ. More specific answers are available once we add auxiliary assumptions about which facts are reasons.

Regarding the second topic, the paper identified the following virtues of Dual Scale:

1. It is entailed by the idea that reasons have weight (force, pressure);
2. It can represent normative theories that allow justifying and requiring weight to be independent variables, such as the standard account of supererogation; and
3. It easily weighs all categories of reasons for and all categories of reasons against.

These three virtues fall short of a complete defense of Dual Scale. I have not explained how to make Dual Scale compatible with the existence of conditions (enablers, disablers) or modifiers (intensifiers, attenuators). Nor have I fully explained how Dual Scale can handle cases with more than two options. Yet these three virtues do commend Dual Scale for further exploration.[24]

References

Archer, Alfred. 2016. "Moral Obligation, Self-Interest, and the Transitivity Problem." *Utilitas* 28: 441-464.

[23] This correlation between local and remote reasons is *not* explained by local reasons grounding remote reasons. The correlation holds simply because remote reasons are always identical to some local reason. A good consequence of φ is a local reason to φ. Yet that same good consequence is also a remote reason against the alternative.

[24] Especially enthusiastic gratitude goes to Justin Snedegar who commented on multiple versions of this paper at various stages of convolution and to William & Mary who provided the generous research grant and sabbatical that made this paper possible. I am also enthusiastically grateful for the comments provided by Anne Jeffrey, Benjamin Kiesewetter, Daniel Muñoz, as well as the audiences at the 15th Annual *MadMeta* and 13th Annual *RoME* conferences.

Clark, Michael. 1978. “The Meritorious and the Mandatory”. *Proceedings of the Aristotelian Society* 79: 22-33.
Curtis, Barry. 1981. “The Supererogatory, the Foolish, and the Morally Required.” *Journal of Value Inquiry* 15: 311-318.
Dancy, Jonathan. 2004. *Ethics Without Principles*. Oxford University Press.
Dorsey, Dale. 2016. *The Limits of Moral Authority*. Oxford University Press.
Gert, Josh. 2007. “Normative Strength and the Balance of Reasons.” *Philosophical Review* 116: 533-562.
_____. 2004. *Brute Rationality: Normativity and Human Action*. Oxford University Press.
Hurka, Thomas and Shubert, Esther. 2012. “Permissions to Do Less than the Best: A Moving Band.” *Oxford Studies in Normative Ethics* 2: 1-27.
Kamm, Frances. 1996. *Morality, Mortality, Volume II: Rights, Duties, and Status*. Oxford University Press.
Kiesewetter, Benjamin. 2018. “Contrary-to-Duty Scenarios, Deontic Dilemmas, and Transmission Principles.” *Ethics* 129: 98-115.
_____. 2015. “Instrumental Normativity: In Defense of the Transmission Principle.” *Ethics* 125: 921-946.
Massoud, Amy. 2016. “Moral Worth and Supererogation.” *Ethics* 126: 690-710.
Muñoz, Daniel. Forthcoming. “Three Paradoxes of Supererogation.” *Nous*. Manuscript version available at: https://philpapers.org/rec/MUOTPO-3.
Portmore, Douglas. 2011. *Commonsense Consequentialism: Wherein Morality Meets Rationality*. Oxford University Press.
_____. 2008. “Are Moral Reasons Morally Overriding?” *Ethical Theory and Moral Practice* 11: 369-388.
Raz, Joseph. 1990. *Practical Reason and Norms*. 2nd ed. Oxford University Press.
Scheffler, Samuel. 1982. *The rejection of consequentialism: A philosophical investigation of the considerations underlying rival moral conceptions*. Oxford University Press.
Snedegar, Justin. 2018. “Reasons for and Reasons Against.” *Philosophical Studies* 175: 725-743.
_____. 2017. *Contrastive Reasons*. Oxford University Press.
Tucker, Chris. Forthcoming. “The Dual Scale Model of Weighing Reasons.” *Nous*.
_____. 2017. “How to Think About Satisficing.” *Philosophical Studies* 174: 1365-1384.
_____. Manuscript. *The Weight of Reasons*. [Tentative Title.]
White, Stephen. 2017. “Transmission Failures.” *Ethics* 127: 719-732.