Partial Aggregation: What the People Think

Markus Kneer, Juri Viehoff

Abstract

This article applies the tools of experimental philosophy to the ongoing debate about both the theoretical viability and the practical import of partially aggregative moral theories in distributive ethics. We conduct a series of three experiments (N=383): First, we document the widespread occurrence of the intuitions that motivate this position. Our study then moves beyond establishing the existence of partially aggregative intuitions in two dimensions: First, we extend experimental work in such a way as to ascertain which amongst existing versions of partial aggregation (localised vs. global) chimes more fully with moral common sense. Specifically, we document how, in tie-breaking cases, 'irrelevant goods' judgments (Kamm) are just as robust as the original aggregative/non-aggregative pair of judgments that constitute partial aggregation. Second, by pairing laypeople's moral judgments in standard cases with their intuitions about the limits of permissible self-prioritisation, we investigate whether one prominent explanation of why irrelevant claims may not be aggregated (Voorhoeve's 'personal prerogative' argument) can be said to underpin people's intuitions about the (ir)relevance relation of claims in conflict cases. We close with a discussion of our findings' practical and theoretical import and highlight avenues for future research.

Keywords: experimental philosophy, partial aggregation, distributive ethics

Introduction

Conflicts of beneficence are situations where an agent must choose between separate courses of action and, depending on what is done, some people will receive morally significant benefits, whilst others, who could receive them, will not. We face a '*pure*' number conflict where the only difference between courses of action consists in the number of people on each side. For example: on a particular day, a surgeon faces the choice of either conducting life-saving surgery that takes a long time or five equally urgent, equally life-saving ones that take less time. The conflict is 'pure' in that the only difference between separate courses of action consists in the *number of people* that can be benefitted, but not the benefit that each person will receive.[1]

[1] Throughout this article, we assume that all individuals who stand to suffer harm and whose well-being a choice affects are placed symmetrically: they are at equal levels of well-being bar the impairments that can be alleviated, and neither of them has any antecedent moral claim to be treated preferentially (i.e. they are all equally deserving, the agent has no special duties or fiduciary obligations towards some people that they do not have towards all etc.). We also assume that no harm is imposed intentionally.

Whether one is required to let the number of people on different sides count for how we should act in pure number conflicts is of great significance for moral theory, and has been the subject of extensive debate in normative ethics (Taurek 1977; Parfit 1978; Kamm 1993; Hirose 2001; Otsuka 2004; Hsieh, Strudler, and Wasserman 2006; Hirose 2014). Yet far more realistic, and hence more relevant from a standpoint of *practical* ethical decision-making, are '*mixed' or 'impure'* number conflicts. These are situations where an agent must choose between competing options of benefitting others where not only the number of affected people, but also the severity of the impairments to be alleviated (and correspondingly the size of the benefit each person would receive) differ. Mixed conflicts of beneficence are the basic currency of everyday life and healthcare policy: At the micro-level, for example, doctors often must decide whether to treat *one* patient with a *severe impairment* or, instead, treat *several* patients with *lesser impairments*. At the macro level, policy-makers and healthcare administrators typically must reach decisions, against limited budgets, of whether to fund treatment centers that allow addressing *rare* but *severe* diseases or *common* but *less severe* ones.

In actual healthcare policy making, by far the most widely adopted strategy for addressing such conflicts systematically is in terms of cost-benefit analysis, i.e. a form of applied consequentialism. One standard model of this sort is the assessment of options in terms of QALYs/DALYs (Weinstein, Torrance, and McGuire 2009; Loomes and McKenzie 1989). Such a model is *fully aggregative*: For each option whose decision-relevant value V is calculated as the product of severity of the impairment [α] and the number of persons [N] to be benefited ($V_1=\alpha*N_1$), there can exist an alternative option ($V_2=\beta*N_2$) of which it is true that $\alpha>\beta$ and yet, because $N_2>N_1$, V_2 is more choice-worthy than V_1. In other words, *fully aggregative* views hold that we can, in principle, always conceive of an alternative in which the individual impairments that each person stands to have alleviated is less severe, but, because the number of people is larger, we ought to help those with the less severe impairments. Crucially, this holds *no matter how far apart* impairments α and β are in terms of severity.

Full aggregation delivers plausible results in cases where the number of people on different sides of the conflict differ significantly, and impairment levels are relatively close. For example:

> *Death vs Paraplegia:* We can save either (i) one person from death or (ii) one thousand people from paraplegia

Here, it seems clear intuitively plausible that we should follow (ii), even if avoiding *one* death is, to most people, more important than avoiding *one* person's paraplegia (because death is a more severe impairment than paraplegia). But, notoriously, full aggregation also implies that alleviating some relatively unimportant impairment for a very large number of people can outrank the treatment of life-threatening conditions to a few. For example:

> *Death vs Migraines:* We can (i) save one person from death or (ii) save a large number N of people from suffering a mild, short migraine

Provided that N is very large, full aggregation accepts that we should do (ii) (Norcross 1997; Horton 2021). But many philosophers find this implication deeply implausible, in fact so implausible that we ought to reject any fully aggregative view because of it (Tadros 2019).

In response, some theorists have in recent years proposed *partially aggregative* views (Lefkowitz 2008; Kelleher 2014; Voorhoeve 2014; 2015; 2016; Tomlin 2017; Tadros 2019; Gils and Tomlin 2020; Rüger 2020; Steuwer 2021; Mann 2021; Hart 2022; Mann 2022). These are intended to respect the pair of intuitive judgments just described: where the benefits to each person are sufficiently 'close' or 'relevant' to one another (as between death and paraplegia), we choose what to do based on aggregating (relevant) benefits and comparing the total 'deontic value' of different options. Where the individual benefits are not close to each other (as between death and brief migraine), aggregating lesser benefits is impermissible: no number of avoided migraines could prevent us from saving the one life. Partially aggregative views seem *prima facie* attractive because both the aggregative judgment in *Death vs Paraplegia* and the non-aggregative *Death vs Migraines* strike many as very robust, that is, few philosophers are prepared to revise them, even if it is shown that accepting them creates complications and theoretical difficulties.

But far from being a matter of theoretical import only, the truth of partially aggregative moral theories could have significant implications for, amongst other things, healthcare policy: In line with the first, *aggregative*, intuition, governments, in an attempt to enact the correct moral view, might recalibrate policy in such a way as to prioritize the treatment of widespread,

severe impairments over rare, deadly ones. Equally significantly, and in line with the second, *anti-aggregative* intuition, public healthcare provision might refrain from treating minor impairments in cases where invested resources could be applied to the treatment of a single, significantly more serious impairment.

The aim of this article is to apply the tools of experimental philosophy to the ongoing debate about both the theoretical viability and the practical import of a partially aggregative moral theories. We do this by reporting a series of empirical studies. The first one provides evidence in favour of partially aggregative theories. In this way, we advance our understanding of what constitutes moral common sense in this domain. However, we go beyond establishing the existence of *partially aggregative* intuitions in two regards: First, we explore which amongst existing *versions* of partial aggregation chimes more fully with moral common sense. Our discovery is that 'irrelevant goods' judgments in tiebreaking cases are just as robust as the aggregative/non-aggregative pair of judgments in the other central cases. As we explain below, this matters for determining which version of partial aggregation more closely tracks common sense. Second, we test whether one prominent *explanation* of why irrelevant claims may not be aggregated (Voorhoeve's 'prerogative' view) is consistent with people's relevance-intuitions.

The article is structured as follows: Section one lays out one prominent partially aggregative view and its underlying motivation. Section two present some criticisms of partial aggregation and documents how defenders have responded by developing more determinate variants. Section three highlights the significance of claims about 'common sense' and 'intuitive plausibility' for questions of partial aggregation, distinguishing between the theoretical and practical-institutional relevance of empirical findings in this domain. Section four discusses existing findings in the empirical literature on the topic of harm aggregation in conflict cases and highlights some of their shortcomings. Section five presents our findings regarding the general prevalence of the intuitions that underpin partially aggregative views. In order to better understand which version of 'partial aggregation' is most consistent with common sense, section six extends existing empirical research on the topic by investing anti-aggregative intuitions in related cases regarding the

distribution of fair chances of treatment and Kamm's 'irrelevant goods.' Section seven presents our finding about whether there is evidence that anti-aggregative common sense judgments derive from, or are congruent with, widely held common sense assumptions about the existence of (and limits to) personal prerogatives. We show that, in 'involved' conflict cases (that is, cases where the decision-maker themselves stands to suffer harm), folk intuitions about where to draw the line between permissibly favouring oneself vs. others do not significantly correlate with intuitions about the (ir)relevance of claims in 'uninvolved' cases. We conclude that although the anti-aggregative judgments that support partially aggregative moral theory are widespread and robust, much theoretical work remains to be done.

1. Partial Aggregation in Moral Theory

Though earlier statements of partially aggregative views can be found in the works of both Thomas Scanlon (1998) and Frances Kamm (1993; 2008), much of the current debate is indebted to Alex Voorhoeve's effort to systematically develop such an account. His conception of partial aggregation, which he calls 'Aggregate Relevant Claims' (ARC) takes the following form (2014, 66):

1. Each individual whose well-being is at stake has a claim on you to be helped. (An individual for whom nothing is at stake does not have a claim.)
2. Individuals' claims *compete* just in case they cannot be jointly satisfied.
3. An individual's claim is *stronger*:
    a. the more her well-being would be increased by being aided; and
    b. the lower the level of well-being from which this increase would take place.
4. A claim is *relevant* if and only if it is sufficiently strong relative to the strongest competing claim.
5. You should choose an alternative that satisfies the greatest sum of strength-weighted, relevant claims.

Voorhoeve's idea is that claims with varying strength capture the relative importance that each person's wellbeing should have from the standpoint of the decision-maker in conflict situations. In combination, clauses (4.) and (5.) of ARC can yield the pair of intuitions that underpin the widespread judgments in *Death vs Paraplegia* and *Death vs Migraine*: it is *because* paraplegia is a claim that is relevant to death, but migraine isn't, that we should sum instances of paraplegia-prevention but not instances of migraine-prevention when the alternative is that of preventing a death.

What Voorhoeve's account makes explicit is the idea that in conflict assessment two distinct kinds of considerations are at play: One type, which corresponds to the notion of a claim's *strength*, is essentially aggregative: we reflect on which option is more choice-worthy by considering those claims that can be satisfied together as an aggregate and, in reaching our decision, we reflect on which option's aggregate of claims *outweighs* another option when they conflict. The other type, corresponding to the notion of a claim's *relevance*, is essentially non-aggregative: we reflect on which kind of claim when it competes with a stronger claim, lacks any weight in our deliberation about which claims to satisfy. This is not because some claim (say to avoid a mild migraine) has no weight at all in general, but because the fact that this weaker claim competes with a stronger claim (say to have one's life saved) *cancels* that weaker claim.

Of course a decision procedure, even an intuitive and elegant one like ARC, is not of much value unless we can point to reasons for why we should follow it. Amongst defenders of partially aggregative accounts, two positions can be discerned (Tadros 2019, 178; Brown 2022, 735). One position takes a non-reductive stance: the fact that some weaker claims are cancelled in the presence of much stronger claims is morally basic and cannot be further explained or derived from additional moral considerations.[3] By contrast, others have suggested that the condition of moral (ir)relevance can be further justified by bringing into play other aspects of non-consequentialist ethics (Voorhoeve 2014; Hart 2022). Most prominently, Voorhoeve has proposed that what 'grounds' the relevance/irrelevance boundary is explained by appeal to an individual's personal prerogatives against maximizing the good in cases where the deciding agent's own wellbeing is at stake (Voorhoeve 2014, 74). The idea is this: Whilst what we ought to do morally is shaped by the impersonal, aggregative perspective of which action would produce the best outcome, we are sometimes permitted to grant our own wellbeing some stronger concern in distributive conflicts. So, for example, it is morally permissible to refuse losing a leg in order to save another person's life, even

[3] As Tadros puts it: "The contextual interaction which determines whether a duty-grounding type of fact has duty-grounding force in a particular case is morally basic in just the same way that it is basic that duties are grounded in duty-grounding facts." (Tadros 2019, 178)

if saving a life is impersonally better than saving a leg. But morality also sets limits to how much additional weight we are permitted to give to our own wellbeing: we are not, for example, permitted to let another person die if the cost of avoiding this outcome is for us to suffer a mild migraine.

Voorhoeve connects this element of non-consequentialist moral common sense to partial aggregation: the deeper reasons for where and why the boundary between relevant and irrelevant harms should be drawn in conflict cases flows from (or is explained by) the benefits we may permissibly claim for ourselves in defiance of the greater good. Thus, the loss of two legs is 'relevant' to death in mixed conflict cases *because* we are each permitted to refuse to lose our legs in order to save another person's life. So this explanation 'extrapolates' from personal prerogatives: aggregation is permissible in cases where the decision-maker (correctly) judges that each person with the lesser harm would have a personal prerogative that permitted her to refuse rescuing those that would otherwise suffer the more serious harm. Conversely, we can think about cases of 'irrelevant' smaller harms that each person would suffer as instances where each of those suffering the smaller harm would *not* have a personal prerogative to favour herself: when a life is at stake and we only stand to avoid a migraine, we have no right that we not suffer that smaller harm (Voorhoeve 2014, 68–70; 2017, 149–51).

2. Criticism and Variants

2.1 Criticism of Partially Aggregative Accounts

Partially aggregative accounts have not gone unchallenged: one structural difficulty they face is that they must, so to speak, fight a two-front battle. Against thoroughly *anti-aggregative* views (according to which we must always improve the position of the person with the strongest individual claim) they have to explain why aggregation is permissible when potential benefits to beneficiaries are 'close enough', that is, relevant. Yet, against advocates of *full aggregation*, partial aggregation theorists must argue that some instances of interpersonal aggregation are impermissible.

Specific challenges that have been levelled against partial aggregation are the following: First, that the view leads to inconsistent (and hence irrational) or deeply implausible results once we apply it to more realistic scenarios that (i) go beyond *binary* choice situations, and/or (ii) contain

'mixed groups' of individuals with impairments of differing degrees of severity under each of the conflicting options, and/or (iii) involve risk of benefitting/failing to benefit rather than certainty. Once we look at cases that involve multifold options, or mixe groups, or contain risk, partial aggregation is said to violate plausible principles of rational decision-making, including transitivity, basic contraction consistency, and dominance (Halstead 2016; Tomlin 2017; Horton 2018; 2020; 2021), for detailed responses, see: (Voorhoeve 2014; 2015, 201; Tadros 2019; Rüger 2020; Mann 2021; 2022).

Second, critics have argued that partial aggregation does not yield determinate results that could be translated into real policy: While fully aggregative views have a clear and established methodology that moves from moral theory to policy in the form of established QALY/DALY rules, partial aggregation remains vague in relation to two issues that are needed for any actual implementation (Fried 2020): First, the 'valuation' question how those harms that can permissible be weighed against each other in conflict cases are to be understood (e.g. how many people with paraplegia outweigh one death?). Second, the 'relevance' question: which impairments are or are not relevant to each other?[4]

2.2 Variants of Partial Aggregation

Even if the responses to the theoretical criticisms just mentioned are persuasive, they reveal that, rather than being a single view, partial aggregation is a family of views and principles and, depending on which variant is chosen, some criticisms pose more powerful challenges than others. We can use the challenge from 'mixed groups'/'stage-wise decisions' to make this point and distinguish different variants.

That Voorhoeve's ARC describes one particular specification within a broader family of partially aggregative views is easily understood if we think about how one should understand the crucial criterion of 'relevance', specifically, what it takes for a claim to be relevant/irrelevant. Relevance is a *relational property*: For some claim C_2 to be relevant in a conflict, it must be relevant to *some other claim* C_1 in a given choice situation. But to understand

[4] For a detailed analysis of how partial aggregation could/is implemented in healthcare priority setting contexts, see: (Voorhoeve 2018b).

how to evaluate relevance in cases where there are more than one (kind of) claim on each side, we need a more perspicuous account of this relation (Tomlin 2017; Mann 2021; 2022). What exactly need a claim be relevant to in order to be relevant? Mann and Tadros have both convincingly argued that, in order to be relevant, a claim must be sufficiently strong relative to the strongest claim *with which it competes* (Tomlin 2017, 239; Mann 2021, 103).

The main worry about this precisification of the relevance relation is that it seems to violate the following plausible principle when we move to multiple options / stagewise cases:

> *Equal Consideration for Equal Claims:* All claims of equal strength ought to be given equal weight in determining which group to save.[6]

Faced with this and related challenges, defenders of partial aggregation have developed increasingly complex views (Tadros 2019; Mann 2021; 2022; Gils and Tomlin 2020; Hart 2022; Rüger 2020). One prominent one, put forward by Tadros introduces a significant distinction: whether, according to the best version of the relevance view, the (ir)relevance relation should be considered *local* or *global*. Local relevance implies that a claim whose choice-worthiness can be silenced as a result of it being irrelevant to a stronger claim with which it competes is only 'locally' silenced: it may still 'counterbalance' other competing claims in the overall choice situation. Tadros favours this 'localized' understanding of (ir)relevance. The resulting view—further refined by Gils and Tomlin (2020)—is a form of sequential claim matching: Claims of the highest relevance level R1 are *matched/cancelled* with claims that are relevant to it; remaining claims are then taken to establish the new highest relevance level R2 and remaining claims are matched against it, and so forth, until one option is the only one with remaining claims.

---

[6] This can be seen from considering Tomlin's *Case 5* (p.242): "*Stage 1*: You can save Group A or Group B. Group A consists of one dying person. Group B consists of ten people facing severe impairment. The Relevance View requires you to save Group B. *Stage 2*: One person facing mild impairment is added to Group A. One billion people facing mild impairment are added to Group B." If the choice of B over A at stage 1 was very close, then this would mean that you should now switch to A even though 999,999,999 more weak impairment were added to B—the reason being that only the single added weak impairment claim is relevant according to *anchor by competition*.

Our goal here is not to assess the merit or demerit of 'global' vs. 'local' versions of partial aggregation. But what becomes clear if we compare this form of local relevance to its competitor is that, importantly, it runs counter to one important other moral intuition that one might have thought to go hand in hand with the anti-aggregative judgment in *Death vs. Migraines*. This is the idea that the alleviation of some harms is too insignificant to play a role in our decisions whom to save. The best case to illustrate this is Frances Kamm's sore throat case (1993, 146):

- *Sore Throat*: I have enough medicine to save either A or B from a lethal illness. If I save B I will have a little medicine left which I can use to cure C's sore throat. Otherwise, C will suffer the sore throat for a week.

According to *locally restricted aggregation*, having matched the two lethal illnesses between A and B, C's sore throat should be allowed to determine what we do (it is, after all, an 'unmatched' claim). But this seems to violate the anti-aggregative intuition that minor improvements (Kamm calls them "irrelevant goods" (1993, 144)) cannot prevent us from giving each of A and B an equal chance of survival by, for example, flipping a coin between them (Hart 2022).

To sum up then, recent years have seen the development of variants of partially aggregative views. The most sophisticated ones have opted for 'anchor by competition' (rather than overall strength) and have shifted from 'global' relevance to a 'local' relevance condition. According to this latter approach, because 'lower order' claims are only *locally* cancelled, parity of strength-weighted claims amongst more serious claims can still leave it up to 'lower order' claims to play a decisive role in determining what to do.

3. Appeal to Intuitions and the Significance of 'moral common sense'

In the debate about partial aggregation, appeals to intuitive judgments and claims that moral theory should respect common sense abound. Just to offer some examples: Lazar aims to develop "principles that can preserve common sense in *Life for Headaches* […]" (Lazar 2018, 118), whilst Kelleher contends that when it comes to (ir)relevance of some harms "[nothing] other than unshakeable intuitions can be marshalled to support it." (Kelleher 2014, 388). Rüger holds that fully aggregative views' inability to "explain the common

intuition in *Death vs. Headaches*" speaks prima facie against accepting them (Rüger 2020, 455), whilst Tadros states that "only restricted [partial] aggregationist views have intuitive implications in both cases, and this is some reason to believe that some such view is true" (Tadros 2019, 172).

The comparatively 'heavy' justificatory load that is imposed on intuitive judgments about particular cases stems in part from the absence of fully worked out theoretical accounts of partial aggregation as a component in non-consequentialist ethics (see section 7). We think that, for this reason, understanding what *actual* 'moral common sense' and 'widely held intuitions' amount to matters within this debate. Specifically, we appeal to two kinds of considerations that matter here:

First, as authors on all sides of this debate implicitly or explicitly rely on the method of reflective equilibrium (whereby the theorist seeks to develop principles that confirm their intuitive judgments about particular cases and vice versa), each side claims that intuitive judgments do have at least *some* prima facie authority in theory construction.[7] But if it turned out that intuitive judgments by theorists differed dramatically from those of the wider population , or if empirical studies showed that judgments about core cases were shaped by idiosyncratic factors (such as e.g. differing cognitive or ideological dispositions, then may we not justifiably worry about the normative status of the principles that are derived from reflective equilibrium? Miller couches the relevant worry that we should have when moral theorists' judgments are disconnected from prevailing folk intuitions in terms of 'bias' and 'theoretical contamination', that is, seemingly 'intuitive' judgments are post-hoc rationalisations of theoretical commitments (Miller 2020, 274; Baderin 2016; Pölzer 2023).[8]

---

[7] For a detailed discussion of empirical research's importance to (narrow and wide) reflective equilibrium in political philosophy, see: (Allard and Cova forthcoming) and Pölzler (2023).

[8] There is a debate regarding whether experimental-philosophical work matters to philosophical debate in the first place. The two main lines of argument regard (i) whether lay people have the relevant *expertise*, and (ii) whether the empirical studies generate *reflective* judgments rather than merely crude shots from the hip; for discussion see, (Kauppinen 2007; Liao 2008; Horvath 2010) and (Horvath and Wiegmann 2016). As empirical work has shown, neither objection is appropriate . The judgments of experts standardly do not vary much from lay subjects – be it with regards to questions in ethics (see e.g. (Schwitzgebel and Cushman 2012; 2015; Wiegmann, Horvath, and Meyer 2020; Horvath and Wiegmann 2022), epistemology (Gerken 2017)) or central philosophical concepts in law (Kneer and Bourgeois-Gironde 2017; 2018; Hannikainen et al. 2022). Furthermore, as (Kneer et al. 2021)have shown, encouraging reflective judgment and controlling for reflective-analytic dispositions (Type 2 thinking) has no relevant impact on

It seems that defenders of partially aggregative views *especially* should worry about the status of intuitions, given how central the appeal to intuition about cases like *Death vs. Migraines* figures in their overall reasoning and how close to 'moral rock bottom' these are taken to be. Or put differently: defenders of *full aggregation* have, amongst other things, sought to undermine the confidence we should have in judgments like *Death vs. Migraines,* for example by explaining that we should be sceptical about our intuitions regarding very large numbers (Parfit 2013, 248; Broome 2004, 56). They have done so because full aggregation contradicts what most theorists take to be moral common sense in cases like *Death vs. Migraines*. The case for partial aggregation would weaken considerably if we were to lose our conviction in this particular moral judgment, for example as a result of realising that it is only widely held amongst a particular social group, or theorists of a particular theoretical conviction and that, perhaps, the fact that it is held can be explained in a debunking manner.

Beyond its theoretical significance, moral common sense matters greatly for what follows *practically* for the political domain of healthcare policy: As Luptakova and Voorhoeve (forthcoming) point out, being more theoretically appealing than competing theories may not be enough to establish that a moral theory should govern actual healthcare policy. Actual policy must also satisfy demands of *political legitimacy* and such legitimacy depends at least to some degree on—perhaps *reasonable*—actual public acceptance (i.e. if many citizens could not accept principles embedded in public institutions, this would reduce legitimacy). (Luptakova and Voorhoeve forthcoming, 14; Lindauer 2020).

Such a concern with political legitimacy, should, we think, be of particular importance to defenders of partially aggregative views, at least when compared to defenders of full aggregation. Here is why: It is plausible to believe that both, the requirement of political legitimacy just mentioned, and the reasons for refusing full aggregation spring from something like a common source, namely a conviction that social arrangements must be justifiable to each person, and that, moreover, a policy's bringing about the

---

judgments elicited by experimental philosophy studies. In short, the *expertise defense* and the *reflection defense* of armchair philosophy are largely misconceived.

best consequences is neither necessary nor sufficient for its being justifiable to each (Waldron 1993; Scanlon 1998). But despite the fact that partially aggregative views derive a considerable aspect of their appeal from the fact that they alone manage to respect moral common sense when it comes to conflict cases like *Death vs. Paraplegia* and *Death vs. Migraines*, few empirical studies have attempted to systematically uncover what 'common sense' in relation to such scenarios actually is. And, as we show below, those studies that do exist suffer from a number of problems.

4. Existing Empirical Studies and their Shortcomings

The existing empirical literature reports evidence in favour of *non-aggregative approaches*, according to which certain types of moderate harms are deemed irrelevant when pitted against severe harms. According to Ubel et al. (1996), Damschroder et al (2007 Study 1) and Cowell et al. (2015) this is the view of roughly one third of the population, according to Pinto-Prades and Lopez Nicolas (1998) over half. However, existing empirical studies suffer from a number of flaws that fall into one (or more) of (i) statistical insignificance, (ii) lack of analytic clarity in relation to the core theoretical issues at stake, and (iii) a too narrow focus on a small set of highly stylized simple cases of conflict choices. In their critical discussions of the literature, Voorhoeve (2018a, 130) and Luptakova and Voorhoeve (forthcoming) point to a number of shortcomings of different studies that look at number vs. size of benefit (significance of harm avoidance) trade-offs. The relevant points of criticism were the following: A first and second worry are statistical, in that existing experiments relied on insufficient sample sizes and, second, they used non-representative groups of subjects (economics undergraduates in Cowell et.al. (2015), philosophers and philosophy students likely already interested in the topic of aggregation in distributive ethics in Rüger (2015)). These experiments cannot be said to truly establish much about the actual prevalence of such views in society.

Third, existing experiments suffer from important framing concerns: For example, some questions were posed in such a way that subjects were prompted to determine the axiological value of resulting states of affairs rather than evaluating the relative deontic status of alternative choice options (what number of less severe harms cured would provide 'the same benefit' as

averting ten deaths? vs. what should a person faced with this kind of choice do?) Problematically, a certain choice could provide greater benefits (i.e. be more valuable than another), yet still be morally impermissible; the question thus requires framing in terms of a suitable moral vocabulary. Our experiments consistently rely on the notion of what an agent *should* choose.

Finally, fourth, those experiments (Pinto-Prades and Lopez-Nicolás 1998) that did frame the question set non-axiologically failed to clearly define the consequent health benefits that accrue to members of each group (Voorhoeve 2018a, 129). Our experiments also corrects for this shortcoming.

5. *Experiment 1*: Prevalence of Partially Aggregative Intuitions

The aim of our first experiment was to establish, against the background of insufficient statistically sound, robust evidence, (a) *whether* there is indeed widespread support for the pair of aggregative and non-aggregative intuitions that make up the partially aggregative view. Moreover, (b) we wanted to explore whether holding partially aggregative views correlates with a number of personal characteristics.

5.1 Participants

127 Participants were recruited on Amazon Mechanical Turk to complete a paid Qualtrics online survey. The IP address location was restricted to the USA. Participants who failed an attention check, or whose responses violated the comprehension check were excluded. 77 participants remained (age M=39 years, SD=12 years).

5.2 Methods and Materials

All participants completed three types of tasks: First, they had to rate the severity of four different medical conditions (adapted from Ubel et al. 1996, experiment 1). Second, they faced six comparative scenarios, in which they had to judge tradeoffs between the four types of medical conditions. Third, and following similar explorations of experimental philosophers regarding participants' cognitive and moral dispositions (Pinillos et al. 2011; Byrd 2021; Kneer et al. 2021), all participants responded to the 10 item Rational Experiential Inventory (Epstein et al. 1996; Pacini and Epstein 1999), the 20 item Moral Foundations Questionnaire (Graham, Haidt, and Nosek 2009; Graham et al. 2011) and a demographic questionnaire to investigate potential

cognitive, ideological, and moral predictive factors of their tradeoff judgments. The order of the three blocks was fixed, the order of the individual sub-tasks in the first and second blocks was randomized. For the first block, we first familiarized participants with four types of illness (presented without the titles here in bold):

> **Ganglion Cyst:** Imagine that you are in the following state of health: You have a ganglion cyst on one hand. This cyst is a tiny bulge on top of one of the tendons in your hand. It does not disturb the function of your hand. You are able to do everything you could normally do, including activities that require strength or agility of the hand. However, occasionally you are aware of the bump on your hand, about the size of a pea. And once every month or so, the cyst causes mild pain, which can be eliminated by taking an aspirin.
>
> **Chronic Knee Damage:** Imagine that you are in the following state of health: You have knee damage that prevents you from engaging in athletic activity and frequently causes your knee to give out on you when you walk. In addition, your knee aches for approximately one hour every day, to the point that it is hard to concentrate. The rest of the time, you are able to function normally.
>
> **Meningioma:** Imagine that you are in the following state of health: You have growth in the tissue lining the brain called a meningioma. It causes you to have constant headaches. The pain is often severe. It can be decreased with medicines, but it cannot be eliminated without interfering with your ability to concentrate. You must take pain medications to sleep at night. The meningioma is not cancerous and will not affect how long you live.
>
> **Acute Appendicitis:** Imagine that you are in the following state of health: You have acute appendicitis. Acute appendicitis is an inflammation of part of the colon that, if untreated, will cause death within hours to days. With treatment, people can be cured of appendicitis and will return to their previous states of health.

Following the four specifications (each presented on a separate screen), participants responded to the following question on a Likert scale: "On a scale from 0 to 100, where 0 is as bad as death and 100 is perfect health, how would you rate this condition?"[9]

For the second block, participants faced six contrastive pairs (each presented on a separate screen in randomized order). To give an example, the Appendicitis v. Ganglion Cyst condition read:

> Dr Smith manages a hospital in a remote area which has limited resources. It can only afford one new operation room, specializing in one type of illness. Dr Smith has to choose from the following two options:
>
> - **OPTION A**: An operation room which could treat ten patients per quarter with acute appendicitis (if untreated, acute appendicitis causes death within hours).
> - **OPTION B**: An operation room which could treat a larger number of patients per quarter with ganglion cysts (the pea-sized bulges on top of the hand, which cause mild pain once a month).

[9] Data, Qualtrics files and an anonymized Appendix are available under https://osf.io/jpvyf/?view_only=adba5d06980b4476afb03e7ab1226cc8.

Q: How many people with ganglion cysts would the hospital need to cure per quarter so that Dr Smith should choose OPTION B (an operation room for ganglion cysts) over OPTION A (an operation room that could cure ten patients with acute appendicitis per quarter)? Please enter a number OR select 'Dr Smith should always choose Option A'.

- Dr Smith should choose option B if the number of cured ganglion cyst patients per quarter were at least _____.
- No matter how high the number of cured ganglion cyst patients, Dr. Smith should always choose option A (an operation room that could cure ten patients with acute appendicitis per quarter).

In the other five contrastive pairs, the prompts and responses were adapted to include comparisons between all other medical conditions (cf. Appendix 1 for complete vignettes).

5.3 Results

The manipulation check was successful: A one-way ANOVA (type III) revealed a significant effect for the type of illness on the perceived severity of the condition ($F(2.40,182.13)=12.68$, $p<.0001$, $\eta_p^2=.42$, a large effect.) The mean severity ratings for the four health conditions (Figure 1) differed significantly from one another (paired samples t-tests, all *ps*<.001, all corrected *ps*<.001, detailed results for this entire section are reported in the Appendix). The effect across conditions was at least medium in size (*d*>.44 Meningioma v. Knee damage) and substantial for most contrasts (all other *ds*>.78).

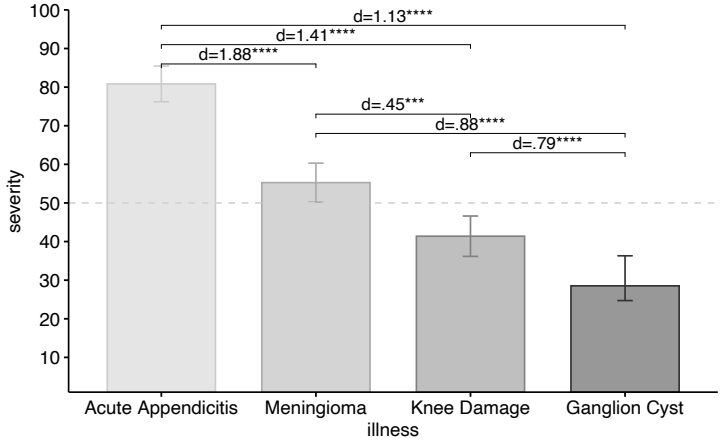

Figure 1. Mean severity for each medical condition and pairwise comparisons. *p<.05; **p<.01; ***p<.001; ****p<.0001.

The percentage of participants who responded with 'no number' in each of the six possible comparisons is displayed in Table 1. All proportions of 'no

number' responses significantly exceeded chance (binomial tests, *ps*<.0031) except the Meningioma v. Knee Damage contrast (56%, *p*=.36). This means that at least half of the participants judged that no number of patients with the less severe condition treated could outweigh the treatment of one patient with the more severe condition. A Cochran's Q test suggests that the ratios of 'no number' responses differ significantly across pairs of conditions (*p*<.001).

	Meningioma	*Knee Damage*	*Ganglion Cyst*
Acute Appendicitis	68%	82%	94%
Meningioma		56%	82%
Knee Damage			86%

Table 1. Percentage of participants who responded that no matter how high the number of patients per quarter with the condition in the row, the doctor should always cure the ten patients per quarter with the conditions in the column.

Overall, the proportion of participants manifesting *partially* aggregative dispositions, i.e. the 58% choosing "no number" in some though not all contrasts significantly exceeded that of those manifesting radically anti-aggregative dispositions (39%, choosing "no number" in all contrastive pairs) and perfect aggregators (3%, never choosing "no number"), binomial tests, *ps*<.001.

The frequency of 'no number' responses was not found to significantly correlate with either of the REI subscales (NFC: Need for Cognition; FI: Faith in Intuition), any of the MFQ subscales (harm, fairness, ingroup, authority, purity, progressivism). Furthermore, no evidence of a correlation with level of education or philosophical expertise was found. The frequency of 'no number' responses correlated positively with gender (females opting more frequently for 'no number', *r*=.18, *p*<.01) and age (older participants opting more frequently for 'no number', *r*=.12, *p*<.01), see Table 2. The lack of interesting correlations regarding the subscales of the REI and MFQ is consistent with a similar absence in recent experimental-philosophical work (Kneer et al. 2021, for further discussion, see Byrd, 2021).

Variable	***r***	***p***
Education Level	.08	.0917
How much philosophy	-.07	.1257

Variable	*r*	*p*
Gender	.18	.0001
Age	.12	.0075
MFQ PURITY	.05	.2965
MFQ PROGRESSIVISM	.04	.3991
MFQ INGROUP	-.03	.5463
MFQ HARM	.08	.0792
MFQ FAIRNESS	.02	.6481
MFQ AUTHORITY	.04	.3949
FI SCORE	.03	.5310
NF SCORE	-.00	.9980

Table 2. Correlations with the choice of "no number would outweigh curing one patient with condition A"

5.4 Discussion

Our experimental work shows that there is robust support for partially aggregative moral views – in our particular study, 58% of our participants judged in a partially aggregative fashion. Furthermore, the very high number that supports the anti-aggregative intuition where impairments are very far apart (94% for *Acute Appendicites* vs. *Ganglion Cyst*) is quite surprising. It therefore seems that moral 'common sense' does in fact overwhelmingly mirror one part of partial aggregation, namely the 'non-aggregative' judgment when impairments are far apart in terms of severity. This strongly speaks against the assumption that full aggregation is a widely shared view.

6. *Experiment 2*: Which version of partial Aggregation?

Recall from our discussion in section two that two versions of partial aggregation can be distinguished in terms of whether minor benefits can act as tiebreakers in conflict cases in which much more important health benefits (e.g. life or death) are at stake: Whereas *locally restricted* partial aggregation views allow that such minor benefits may sway us in the direction of one course of action where higher-order benefits are balanced in a conflict, *global version* of partial aggregation deny this. Our aim in Experiment B was to determine whether moral common sense aligns with one of these judgments. We do this by testing the prevalence of aggregative vs. equal chance favouring reasoning in cases like *Sore Throat*.

6.1 Participants

We recruited 115 participants on Prolific. In line with the preregistered criteria,[10] participants who failed an attention check or took less than twenty-five seconds to answer the main questions (including reading the prompt) were excluded, leaving a sample of 109 participants (female: 53%; age M=43 years, SD=14 years, range: 19–76 years).

6.2 Methods and Materials

The experiment took a between-subjects design in which we manipulated a single factor: The amount of extra utility the agent can generate by choosing one action over another, i.e. curing a mild headache v. a chronic headache. The scenario was loosely based Kamm's *Sore Throat Case* (1993, p. 146). Having completed an attention check, participants were randomly assigned to one of two conditions. The scenarios read:

Headache

Sally is in the jungle and comes to a fork in the path. She knows that Ms Smith is on the left, and Ms Jones is on the right. Both Ms Jones and Ms Smith have been bitten by the King Cobra, whose poison is lethal. They will die if they do not receive an antidote very soon. Sally has only a single unit of antidote to Cobra bite in her pocket.

Ms Smith is with her husband, who is in good health. Ms Jones is also with her husband. He cannot take the heat that well, and currently suffers from a mild headache. This is an ordinary headache, which will disappear after 30 minutes by itself. Sally also has one aspirin in her pocket, with which she can treat the headache of Mr Jones.

If Sally goes left, she can save Ms Smith's life. If she goes right, she can save Ms Jones' life and cure Mr Jones' mild headache with the aspirin.

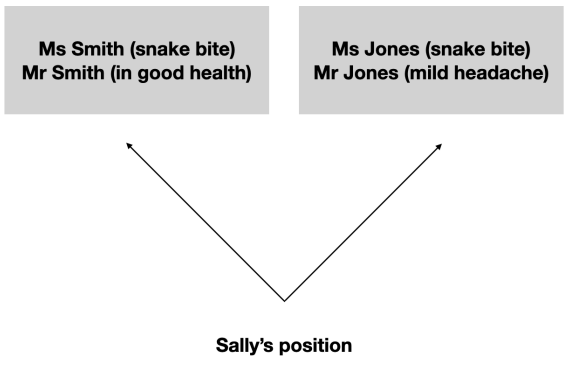

[10] The preregistration can be found under https://aspredicted.org/4YW_H5C

Q. Which of the following options should Sally choose?

- Sally should go left and save the life of Ms Smith. Ms Jones will die, and Mr Jones will suffer from a mild headache for 30 minutes.
- Sally should go right, save the life of Ms Jones and cure Mr Jones from a mild 30-minute headache. Ms Smith will die.
- Sally should flip a coin to determine whom to save. If the coin lands heads, she goes left. If it lands tails she goes right.

In a second condition, Mr Jones does not suffer from a mild headache, but has been bitten by a mosquito which transfers a chronic disease, *Tolerabilis*, that causes a mild headache about once a week (it has no other effects). If treated immediately with an antidote – which Sally has in her pocket – it will not become chronic. The question and response options, adapted to the situation, are the same.

The order of the responses was randomized. Following the main task, participants had to rate the three medical conditions (cobra bite, chronic mild headaches, a single mild headache) to ensure that the participants ranked them in the expected way. Participants also had to complete a brief demographic questionnaire.

6.3 Results

Our manipulation check confirmed that mean ratings for death by snake bite was considered significantly more severe than *tolerabilis*, the chronic headache, which in turn was deemed significantly more severe than a mild headache, see Figure 2.

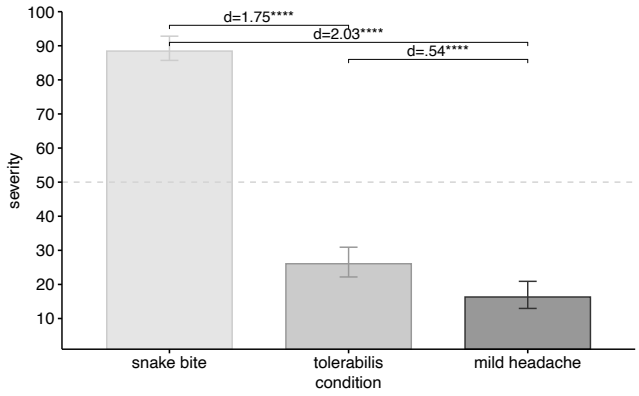

Figure 2. Mean severity for each type of condition and pairwise comparisons. *p<.05; **p<.01; ***p<.001; ****p<.0001

A Pearson Chi-squared test revealed a significant effect of added utility on choice proportions (for N=109, $\chi^2(2) = 13.492$, $p = .0012$). The results are

graphically represented in Figure 3 (see also Appendix, section 2). Whereas in the *mild headache* condition the proportion of participants who prefer a coinflip (55.2%) significantly exceeds chance (binomial test, $p<.001$), the same choice is below chance in the *tolerabilis* (chronic headache) condition (21.6%, no significant difference, $p=.1024$). In the latter condition, going right, so as to save the life of one and cure the chronic headache of a second person (66.7%), was by far the most popular response, significantly exceeding chance ($p<.001$). For details, see Appendix, Section 2.

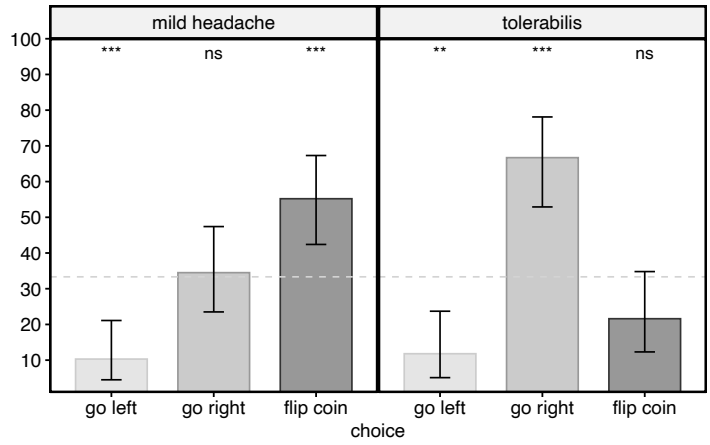

Figure 3. Proportion of choices. Error bars denote 95% confidence intervals. Significance levels for the difference from chance (33%). *p<.05; **p<.01; ***p<.001; ****p<.0001.

In the *tolerabilis* condition, we see a significant and medium-sized correlation between age and the willingness to flip a coin ($r=.35$, $p=.015$) All other correlations were nonsignificant, see Table 3.

added utility	**variable**	***r***	***p***
mild headache	age	.19	.1429
	gender	-.01	.9141
tolerabilis	age	.35	.0150
	gender	.19	.1857

Table 3. Correlations with probabilities of choosing "flip coin" grouped by added utility.

6.4 Discussion

The study confirms the intuition common among non-consequentialist philosophers, that there can be 'irrelevant goods' (or 'irrelevant utilities') in 'mixed' conflict cases, that is, cases where the severity of the harm to be suffered differs amongst potential victims: A majority of participants thought

it more important to give people who each stand to suffer a very serious harm equal chances of having their plight addressed than opting for the option that generates most value because we can also provide a small benefit to a third person. As the severity of the additional impairment that can be alleviated goes up, subjects increase the weight that this lesser impairment has on the decision, i.e. they become increasingly minded to maximize overall health-related outcomes. Conversely, subjects are increasingly unlikely to distribute fair (equal) chances to those facing life or death by flipping a coin.

The upshot of this is not only that there is widespread support for an additional anti-aggregative judgment in the population, but also, interestingly, that some recent accounts of locally restricted aggregation (Tadros 2019; Gils and Tomlin 2020; Hart 2022) come at a significant cost in terms of matching people's overall set of non-consequentialist intuitions: this is the case because following these authors' 'local only' relevance variant of partial aggregation, Mr. Jones' headache is an 'unmatched' claim and would, therefore, mandate that we always save the Jones'.

7. *Experiment 3*: Support from other elements of non-consequentialist common-sense morality?

In section one we explained that even though there now is a significant range of partially aggregative accounts seeking to make sense of various aggregative and non-aggregative intuitions by way of proposing decision-rules, only few authors have attempted to argumentatively support partial aggregation by going beyond the mere appeal to 'moral bedrock' intuitions about cases. In our final experiment, the aim was to assess whether patrial aggregation, at a justificatory level, finds support in other parts of a wider non-consequentialist outlook or whether the fallback non-reductionist position is the best we can hope for.[11]

As Voorhoeve presents the most fully developed attempt that goes beyond there mere appeal to intuitive judgments, our experiment was specifically designed to assess his variant of reductionist reasoning, which, recall from section one, grounds judgments about the relevance relation in

[11] For a general outline of the features of non-consequentialism, see: (Kamm 2008 chapter 1)

views regarding our permissible *personal prerogatives* (a further and distinct element of non-consequentialist moral common sense).

To test this, we engaged in a within-subject experiment where we compared the 'breadth' of the relevance condition that people propose/accept in 'uninvolved' scenarios (i.e. scenarios where the decision-maker's interests are not at stake) with subjects' reported views on permissible levels of self-prioritization in 'involved' conflict cases (i.e. cases where the agent must choose whether they themselves or somebody else will suffer some loss to wellbeing).

7.1 Participants

We recruited 141 participants on Prolific. In line with the preregistered criteria,[12] participants who failed an attention check or took less than fifteen seconds to answer the first question (including reading the prompt) were excluded, leaving a sample of 138 participants (female: 51%; age M=43 years, SD=14 years, range: 18–92 years).

7.2 Methods and Materials

The study had a single condition which came in two steps. In a first step, participants read about Dr. Smith, who manages a hospital nearby a sawmill. Due to the latter, there are *many* patients with hand injuries (severed fingers) and very few potentially deadly head traumas from accidents. Extra resources have become available which can be used towards an increased treatment of *many* fingers per quarter or *one* head trauma patient (who would otherwise die). Participants must decide what Smith should do:

> **[A – Aggregation]** If the numbers that could be saved over time is very high, Smith should extend the hand surgery facilities (sacrificing one head trauma victim's life per quarter).
>
> **[B – No Aggregation]** No matter how high the number of fingers saved, Smith should devote the resources to saving one head trauma patient's life per quarter.

We will call the first step *Uninvolved Judgment* task. It contrasts with Step 2, the *Involved Judgement* task (always presented afterwards), where judgments that involve personal interests or costs to the participants' wellbeing are solicited. The prompt and question read:

---

[12] https://aspredicted.org/6CD_W6B

Burning Building
Jones is trapped in a burning building and is unable to escape due to debris blocking their path. You are the sole person around. The only way for you to save Jones is to use your body to break into the house and clear a path for them to escape. This will cause you to suffer some health damage.

Q.: In order to save Jones' life, what kind of sacrifice are you morally required to make (if any)? (Please tick the *most serious harm* you must morally accept so that Jones will live.)

Responses are collected on a nine-point Likert scale labelled thus: 1: nothing; 2: a mild, brief headache; 3: a bruised ankle; 4: the loss of a toe; 5: the loss of a finger; 6: the loss of three fingers; 7: the loss of one hand; 8: the loss of one arm; 9: anything less serious than death. Following the main task, participants completed a brief demographic questionnaire. The materials are stated in full in the appendix on the project's OSF page.

7.3 Results

The proportion of participants who, in Step 1, opted *against* aggregation (no number of fingers per quarter outweigh saving a life) was 65.9%, which significantly exceeds the proportion of participants opting *for* aggregation, as well as chance (binomial tests, *ps*<.001).

The histograms (Figure **4**) represent what participants considered the morally required self-sacrifice on their behalf to save one person in the second step of the study (*left panel:* anti-aggregators; *right panel*: aggregators). They show a similar bimodal distributions peaking at the sacrifice of a bruised ankle (scale point 3) and anything short of death (scale point 9). A rank correlation analysis revealed a very small and nonsignificant correlation between *uninvolved* (step 1) and *involved* (step 2) assessment r_S=.16, p=.0628, a Mann-Whitney-U test also just failed to make the significance threshold (W=1733.5, p=.0630).

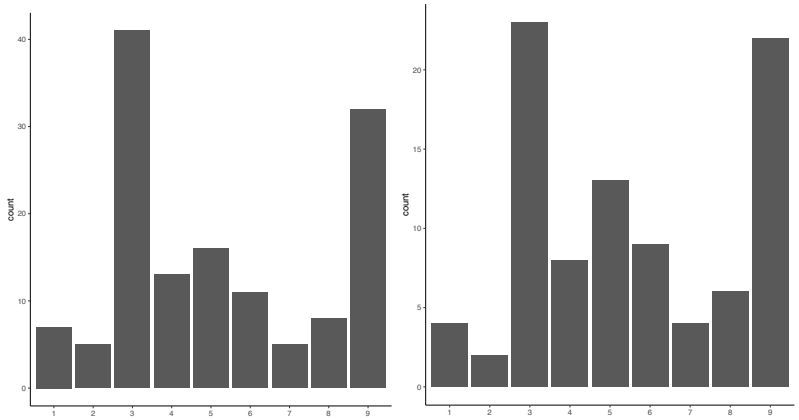

Figure 4. Histograms of morally required sacrifice ranging from 1 (nothing) via 5 (loss of a finger) to 9 (anything short of death) for the involved assessment (Step 2). *Left panel*: Participants who opted *against* aggregation in Step 1; *Right panel*: Participants who opted for aggregation in Step 1.

The proportion of participants who considered the largest morally required sacrifice to be *less* than the loss of a finger (choice of scale point below 5) among those who did not aggregate in Step 1 was 40.7% and significantly below the proportion of those who did aggregate in Step 2 (60.3%), $\chi^2(1)=4.69$, $p=.0304$. Neither proportion significantly differs from chance (binomial tests, *ps*>*.093*, two-tailed), see Figure 5 and Appendix, section 3.

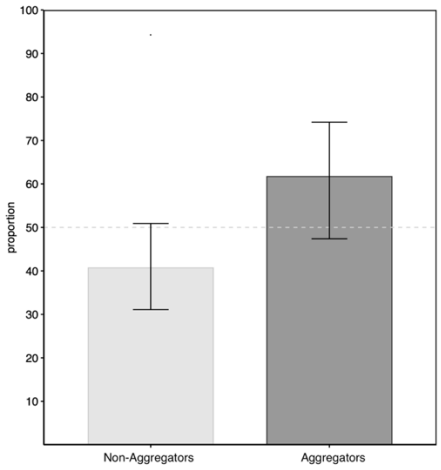

Figure 5. Proportions of participants indicating a less severe sacrifice than the loss of a finger (choosing options below the midpoint) split by uninvolved assessment (non-aggregators v. aggregators). Error bars denote 95% Agresti–Coull confidence interval.

7.4 Discussion

Experiment 3 documents a very broad range of responses when it comes to assessing the limits of permissible personal prerogatives, even amongst all those that think that a claim to losing a finger is *irrelevant* to the claim of losing one's life in standard partial aggregation cases (which we called 'uninvolved'): First, a fair number of respondents (24%) in this 'No Aggregation' group appear to take a *fully impartial* stance towards their own wellbeing, suggesting that they have a duty to forego anything less significant than what the other person would be suffering in a conflict case. Second, an even greater number of 'No Aggregation' respondents in our uninvolved scenario report that the greatest harm that they have a duty to suffer if the result is the saving of a life is *less* than the loss of a single finger (about 41%). Only 14% indicate that the loss of a finger constitutes the maximum harm they have a duty to suffer to save one other life. Beyond this, it strikes us as noteworthy that no significant difference can be discerned when comparing the assumed duty to forgo benefits for another amongst those who report non-aggregative intuitions for the uninvolved task and those who do.

8. General Discussion

We close by recapitulating our core findings and explaining their relevance. Overall, the study has extended existing experimental work on mixed conflicts of beneficence along several dimensions:

First, (through *Experiment 1*) we have added to the growing evidence that a significant percentage of subjects do not follow unrestricted aggregation views in distributive conflict that contain claims and values of very different magnitudes. Instead, they report judgments that are clearly aligned with partially aggregative views. Our study is more theoretically persuasive and representative than previous work on this subject.

Second, (through *Experiment 2*) we establish that a significant number of subjects report intuitions that closely align with the idea that in matters of life and death, minor improvements in outcome value constitute 'irrelevant goods', that is, values that are outranked by distributing fair chances. Whilst our finding is of independent interest to the wider issue of how prevalent non-consequentialist moral convictions are, it matters specifically for the "intra-partial-aggregation debate" amongst those favoring *localized* or *global*

relevance conditions. Whereas *global relevance* normatively predicts and endorses the moral intuition we document, *localized relevance* depends, at least prima facie, on rejecting it. Thus, our finding at least increases the argumentative burden on *localized* partial aggregation: Just like those defending full aggregation, advocates of localized versions must now insist that we should abandon at least one widely held non-consequentialist intuition in order to deliver a moral theory that is more coherent overall.

Third, we have shown (through *Experiment 3*) that there remains much work to be done insofar as defenders of partial aggregation want to move beyond merely appealing to 'bedrock' moral intuitions about particular cases (i.e. the non-reductionalist strategy (Tadros 2019, 176–79)). However, let us be clear about what our findings in Experiment 3 do not establish, namely a strong refutation of Voorhoeve's reductivist argument. The reductionist argument is a claim about how the *morally correct* view about personal prerogatives explains the *morally correct* view about the relevance of claims in aggregative conflict cases, so experimental data from actual intuitions amongst laypeople cannot disprove (or prove!) this argument.[13]

Nonetheless, we believe that our findings at least hint at the fact that the task of systematically connecting two domains of non-consequentialist common sense (anti-aggregative judgments and personal prerogatives) is complex: a substantive number of subjects report a pair of involved and uninvolved judgments that seems *prima facie* to contradict the idea that one aspect of common sense morality derives (or is systematically tied to) another. Our data does not clearly support the idea that intuitions in partial aggregation cases derive from individual permissions to refuse harms to oneself in defiance of the greater good. Specifically, a fair number of subjects that explicitly claim that fingers are *irrelevant* to the harm of death in the standard case also report that one has a permission to save one's own finger when giving it up could save a life. These laypeople follow intuitive judgments put forward by Kamm (2015, 693) and Tadros (2019, 177) that are meant to dispel the relation between personal prerogatives and partial aggregation postulated by Voorhoeve.

---

[13] Arguments critical of Voorhoeve's strategy can be found in (Kamm 2015; Tadros 2019; Brown 2022).

9. Conclusion

In closing we want to highlight some areas for future research. Overall, we see a clear need for further systematic empirical investigations in this field and want to suggest some potential avenues for how to proceed.

First, our Experiment 2 is amongst the first to investigate trade-offs that do not occur exclusively between claims that individuals have to certain *outcomes*, but also between *claims* to outcomes and *claims* to (fair) chances to have one's claims satisfied. Though our findings indicate that people care enough about giving fair chances to those with claims that are of equal weight to disregard some minor improvements in outcomes, it would be important (both theoretically and for practical healthcare contexts) to better understand *how much* weight moral common-sense places on giving fair chances and, conversely, where overall betterness outweighs such concerns.

Second, Experient 3 breaks new ground in seeking to uncover connections between different elements of non-consequentialist moral common sense. Our results show quite generally that, amongst laypeople, we do not find some neat alignment of beliefs that are often theoretically grouped together as constituting a non-consequentialist core. For example, (im)partiality with regards to one's own life vs. that of others does not seem to correlate in any straightforward or systematic way with an (un)willingness to fully aggregate in cases like *Life vs. Fingers*. Future research should attempt to uncover more systematically the relationship between personal prerogatives and anti-aggregative judgments.

Third, and more generally, we have employed a limited number of scenarios and experimental implementations. Future work on this topic should attempt to replicate the reported findings with a broader variety of cases, both within the domain of health and bodily harm but also beyond, and strive for different experimental designs so as to increase the external validity of the results.

Finally, we would like to stress that the data here reported was collected for US Americans, and one must refrain from generalizing to other societies,

in particular non-WEIRD ones.[14] As in other areas of experimental philosophy psychology, and cognitive science, replications with subjects from different cultures are key to drawing more conclusions about general human moral and cognitive dispositions. [15]

[14] WEIRD countries are *W*estern, *E*ducated, *I*ndustrialized, *R*ich and *D*emocratic, see (Henrich, Heine, and Norenzayan 2010). For a recent overview of findings regarding the impact of culture on cognition, see (Block and Kelly 2022)

[15] We would like to thank Thomas Pölzler, Veronika Luptakova, Alex Voorhoeve and the participants of Peter Schaber's Colloquium at the University of Zurich for excellent feedback. We do not want to suggest that any of them agree with our views. This work was supported by an SNF Ambizione Grant (PZ00P1_179912, PI: Kneer).

Bibliography

Allard, Aurélien, and Florian Cova. forthcoming. "What Experiments Can Teach Us about Justice and Impartiality: Vindicating Experimental Political Philosophy." In *Issues in Experimental Moral Philosophy*. London: Routledge.

Baderin, Alice. 2016. "Reflective Equilibrium: Individual or Public?" *Social Theory and Practice* 43 (1): 1–28. https://doi.org/10.5840/soctheorpract20174311.

Block, Andreas De, and Daniel Kelly. 2022. "Culture and Cognitive Science." *Stanford Encyclopedia of Philosophy*.

Broome, John. 2004. *Weighing Lives*. Oxford: Oxford University Press.

Brown, Campbell. 2022. "Aggregation and Self-Sacrifice." *Ethics* 132 (3): 730–35. https://doi.org/10.1086/718073.

Byrd, Nick. 2021. "Reflective Reasoning & Philosophy." *Philosophy Compass* 16 (11): 1–12.

Cowell, Frank A., Marc Fleurbaey, and Bertil Tungodden. 2015. "The Tyranny Puzzle in Social Preferences: An Empirical Investigation." *Social Choice and Welfare* 45 (4): 765–92. https://doi.org/10.1007/s00355-015-0880-9.

Damschroder, Laura J., Todd R. Roberts, Brian J. Zikmund-Fisher, and Peter A. Ubel. 2007. "Why People Refuse to Make Tradeoffs in Person Tradeoff Elicitations: A Matter of Perspective?" *Medical Decision Making* 27 (3): 266–80. https://doi.org/10.1177/0272989X07300601.

Epstein, Seymour, Rosemary Pacini, Veronika Denes-Raj, and Harriet Heier. 1996. "Individual Differences in Intuitive–Experiential and Analytical–Rational Thinking Styles." *Journal of Personality and Social Psychology* 71: 390–405. https://doi.org/10.1037/0022-3514.71.2.390.

Fried, Barbara H. 2020. *Facing Up to Scarcity: The Logic and Limits of Nonconsequentialist Thought*. Oxford University Press.

Gerken, Mikkel. 2017. *On Folk Epistemology: How We Think and Talk about Knowledge*. Oxford University Press.

Gils, Art van, and Patrick Tomlin. 2020. "Relevance Rides Again." In *Oxford Studies in Political Philosophy*, edited by David Sobel, Peter Vallentyne, and Steven Wall. Vol. 6. Oxford University Press.

Graham, Jesse, Jonathan Haidt, and Brian A. Nosek. 2009. "Liberals and Conservatives Rely on Different Sets of Moral Foundations." *Journal of Personality and Social Psychology* 96: 1029–46. https://doi.org/10.1037/a0015141.

Graham, Jesse, Brian A. Nosek, Jonathan Haidt, Ravi Iyer, Spassena Koleva, and Peter H. Ditto. 2011. "Mapping the Moral Domain." *Journal of Personality and Social Psychology* 101: 366–85. https://doi.org/10.1037/a0021847.

Halstead, John. 2016. "The Numbers Always Count." *Ethics* 126 (3): 789–802.

Hannikainen, Ivar R., Kevin P. Tobia, Guilherme da F. C. F. de Almeida, Noel Struchiner, Markus Kneer, Piotr Bystranowski, Vilius Dranseika, et al. 2022. "Coordination and Expertise Foster Legal Textualism." *Proceedings of the National Academy of Sciences* 119 (44): e2206531119. https://doi.org/10.1073/pnas.2206531119.

Hart, James. 2022. "Tie-Breaks and Two Types of Relevance." *Ethical Theory and Moral Practice* 25 (2): 315–34. https://doi.org/10.1007/s10677-022-10270-3.

Henrich, Joseph, Steven J. Heine, and Ara Norenzayan. 2010. "The Weirdest People in the World?" *Behavioral and Brain Sciences* 33 (2–3): 61–83. https://doi.org/10.1017/S0140525X0999152X.
Hirose, Iwao. 2001. "Saving the Greater Number without Combining Claims." *Analysis* 61 (4): 341–42.
———. 2014. *Moral Aggregation*. New York: Oxford University Press Inc.
Horton, Joe. 2018. "Always Aggregate." *Philosophy & Public Affairs* 46 (2): 160–74. https://doi.org/10.1111/papa.12116.
———. 2020. "Aggregation, Risk, and Reductio." *Ethics* 130 (4): 514–29. https://doi.org/10.1086/708534.
———. 2021. "Partial Aggregation in Ethics." *Philosophy Compass* 16 (3): e12719. https://doi.org/10.1111/phc3.12719.
Horvath, Joachim. 2010. "How (Not) to React to Experimental Philosophy." *Philosophical Psychology* 23 (4): 447–80. https://doi.org/10.1080/09515089.2010.505878.
Horvath, Joachim, and Alex Wiegmann. 2022. "Intuitive Expertise in Moral Judgments." *Australasian Journal of Philosophy* 100 (2): 342–59. https://doi.org/10.1080/00048402.2021.1890162.
Hsieh, NH, A Strudler, and D Wasserman. 2006. "The Numbers Problem." *Philosophy and Public Affairs* 34 (4): 352–72.
Kamm, F. M. 1993. *Morality, Mortality: Volume 1: Death & Whom to Save from It*. Revised. Oxford: Oxford University Press.
———. 2008. *Intricate Ethics: Rights, Responsibilities, and Permissible Harm*. New York: Oxford University Pres.
———. 2015. "Bioethical Prescriptions." *Journal of Medical Ethics* 41 (6): 493–95.
Kauppinen, Antti. 2007. "The Rise and Fall of Experimental Philosophy." *Philosophical Explorations* 10 (2): 95–118. https://doi.org/10.1080/13869790701305871.
Kelleher, J. Paul. 2014. "Relevance and Non-Consequentialist Aggregation." *Utilitas* 26 (04): 385–408.
Kneer, M., E. Alexander, L. Colaco, and E. Machery. 2021. "On Second Thought: Reflections on the Reflection Defense." In *Oxford Studies in Experimental Philosophy*. Vol. 4. Oxford: Oxford University Press.
Kneer, M., and S. Bourgeois-Gironde. 2017. "Mens Rea Ascription, Expertise and Outcome Effects: Professional Judges Surveyed." *Cognition* 169 (December): 139–46. https://doi.org/10.1016/j.cognition.2017.08.008.
———. 2018. "Intention, Cause, Et Responsabilité: Mens Rea Et Effet Knobe." In *Causalité, Responsabilité et Contribution à La Dette*, edited by S. Ferey and F. G'Sell, 117–44.
Lazar, Seth. 2018. "Limited Aggregation and Risk." *Philosophy & Public Affairs* 46 (2): 117–59. https://doi.org/10.1111/papa.12115.
Lefkowitz, David. 2008. "On the Concept of a Morally Relevant Harm." *Utilitas* 20 (04): 409–23.
Liao, S. Matthew. 2008. "A Defense of Intuitions." *Philosophical Studies* 140 (2): 247–62. https://doi.org/10.1007/s11098-007-9140-x.
Lindauer, Matthew. 2020. "Experimental Philosophy and the Fruitfulness of Normative Concepts." *Philosophical Studies* 177 (8): 2129–52. https://doi.org/10.1007/s11098-019-01302-3.
Loomes, Graham, and Lynda McKenzie. 1989. "The Use of QALYs in Health Care Decision Making." *Social Science & Medicine* 28 (4): 299–308. https://doi.org/10.1016/0277-9536(89)90030-0.

Luptakova, Veronika, and Alex Voorhoeve. forthcoming. "How Do People Balance Death against Lesser Burdens?" In *Advances in Experimental Political Philosophy*, edited by Michael Landauer. London: Bloomsbury.

Machery, Edouard. 2017. *Philosophy Within Its Proper Bounds*. Oxford University Press.

Mann, Kirsten. 2021. "The Relevance View: Defended and Extended." *Utilitas* 33 (1): 101–10. https://doi.org/10.1017/S095382082000028X.

———. 2022. "Relevance and Nonbinary Choices." *Ethics* 132 (2): 382–413. https://doi.org/10.1086/716873.

Miller, David. 2020. "Justice beyond the Armchair." *Contemporary Political Theory* 17 (1): 92–104.

Norcross, Alastair. 1997. "Comparing Harms: Headaches and Human Lives." *Philosophy & Public Affairs* 26 (2): 135–67.

Otsuka, Michael. 2004. "Skepticism about Saving the Greater Number." *Philosophy and Public Affairs* 32 (4): 413–26.

Pacini, Rosemary, and Seymour Epstein. 1999. "The Relation of Rational and Experiential Information Processing Styles to Personality, Basic Beliefs, and the Ratio-Bias Phenomenon." *Journal of Personality and Social Psychology* 76: 972–87. https://doi.org/10.1037/0022-3514.76.6.972.

Parfit, Derek. 1978. "Innumerate Ethics." *Philosophy & Public Affairs* 7 (4): 285–301.

———. 2013. *On What Matters Volume 2*. Oxford University Press.

Pinillos, N. Ángel, Nick Smith, G. Shyam Nair, Peter Marchetto, and Cecilea Mun. 2011. "Philosophy's New Challenge: Experiments and Intentional Action." *Mind & Language* 26 (1): 115–39. https://doi.org/10.1111/j.1468-0017.2010.01412.x.

Pinto-Prades, Jose-Luis, and Angel Lopez-Nicolás. 1998. "More Evidence of the Plateau Effect: A Social Perspective." *Medical Decision Making* 18 (3): 287–94. https://doi.org/10.1177/0272989X9801800306.

Pölzer, Thomas. 2023. "How to Do Empirical Political Philosophy: A Case Study of Miller's Argument for Needs-Based Justice." *Ms*.

Rüger, Korbinian. 2015. *The Close Enough View in Distributive Ethics and Two Dominance Objections.* (MSc dissertation). London: London School of Economics and Political Science.

———. 2020. "Aggregation with Constraints." *Utilitas* 32 (4): 454–71. https://doi.org/10.1017/S095382082000014X.

Scanlon, T.M. 1998. *What We Owe To Each Other*. Cambridge, MA: Belknap Press.

Schwitzgebel, Eric, and Fiery Cushman. 2012. "Expertise in Moral Reasoning? Order Effects on Moral Judgment in Professional Philosophers and Non-Philosophers." *Mind & Language* 27 (2): 135–53. https://doi.org/10.1111/j.1468-0017.2012.01438.x.

———. 2015. "Philosophers' Biased Judgments Persist despite Training, Expertise and Reflection." *Cognition* 141 (August): 127–37. https://doi.org/10.1016/j.cognition.2015.04.015.

Steuwer, Bastian. 2021. "Aggregation, Balancing, and Respect for the Claims of Individuals." *Utilitas* 33 (1): 17–34. https://doi.org/10.1017/S0953820820000217.

Tadros, Victor. 2019. "Localized Restricted Aggregation." In *Oxford Studies in Political Philosophy*, edited by David Sobel, Peter Vallentyne, and Steven Wall. Vol. 5. Oxford University Press.

Taurek, John M. 1977. "Should the Numbers Count?" *Philosophy & Public Affairs* 6 (4): 293–316.
Tomlin, Patrick. 2017. "On Limited Aggregation." *Philosophy & Public Affairs* 45 (3): 232–60. https://doi.org/10.1111/papa.12097.
Ubel, Peter A., George Loewenstein, Dennis Scanlon, and Mark Kamlet. 1996. "Individual Utilities Are Inconsistent with Rationing Choices: A Partial Explanation of Why Oregon's Cost-Effectiveness List Failed." *Medical Decision Making* 16 (2): 108–16. https://doi.org/10.1177/0272989X9601600202.
Voorhoeve, Alex. 2014. "How Should We Aggregate Competing Claims?" *Ethics* 125 (1): 64–87.
———. 2015. "Why Sore Throats Don't Aggregate against a Life, but Arms Do." *Journal of Medical Ethics* 41 (6): 492–93.
———. 2016. "Why One Should Count Only Claims with Which One Can Sympathize." *Public Health Ethics*, March, phw006. https://doi.org/10.1093/phe/phw006.
———. 2017. "Why One Should Count Only Claims with Which One Can Sympathize." *Public Health Ethics* 10 (2): 148–56. https://doi.org/10.1093/phe/phw006.
———. 2018a. "Balancing Small against Large Burdens." *Behavioural Public Policy* 2 (1): 125–42. https://doi.org/10.1017/bpp.2017.4.
———. 2018b. "Healthy Nails versus Long Lives: An Analysis of a Dutch Priority-Setting Proposal." In *Measuring the Global Burden of Disease: Philosophical Dimensions*, edited by D. Schroeder, N. Eyal, and D. Winkler. Oxford: Oxford University Press.
Waldron, Jeremy. 1993. *Liberal Rights: Collected Papers 1981-1991*. Cambridge: Cambridge University Press.
Weinstein, Milton C., George Torrance, and Alistair McGuire. 2009. "QALYs: The Basics." *Value in Health* 12: S5–9. https://doi.org/10.1111/j.1524-4733.2009.00515.x.
Wiegmann, Alex, Joachim Horvath, and Karina Meyer. 2020. "Intuitive Expertise and Irrelevant Options." *Oxford Studies in Experimental Philosophy* 3 (3): 275.