

# UC Merced

## Proceedings of the Annual Meeting of the Cognitive Science Society

### Title

Are there Irrelevant Utilities? What the Folk Think (and Why This is Relevant)

### Permalink

<https://escholarship.org/uc/item/66m3w8x7>

### Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

### Authors

Kneer, Markus

Viehoff, Juri

### Publication Date

2023

### Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

# Are there *Irrelevant Utilities*? What the Folk Think (And Why it Matters)

Markus Kneer ([markus.kneer@gmail.com](mailto:markus.kneer@gmail.com))

University of Zurich, Department of Philosophy, Zollikerstrasse 117  
8008 Zurich, CH

Juri Viehoff ([juri.viehoff@manchester.ac.uk](mailto:juri.viehoff@manchester.ac.uk))

University of Manchester, MANCEPT, Oxford Road  
M13 9PL Manchester, United Kingdom

## Abstract

We test the importance people attribute to the realization of small gains in outcome value for cases where the decision-maker must competitively distribute significant harm between separate groups. We find that, in line with recent non-consequentialist moral theories, subjects (i) sometimes rank giving those that stand to suffer harm equal chances *above* maximizing outcome value and (ii) that whether they opt for equal chance procedures ('coin flips') depends on the magnitude of the value that can be secured by not offering them. Our findings vindicate the idea that there can be 'irrelevant utilities' in cases of competing claims to avoid harm. Our study thus extends existing work on decision-making in conflict of harm cases along several dimensions, and we demonstrate their import for determining which version of 'partially aggregative' accounts in normative ethics aligns best with common sense.

**Keywords:** irrelevant utilities; aggregation; distributive ethics; harm; moral decision-making

## 1. Introduction

Doctors and healthcare professionals frequently have to choose between competing courses of action that lead to different individuals or groups suffering harms of differing magnitude: confronted with scarce medical resources and limited time, doctors must sometimes competitively choose whether to operate on one person with a serious condition or several with less serious ones. On a macro-level, healthcare administrators must decide whether to provide resources to treat common, but less harmful conditions at the expense of more serious, but rare conditions. Relevantly similar trade-offs also arise in non-medical contexts: For example, military commanders and anti-terrorist police may face scenarios where they can only avert an imminent, catastrophic threat by choosing one amongst several options that will all foreseeably cause harm of varying magnitudes to different-sized groups of innocent civilians.

Given the high stakes in many such decisions, understanding our process of reasoning, and improving accounts of how such decisions should be reached is of immense practical importance. Moreover, questions of what to do when moral claims of differing weights conflict with each other is of crucial importance in normative ethics, where competing theories make very different prescriptions as to how such conflicts should be resolved. And yet, experimental

work that systematically maps the prevalence of central moral intuitions to which theorists appeal when arguing about how decision-making in conflicts ought to be resolved is still in its infancy (Dolan et al. 2003, Shah 2009).

Our study contributes to this body of work: we investigate under what conditions subjects are prepared to forego beneficial outcomes in order to offer different groups that stand to be harmed by a decision equal chances of avoiding the harm. We demonstrate that a significant number of subjects judge *as if* some benefits were 'irrelevant utilities' (Kamm 1993). We also show how the decision to offer equal chances rather than maximize outcome value is conditioned by the magnitude of the value that offering equal chances would sacrifice, even if the additional value is primarily of an aesthetic kind. Our contribution thus extends existing work on anti-aggregative moral reasoning: First, we confirm, that 'irrelevant utility'-thinking is widespread, something that only recently has been tested for the first time (Kneer & Viehoff 2023). Second, we show that 'irrelevant utility'-thinking also has purchase for aesthetic value and in non-medical contexts. In our discussion, we assess the importance of our robust finding of 'irrelevant utility' thinking for the question *which of two variants* of partially aggregative moral theories best coheres with common sense.

## 2. Background: When to Aggregate?

The historically most prominent approach to resolving conflicts of the kind described above in bioethics and policy-making has been cost-benefit analysis (CBA) measured in terms of QALYs/DALYs or some other metric of outcome value (Loomes and McKenzie 1989; Hoedemaekers and Dekkers 2003; Weinstein, Torrance, and McGuire 2009). To illustrate, consider *Death vs Paraplegia*: We can save either (i) one person from death or (ii) one thousand people from paraplegia. The CBA method will allocate a QALY/DALY based value for the avoidance of death and the avoidance of paraplegia and will then calculate the 'net value' of (i) and (ii) by summing the QALYs/DALYs between patients on each side. Since the avoidance of paraplegia, though less important than avoiding death, is still very weighty, we likely save the 1000 from suffering it, thus leaving the one to die.

Though CBA here coheres with the intuitive judgment of a majority of subjects (Damschroeder et.al. 2007, Kneer & Viehoff 2023), it has faced serious theoretical objections

from non-consequentialist philosophers: being *unrestrictedly aggregative*, CBA also mandates that, as a matter of choice-worthiness, decision-makers rank very small improvements in health outcomes to groups made up of very many individuals *above* very significant improvements in health outcomes to a smaller number of people. To illustrate, consider *Death vs Migraines*, a scenario where we can only (i) save one person from death or (ii) save a huge number N of people from suffering a mild, short migraine (Norcross 1997, 1998; Dorsey 2009; Schönherr 2018; Horton 2021). Unrestricted aggregative approaches like CBA suggest that for some (potentially very large) N, we should prevent that number of migraines rather than save a person's life. But this, many have insisted, runs counter to moral common sense and fails to appropriately take into account the importance of the claim by the single person whom we could save (cf. Kamm 1993, 2008; Scanlon 1998; Otsuka 2006; Frick 2015).

### 'Partial aggregation' as an alternative to CBA

Faced with the judgments that we should aggregate claims in cases like *Death vs Paraplegia* but not do so in cases like *Death vs Migraines*, philosophers have recently begun to formulate 'partially aggregative' views that accommodate both these data points and offer some coherent moral explanation when aggregation is and isn't permissible (and why). Partial aggregation is a class of approaches according to which decisions amongst claims of a similar or 'relevant' weight or magnitude should be judged in accordance with the aggregative model of CBA (e.g. *Death vs Paraplegia*), whilst cases where the harms that we can avoid on one side of a binary decision are much less weighty for each affected person (e.g. *Death vs Migraines*) are judged in accordance with a non-aggregative model such that there is *no number* of short migraines the avoidance of which can justify not saving a person's life.

Though earlier statements of partially aggregative views can be found in the works of both Scanlon (1998) and Kamm (1993, 2008), the current debate is indebted to a recent attempt by Voorhoeve to systematically lay out such an approach, which we call 'Aggregate Relevant Claims' (ARC). It goes as follows (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.

Using the notion of moral claims with the two dimensions of (i) strength and (ii) relevance can help to capture the relative importance that each person's wellbeing has on the decision-maker. Specifically, (4.) and (5.) in combination can yield the pair of intuitions that underpinned the judgments in *Death vs Paraplegia* and *Death vs Migraine*: it is *because* paraplegia is, but migraine isn't, relevant to death that we should sum ('aggregate') instances of paraplegia but not migraines when the alternative is that of saving a life. The difference between a claim merely being *less strong* than another and that claim being *irrelevant* to the other is perhaps best spelled out in terms of the distinction between outweighing and disabling, e.g. in relation to reasons (Dancy 2006; Tadros 2019): A less weighty harm is merely outweighed because of the severity of what satisfying this claim would cost. An *irrelevant* claim, by contrast, is disabled in that it simply cannot have any duty-grounding force in the deliberation when its realization would prevent satisfying a much, much more important claim.

**Variants of Partial Aggregation.** Voorhoeve's pioneering study of ARC left open a number of issues that subsequent accounts have sought to address, most notably the question how partially aggregative views apply to scenarios (i) where the two competing option sets contain both claims that are relevant and those that are not relevant to the strongest competing claims and (ii) situations where the option set contains more than two options and we therefore (seem to) run into potential cycling/intransitivity issues (for respond to these, see: Voorhoeve 2018; Mann 2021; Lazar 2018; Tadros 2019; Harney & Khawaja 2023). Amongst those favoring partially aggregative views, at least two ways of understanding the crucial criterion of 'relevance' have crystalized, leading to divergent strategies of addressing potential difficulties for more complex scenarios. According to *anchor by strength* (AS), a claim, in order to be relevant, must be sufficiently strong relative to the *strongest overall claim* in the conflict. (AS) appeals to the idea that whatever we end up doing, it would be deeply problematic for a very weak claim or value gain to play a role in our deliberation about whether to help or refrain from helping somebody who has a very strong claim not to suffer serious harm. Alternatively, we may conceive of the relevance criterion in terms of *anchor by competition* (AC): a claim, in order to be relevant, must be sufficiently strong relative to the strongest claim *with which it competes* (Tomlin 2017; Mann 2021).

AS and AC lead to differences when we try to extend partially aggregative views to more complex scenarios. One prominent version is the one recently put forward by Tadros (2019): after introducing a further distinction (whether irrelevance is *local* or *global*) he opts for local AC. (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. Put differently: where two equally weighty claims can be found on each side of a decision, even very minor claims—say, the alleviation of a headache, could potentially determine which of two options we should

pursue) The resulting overall view, further discussed in 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 ‘outstanding’ claims. For visualization, consider a schematic proposed by Gils and Tomlin (2020: 22):

LEVEL (R=N+2)	GROUP A	GROUP B
1	1	
3		5
5	3	

**Table 1:** Sequential Claim Matching and Relevance Levels (Gils & Tomlin, 2020: 22)

In this example, if we assume that relevance extends ‘down’ for two levels, but no more, then, according to Tadros’ view, we start with ‘matching/cancelling’ the single claim in group A at level 1 with the 5 claims at level 3; suppose 5 claims at level 3 *just* outweigh the one claim at level 1, say by 4 to 1. Now, in the next step, the single remaining level 3 claim is weighed against the 3 level 5 claims. Since these outweigh this single claim, we ought to save group A. By contrast, if we follow AS as the proper conception of the relevance criterion, then the 3 claims by those at level 5 would be excluded *ab initio* and, consequently, we would opt to save group B instead.

What becomes clear if we compare this form of local, anchor-by-competition (AC) relevance to the AS alternative is that, importantly, it *runs counter* to one moral intuition that one might intuitively have thought to go hand in hand with the anti-aggregative judgment in *Death vs. Migraines*. This is the judgment that the alleviation of some harms is too insignificant to play a role in our decision whom to save.

The relevant intuition comes out clearest in a case designed by France Kamm, *Sore Throat*: a decision-maker has only enough medicine to save either person A or person B from a lethal illness, but if the decision-maker saves B, then they will have a little medicine left which can be used to cure person C’s sore throat. (If A is cured, C will suffer from the sore throat for a week.) Kamm thinks that we should not be influenced by such minor ‘sweetening’. Rather, the appropriate course of action, many would think, is a coin flip—in the context of our decision-situation, the additional benefit of avoiding C’s sore throat is an ‘irrelevant utility’.

According to Tadros’ *locally restricted aggregation* and other versions of AC, having matched the two lethal illnesses between A and B, we *should* allow C’s sore throat to determine what we should do (it is, after all, an unmatched claim). But this seems to violate the anti-aggregative intuition that minor improvements cannot prevent us from giving each of A and B an equal chance of survival by, for example, flipping a coin.

## The significance of ‘moral common sense’

Partial aggregation as a topic has attracted a huge deal of scholarly attention over recent years, both from advocates and critics. Importantly, we can distinguish two parts to the overall debate:

- (1) One sub-debate focuses on the question whether unrestricted aggregation or partial aggregation are more convincing overall approaches to distributive ethics (Halsted 2016; Horton 2017, 2018, 2020; Privitera 2017; Lazar 2018).
- (2) A more recent second debate concerns the issue which of a number of possible variants of the partial aggregation position (notably AS or AC) is most convincing (Voorhoeve 2014, 2018; Tomlin 2017; Tadros 2019; Gils & Tomlin 2020; Ruger 2020; Steuwer 2021; Mann 2021, 2022; Brown 2020, 2022; Hart 2022; Harney & Khawaja 2023).

Both in relation to (1) and (2), appeals to intuitive judgments to seemingly widely shared moral intuitions abound. We therefore think that understanding what *actual* moral common sense amounts to matters a great deal: First, 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 conform their intuitive judgments about particular cases (Rawls 1971; Daniels 1996), and all parties agree that intuitive judgments have considerable *prima facie* authority for theory construction.

But if it turned out that intuitive armchair judgments by philosophers differed dramatically from those of the wider population, or if empirical studies showed that judgments about core cases were frequently shaped by idiosyncratic factors (such as psychological dispositions or situational factors), then we should worry about the normative status of the principles that derive from reflective equilibrium (Tobia, Buckwalter, Stich 2013; Kneer et al. 2021).

In their debate with unrestricted aggregationists, 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 rationale for rejecting the former view. Defenders of *unrestricted aggregation* (like CBA) have in turn 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 (Broom 2004; Parfit 2013). They have done so because unrestricted aggregation here contradicts what most theorists take to be moral common sense. If we were to lose our conviction in this particular moral judgment, the case for partial aggregation would weaken significantly.

In relation to the second sub-debate, namely *which variant* of partial aggregation is most persuasive, it has recently become clear that at least AC and prominent accounts based on it, must, in order to fend off objections by those arguing for unrestricted aggregation, reject Kamm’s intuition from *Sore Throat* above. Now if AC-based accounts must, like

unrestricted aggregative views, reject a—supposedly widely held—anti-aggregative moral intuition, then it matters how prevalent and robust said intuition actually is amongst folk.

Luptakova & Voorhoeve (2023, 3) present a second reason why establishing the prevalence of public intuitions about conflict cases matter greatly beyond theory construction: whilst philosophical theories in normative ethics should exclusively aim for truth, those public institutions that will ultimately be required to implement the correct theories must *also* aim to be legitimate, that is, acceptable to those subject to their decisions. Where public policy diverges significantly from the moral beliefs of those to whom they apply, they not only become potentially impossible to be implemented, but may also violate a moral constraint on public justifiability or legitimacy. So if we want our institutions to conform to both standards of moral correctness and public legitimacy, uncovering moral common sense is of importance.

## 2. Existing Experimental Work

### Evidence of anti-aggregative intuitions

Philosophical suspicion of CBA is to some extent vindicated by recent experimental work. In various studies, a substantial number of subjects indicate a ‘mixed view’ that combines aggregative and non-aggregative judgments. On the one hand, subjects are prepared to offer equivalence ratios for harms that are close to each other in severity. (In other words, they offer a finite number or range in response to the question “*How many more patients with less serious impairment y would need to be cured for it to be the case that the decision-maker should choose this course of action rather than curing n number of patients with the more serious impairment x?*”). But, crucially, a significant number of subjects combine such aggregation-implying responses with what is called ‘off-scale refusals’ (Damschroeder et al. 2007): when asked how many cases of a much lesser impairment (e.g. ganglion cyst) would need to be treated to outweigh treating a more serious impairment (e.g. acute appendicitis), over 90% of subjects in one recent study reported that ‘no number’ would suffice (Kneer & Viehoff 2023). Other recent experimental work by Luptakova & Voorhoeve (2023) reaches similar, if less definitive, results. Older studies too indicate that significant numbers of subjects reject moral decision-making that is based on unrestricted aggregative approaches to conflict cases (Ubel et al. 1996; Damschroeder et al. 2007; Pinto-Prades & Lopez-Nicolás 1998). The documented mix of responses matches the pair of judgments in *Death vs. Paraplegia* (“Aggregate!”) and *Death vs. Migraines* (“Don’t Aggregate!”) that defines partial aggregation as a view.

### Limits of existing experimental work

Three features of the existing empirical work clearly indicate the need for more extensive, rigorous empirical analysis:

First, the more dated studies suffer from a number of shortcomings that limit their usefulness and external validity (for discussion, see: Voorhoeve 2018, Kneer & Viehoff 2023): Some of these studies rely on exceedingly small and

unrepresentative sample sizes combined with unfortunate framings in axiological terms (Ubel et al. 1996, for example has a sample size of  $N=42$ , all college students) whilst others do not permit readers to fully gauge the total or relative number of participants that displayed the combination of judgments characteristic of partially aggregative views (Pinto-Prades & Lopez-Nicolas 1998; see Voorhoeve 2018).

Second, existing experimental work elicits subjects’ non-aggregative intuitions only in relation to a particular set of cases and in one specific domain, namely healthcare priority setting in medical contexts where subjects must make trade-offs between competing groups with impairments that differ in severity. What is missing from the literature at present are studies that test a related, but different anti-aggregative intuitions about whether decision-makers ought to diverge from maximizing outcomes in order to distribute fair chances to those standing to suffer harms. Such intuitions play an important role in non-consequentialism and, as we showed above, have recently been shown to matter for the question which *variant* of partial aggregation is most grounded in moral common sense.

Third, existing studies have exclusively focused on medical contexts and trade-offs between individuals’ claims to the avoidance of harm. But in order to increase our confidence in the prevalence of anti-aggregationist intuitions, we should also test such intuitions in non-medical contexts and for value scales that do not exclusively rely on issues of harm and claims to welfare-based conditions.

## 3. Experiment

The principal goal of the experiment was to test for previously unexamined anti-aggregative intuitions in competitive harm trade-off cases that should display ‘irrelevant utility’-characteristics. We also designed the study with a view to extending the corpus of experimental work beyond narrow medical cases to military/law enforcement contexts where decisions about competing harms also loom large. Moreover, we aimed to ascertain whether or not values other than the avoidance of harm or bad health outcomes can matter when subjects are required to choose between maximizing outcomes or distributing fair chances amongst victims. The overall goal was thus to increase the validity of findings regarding anti-aggregationist common sense in several regards. Data and materials are available under <https://osf.io/5ny43/>.

### 3.1 Participants, Methods & Materials

**Participants.** 246 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, responded excessively quickly ( $t < 20s$ ), or whose native language was not English were excluded. 207 participants remained (mean age=37 years,  $SD=12$  years), of which 108 were female.

**Methods & Materials.** The vignette describes a terror drone heading for the market square where thousands are gathered

for a rock concert. Smith, a military commander, can shoot down the drone on Stanton Street, which will lead to the smallest number of civilian casualties. Shooting it down on Stanton Street No. 24 “an ordinary, recently built house” will lead to the death of five civilians. So will shooting it down over Stanton Street No. 27. No. 27 has a minor advantage over No.24 (it has a beautiful front garden) or a moderate advantage (it is a landmark building). Participants had to choose whether Smith should shoot it down over No.24 (ordinary house), No. 27 (beautiful garden v. landmark building) or else take the decision by aid of a *coin flip*.


We used a 2 *contrast* (minor v. moderate difference in utility) x 2 *perspective* (first-person v. third-person) design, to which participants were randomly assigned. In the first-person conditions, participants had to state what they *would* do in the shoes of Smith, in the third-person design they had to choose what Smith *should* do. Perspective is seriously understudied in experimental philosophy and moral psychology. Given that the limited evidence to date is mixed (see Feltz et al. 2012; Kneer 2018; Frisch et al. 2022; Horvath & Wiegmann 2022) it needs considerably more attention.

This *implicit* task was followed by an *explicit* task: participants were asked whether the individuals of the ordinary building had a lesser claim to be saved than inhabit the special building (i.e. the one with a nice front garden or the landmark building). As a manipulation check, we asked participants to rate the severity of harm of killing five to save many when the former live in (i) an ordinary house, (ii) a house with a nice front garden, or (iii) a landmark building on a scale from 0 (not bad at all) to 100 (extremely bad).

### 3.2 Results


**Implicit Task.** As hoped for, the difference in perceived severity between having the drone explode over an ordinary house ( $M=49.85$ ,  $SD=33.31$ ) did not differ significantly from having it explode over a house with a nice front garden ( $M=50.86$ ,  $SD=32.53$ ), paired samples t-test  $t(206)=-1.38$ ,  $p=.167$ , two-tailed. The perceived severity of shooting down the drone over the landmark building ( $M=58.00$ ,  $SD=31.84$ ) significantly exceeded the perceived severity of shooting it down over either of the other buildings ( $ps<.001$ ).

Since *perspective* (first person v. third-person) proved nonsignificant (for  $N=207$ ,  $\chi^2(1)=.239$ ,  $p=.887$ , Cramer’s  $V=.034$ ) we collapsed the results, see Figure 1. *Contrast* (minor v. moderate), i.e. whether to shoot down the drone over No. 24, No. 27 or flip a coin had a significant impact on action preference (for  $N=207$ ,  $\chi^2(2)=30.26$ ,  $p<.001$ , Cramer’s  $V=.382$ , a medium-sized effect).



**Figure 1:** Proportions of responses in the minor difference condition (shoot down drone over the house with beautiful garden v. ordinary house) and the moderate difference condition (landmark building v. ordinary house).

Since of particular importance for our purposes, we aggregated action preferences into *determinate* (shoot down the drone over No. 24 or No.27) and *randomized* ones (flip a coin), see Fig. 2. *Contrast* proved significant again (for  $N=207$ ,  $\chi^2(1)=29.16$ ,  $p<.001$ , Cramer’s  $V=.375$ , a medium-sized effect). In the moderate contrast condition (ordinary building v. landmark building) the proportion of people opting *against* a coinflip (82%) significantly exceeded chance (binomial test, two-tailed,  $p<.001$ ). In the minor contrast condition (ordinary building v. building with beautiful garden), those opting against a coinflip (46%) were in the minority (no significant difference from chance,  $p=.439$ ).



**Figure 2:** Proportions of responses across conditions.

**Explicit Task.** As regards the explicit task, in which participants were asked whether the people in the ordinary house had a lesser claim to be saved than those in the special house (beautiful garden v. landmark building), *contrast* (minor v. moderate) had a significant impact (for  $N=207$ ,  $\chi^2(1)=5.22$ ,  $p=.022$ ), see Fig. 3. Note, however, that the effect size was very small (Cramer’s  $V=.159$ ), and that in both conditions, the vast majority (>91%) of participants considered the claims of the inhabitants of both houses equally strong (significantly above chance,  $ps<.001$ ).




Figure 3: Proportions of responses across conditions in the abstract task

### 3.3 Discussion

Three of our experimental findings are especially noteworthy: First, the data suggests that a significant percentage of subjects does not judge in accordance with the unrestricted aggregative assumption the CBA-model prescribes: they reject the idea that, when weighty claims of individuals conflict, maximizing outcome-value is the only consideration that matters in our overall assessment of an option’s choiceworthiness. Specifically, they do not accept that, where lives are at stake and the additional ‘value’ that can be realized is the continued existence of a beautiful garden or an aesthetically pleasing building, such considerations should play a role in our assessment. Our findings here are in line with the extant literature that documents the existence of significant anti-aggregative sentiment amongst subjects for cases where the harms or values differ dramatically in magnitude (cf. Ubel et al. 1996; Damschroder et al. 2007; Pinto-Prades & Lopez-Nicolás 1998; Voorhoeve 2018; Luptakova & Voorhoeve 2023).

Second, our experiments indicate widespread existence of ‘irrelevant utility’ intuitions amongst laypeople. They coincide with studies that investigate the prevalence of such intuitions in other contexts. For example, recent experiments in relation to where healthcare resources should be allocated indicate that some minor harms appear to be ‘irrelevant utilities’ in such cases (Kneer & Viehoff, 2023).

Third, the data indicates that *how much chances matter* clearly depends on the significance of the cost that can be avoided by refusal. In other words, whilst subjects consider giving affected individuals fair chances of avoiding harm matters to some extent, they also care about producing good outcomes, and, when it would be very costly to give those with equally strong claims equal *ex ante* chances, subjects typically prefer to bring about better outcomes. As our findings show, even relatively minor increases in ‘better outcome’ (from ‘beautiful garden’ to ‘landmark building’) significantly reduces the proportion of subjects that are prepared to grant individuals with equally weighty claims to avoid harm equal chances of avoiding them.

## 4. General Discussion

We close by recapitulating our core findings, explaining their importance, and highlighting areas for future research. Our study has extended existing work along several

dimensions: First, 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. Second, 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 utilities’, that is, values that are outranked by distributing fair chances. Whilst our finding is of independent interest to the wider issue of non-consequentialist ethics, it should be of particular noteworthiness to those interested in the intra-partial-aggregation debate.

Finally, by both using a vignette from a military or law enforcement context and evaluating minor benefits that are of primarily aesthetic rather than welfare-related value, we have been able to establish that anti-aggregative and partially aggregative judgments, which had thus far only been documented in health contexts, also register outside of them.

In closing, we want to emphasize the need for further systematic research in this field and provide some potential avenues for how to proceed. Our study 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 experiment clearly indicates that people care about giving fair chances to those with claims that are of equal weight, there remains much to explore. Two examples must suffice here:

First, we know very little about the relevant thresholds at which even those willing to ‘sacrifice’ outcome value for the sake of fair chances deem these outcomes sufficiently weighty to silence/outweigh procedural considerations. One plausible hypothesis—yet to be tested—is that these thresholds are relative to the absolute strength of the claims that are tied on each side: presumably, protecting a beautiful garden will register more significantly when the choice is between preventing a number of minor injuries on each side than when we are dealing with questions of life and death. To test this, one would need to vary the severity of the tied claims whilst keeping the ‘irrelevant’ value or benefit constant.

Second, research should further explore the relevant characteristics of both those that report ‘irrelevant utility’ intuitions and those that ‘always aggregate’ for outcome value. Do anti-aggregative intuitions about *chances vs. outcomes* translate directly to anti-aggregative and anti-consequentialist judgments in outcome-only cases? Or are the important differences in how people respond to these? Moreover, do respondents who display consequentialist/anti-consequentialist leanings in these scenarios also display matching judgments in relation to other moral questions, e.g. the demands of impartiality between one’s own interests and those of others? Answering these questions would enhance our understanding of the landscape of aggregation-related moral intuitions. Our study has shown that ‘irrelevant utilities’ are relevant after all—now we should move on to when (and why!) they are relevant.

## Acknowledgments

Work on this project was supported by a Swiss National Science Foundation Ambizione grant (PZ00P1\_179912, PI: Kneer) and armasuisse Science and Technology (S+T), via the Swiss Center for Drones and Robotics

## References

- Broome, J. (2004). *Weighing Lives*. Oxford University Press.
- Brown, C. (2020). Is close enough good enough? *Economics & Philosophy*, 36(1), 29–59. <https://doi.org/10.1017/S0266267119000099>
- Brown, C. (2022). Aggregation and Self-Sacrifice. *Ethics*, 132(3), 730–735. <https://doi.org/10.1086/718073>
- Damschroder, L. J., Roberts, T. R., Zikmund-Fisher, B. J., & Ubel, P. A. (2007). Why People Refuse to Make Tradeoffs in Person Tradeoff Elicitations: A Matter of Perspective? *Medical Decision Making*, 27(3), 266–280. <https://doi.org/10.1177/0272989X07300601>
- Dancy, J. (2006). *Ethics without Principles*. Oxford University Press, U.S.A.
- Daniels, N. (1996). *Justice and Justification: Reflective Equilibrium in Theory and Practice*. Cambridge University Press.
- Dolan, P., Olsen, J. A., Menzel, P., & Richardson, J. (2003). An inquiry into the different perspectives that can be used when eliciting preferences in health. *Health Economics*, 12(7), 545–551. <https://doi.org/10.1002/hec.760>
- Dorsey, D. (2009). Headaches, Lives and Value. *Utilitas*, 21(01), 36–58. <https://doi.org/10.1017/S095382080800335X>
- Feltz, A., Harris, M., & Perez, A. (2012). Perspective in intentional action attribution. *Philosophical Psychology*, 25(5), 673–687. <https://doi.org/10.1080/09515089.2011.627540>
- Frick, J. (2015). Contractualism and Social Risk. *Philosophy & Public Affairs*, 43(3), 175–223. <https://doi.org/10.1111/papa.12058>
- Frisch, L. K., Kneer, M., Krueger, J. I., & Ullrich, J. (2021). The effect of outcome severity on moral judgement and interpersonal goals of perpetrators, victims, and bystanders. *European Journal of Social Psychology*, 51(7), 1158–1171. <https://doi.org/10.1002/ejsp.2805>
- Gils, A. van, & Tomlin, P. (2020). Relevance Rides Again. In D. Sobel, P. Vallentyne, & S. Wall (Eds.), *Oxford Studies in Political Philosophy* (Vol. 6). Oxford University Press.
- Halstead, J. (2016). The Numbers Always Count. *Ethics*, 126(3), 789–802.
- Harney, J., & Khawaja, J. (2023). How to Assess Claims in Multiple-Option Choice Sets. *Philosophy & Public Affairs*, 51(1), 60–92. <https://doi.org/10.1111/papa.12227>
- Hart, J. (2022). Tie-breaks and Two Types of Relevance. *Ethical Theory and Moral Practice*, 25(2), 315–334. <https://doi.org/10.1007/s10677-022-10270-3>
- Hoedemaekers, R., & Dekkers, W. (2003). Key Concepts in Health Care Priority Setting. *Health Care Analysis*, 11(4), 309–323.
- Horton, J. (2017). Aggregation, Complaints, and Risk. *Philosophy & Public Affairs*, 45(1), 54–81. <https://doi.org/10.1111/papa.12084>
- Horton, J. (2018). Always Aggregate. *Philosophy & Public Affairs*, 46(2), 160–174. <https://doi.org/10.1111/papa.12116>
- Horton, J. (2021). Partial aggregation in ethics. *Philosophy Compass*, 16(3), e12719. <https://doi.org/10.1111/phc3.12719>
- Horvath, J., & Wiegmann, A. (2022). Intuitive Expertise in Moral Judgments. *Australasian Journal of Philosophy*, 100(2), 342–359. <https://doi.org/10.1080/00048402.2021.1890162>
- Kamm, F. (2008). *Intricate Ethics: Rights, Responsibilities, and Permissible Harm*. Oxford University Press.
- Kamm, F. (1993). *Mortality, Volume 1: Death & Whom to Save from It* (Revised.). Oxford University Press.
- Kelleher, J. P. (2014). Relevance and Non-consequentialist Aggregation. *Utilitas*, 26(04), 385–408.
- Kneer, M. (2018). Perspective and Epistemic State Ascriptions. *Review of Philosophy and Psychology*, 9(2), 313–341. <https://doi.org/10.1007/s13164-017-0361-4>
- Kneer, M., Colaço, D., Alexander, J. & Machery, E. (2021). On Second Thought: Reflections on the Reflection Defense. In T. Lombrozo, J. Knobe, & S. Nichols (Eds.), *Oxford Studies in Experimental Philosophy Volume* (Vol. 4). Oxford University Press.
- Kneer, M., & Viehoff, J. (2023). Partial Aggregation: What the People Think. *Ms*.
- Lazar, S. (2018). Limited Aggregation and Risk. *Philosophy & Public Affairs*, 46(2), 117–159. <https://doi.org/10.1111/papa.12115>
- Lefkowitz, D. (2008). On the Concept of a Morally Relevant Harm. *Utilitas*, 20(04), 409–423.
- Loomes, G., & McKenzie, L. (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](https://doi.org/10.1016/0277-9536(89)90030-0)
- Luptakova, V., & Voorhoeve, A. (forthcoming). How Do People Balance Death against Lesser Burdens? In M. Landauer (Ed.), *Advances in Experimental Political Philosophy*. Bloomsbury.
- Mann, K. (2021). The Relevance View: Defended and Extended. *Utilitas*, 33(1), 101–110. <https://doi.org/10.1017/S095382082000028X>
- Mann, K. (2022). Relevance and Nonbinary Choices. *Ethics*, 132(2), 382–413. <https://doi.org/10.1086/716873>
- Norcross, A. (1997). Comparing Harms: Headaches and Human Lives. *Philosophy & Public Affairs*, 26(2), 135–167.
- Norcross, A. (1998). Great harms from small benefits grow: How death can be outweighed by headaches. *Analysis*, 58(2), 152–158.
- Otsuka, M. (2006). Saving Lives, Moral Theory, and the Claims of Individuals. *Philosophy & Public Affairs*, 34(2), 109–135. <https://doi.org/10.1111/j.1088-4963.2006.00058.x>



- Parfit, D. (2013). *On What Matters Volume 2*. Oxford University Press.
- Pinto-Prades, J.-L., & Lopez-Nicolás, A. (1998). More Evidence of the Plateau Effect: A Social Perspective. *Medical Decision Making*, 18(3), 287–294. <https://doi.org/10.1177/0272989X9801800306>
- Privitera, J. (2018). Aggregate Relevant Claims in Rescue Cases? *Utilitas*, 30(2), 228–236. <https://doi.org/10.1017/S0953820817000279>
- Pummer, T. (2013). Intuitions about large number cases. *Analysis*, 73(1), 37–46.
- Rawls, J. (1971). *A theory of justice*. Belknap Press of Harvard University Press.
- Rüger, K. (2020). Aggregation with Constraints. *Utilitas*, 32(4), 454–471. <https://doi.org/10.1017/S095382082000014X>
- Scanlon, T. M. (1998). *What We Owe To Each Other*. Belknap Press of Harvard University Press.
- Schönherr, J. (2018). Still Lives for Headaches: A reply to Dorsey and Voorhoeve. *Utilitas*, 30(2), 209–218. <https://doi.org/10.1017/S0953820817000152>
- Shah, K. K. (2009). Severity of illness and priority setting in healthcare: A review of the literature. *Health Policy*, 93(2), 77–84. <https://doi.org/10.1016/j.healthpol.2009.08.005>
- Steuwer, B. (2021). Aggregation, Balancing, and Respect for the Claims of Individuals. *Utilitas*, 33(1), 17–34. <https://doi.org/10.1017/S0953820820000217>
- Tadros, V. (2019a). Localized Restricted Aggregation. In D. Sobel, P. Vallentyne, & S. Wall (Eds.), *Oxford Studies in Political Philosophy: (Vol. 5)*. Oxford University Press.
- Tomlin, P. (2017). On Limited Aggregation. *Philosophy & Public Affairs*, 45(3), 232–260. <https://doi.org/10.1111/papa.12097>
- Ubel, P. A., Loewenstein, G., Scanlon, D., & Kamlet, M. (1996). Individual Utilities Are Inconsistent with Rationing Choices: A Partial Explanation of Why Oregon’s Cost-Effectiveness List Failed. *Medical Decision Making*, 16(2), 108–116. <https://doi.org/10.1177/0272989X9601600202>
- Voorhoeve, A. (2014). How Should We Aggregate Competing Claims? *Ethics*, 125(1), 64–87.
- Voorhoeve, A. (2015). Why sore throats don’t aggregate against a life, but arms do. *Journal of Medical Ethics*, 41(6), 492–493. <http://dx.doi.org/10.1136/medethics-2014-102036>
- Voorhoeve, A. (2016). Why One Should Count Only Claims with which One Can Sympathize. *Public Health Ethics*, 10(2), 148–156. <https://doi.org/10.1093/phe/phw006>
- Voorhoeve, A. (2018). Balancing small against large burdens. *Behavioural Public Policy*, 2(1), 125–142. <https://doi.org/10.1017/bpp.2017.4>
- Weinstein, M. C., Torrance, G., & McGuire, A. (2009). QALYs: The basics. *Value in Health*, 12, S5–S9. <https://doi.org/10.1111/j.1524-4733.2009.00515.x>