

Fair Chances vs. Better Outcomes in Complex Moral Trade-Offs

Keywords: moral judgment, aggregation; fairness; distributive ethics; moral decision-making; bio x-phi

Allocation decisions in healthcare and public policy often pit “fair chances” against “better outcomes”: When two parties have equally strong claims to avoid a serious harm, an allocator can either decide *determinately* in the way that secures a slightly better overall outcome, or use a lottery to give each claimant an equal *ex ante* chances. Standard cost–benefit/cost-effectiveness analysis and many consequentialist theories treat fair chances as having no independent value and therefore recommend always choosing the outcome-maximizing option, even when the improvement is small. By contrast, non-consequentialist proposals—especially “partially aggregative” views (e.g. Voorhoeve 2014, 2016; Tomlin 2017; Mann 2022) and Kamm’s notion of “irrelevant utilities” (1993, 2008)—theorize that some small outcome improvements may be treated as morally disabled in the presence of weighty competing claims, so that fair-chance procedures are morally required. Recent work has confirmed the widespread prevalence of such complex non-aggregative intuitions amongst laypeople (Luptakova & Voorhoeve 2023; Kneer & Viehoff 2023, 2025). However, what is missing for a more complete theory of actual moral decision-making in this field is systematic evidence about (a) how sensitive such judgments are to the magnitude of the foregone gain, (b) whether they persist once choices are embedded in institutional roles, (c) how they interact with baseline need/priority differences, (d) whether they extend to non-harm goods, and (e) whether they survive explicit reflection about the strength of claims. Extending existing findings along these dimensions is both practically relevant for real-world decision-making beyond a narrow set of cases in bioethics, and of theoretical importance for an improved understanding of the various factors that shape moral judgment in this domain.

We report a set of vignette experiments ordered by increasing complexity: a simple harm-based rescue case (“Snake”), an institutional healthcare allocation case (“Ambulance”) that varies whether additional benefit accrues to a separate person or to the primary beneficiary, and a non-harm/non-welfare case (“Drone”). In the presentation, we might also include a series of further experiments (total $N=2745$), which replicate our findings across contexts, formulations and experimental designs.

Study 1 (“Snake”; $N=309$) adapts a classic two-claimant rescue case. Participants chose between saving one endangered person, saving the other endangered person while also producing an additional benefit for a third party, or flipping a coin between the two primary claimants. We manipulated the magnitude of the third-party benefit (“mild headache” vs. “tolerabilis”) and perspective (first-person vs. third-person). Added utility had a strong effect on choice ($\chi^2(2)=34.93$, $p<.0001$), while perspective did not ($\chi^2(2)=3.15$, $p=.207$). When the added benefit was small (“mild headache”), the coin flip dominated (65.4%), consistent with treating the extra benefit as an “irrelevant utility”; when the added benefit was larger (“tolerabilis”), the outcome-improving determinate option became the majority choice (58.8%), while coin flips dropped to 36.6%, see Figure 1.

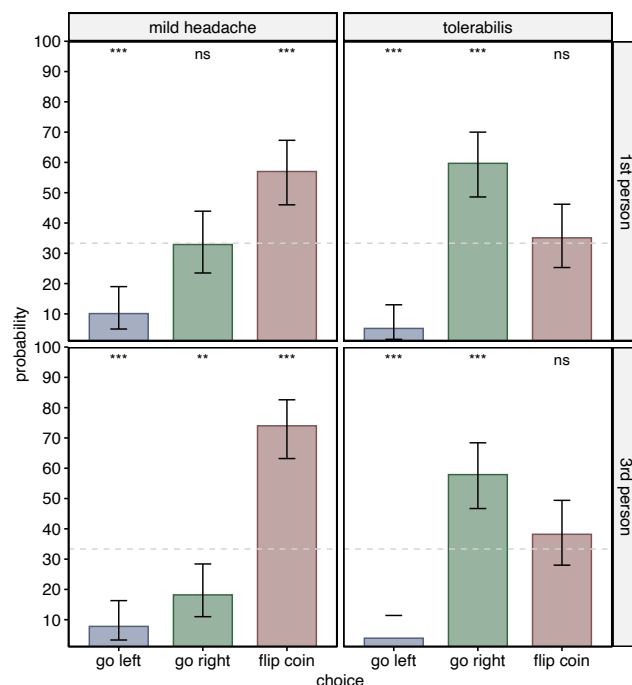


Figure 1 Probabilities of choices split by added utility to a third person and perspective. Error bars denote 95% CIs. Difference from chance: * $p<.05$; ** $p<.01$; *** $p<.001$; **** $p<.0001$.

Studies 2–3 (“Ambulance / Emergency Operator”; N=323) embed the same structural choice in a healthcare allocation role and test interactions with priority/need. Participants read about an emergency-response shift leader who must direct a single ambulance to one of two simultaneous accidents. In both accidents, the driver will become paraplegic unless treated in time. Accident B additionally involves either a marginal injury (a bruised ankle) or a substantial injury (loss of a finger). Crucially, the extra injury either affects a second person (multiple-beneficiary institutional allocation) or affects the driver themselves (intra-personal need variation). In a multinomial model, both contrast and distribution significantly affected choices (contrast $\chi(2)=26.31$, $p<.0001$; distribution $\chi(2)=33.41$, $p<.0001$), while the interaction was not significant ($\chi(2)=1.20$, $p=.5494$). In the multiple-beneficiary/marginal contrast condition, about one in two participants chose the coin flip (exceeding both the other options and chance set at 1/3). When the marginal benefit increased to a substantial benefit (loss of a finger), the majority (about 60%) chose the outcome-maximizing determinate option, though a sizeable minority (about one in three) still chose the coin flip. When the substantial additional harm threatened the same individual (single-person/substantial), nearly everyone prioritized sending the ambulance to that individual rather than flipping a coin—showing that sufficiently large need differences can override fair-chance reasoning.

Study 4 (“Drone”; N=207) tests whether “irrelevant utility” judgments extend beyond welfare and health to cases where the “better outcome” is the preservation of impersonal/non-health value. Participants chose where to shoot down an imminent terror drone, knowing that five civilians would die whichever option is chosen; the only difference was whether the civilians were in an ordinary building or a “special” one (minor contrast: beautiful garden; moderate contrast: landmark building). Contrast significantly altered choice patterns ($\chi^2(2)=30.26$, $p<.001$). Aggregating responses into randomized (coin flip) vs determinate (either building), determinate choices were a minority in the minor contrast condition (46% against coin flip) but dominant in the moderate contrast condition (82% against coin flip). Importantly, the explicit claim-strength question showed that in both contrasts the vast majority (>91%) judged the two groups’ claims to be equally strong, even when they nonetheless favored the determinate, outcome-preserving option—suggesting that many subjects do not rationalize determinate choice by downgrading one group’s claim.

Across contexts, the data support a complex picture of non-consequentialist moral judgment: fair *ex ante* chances have independent moral force; many people treat some gains—small welfare gains in rescue/healthcare cases and modest impersonal goods—as too small to defeat the moral force of equal *ex ante* treatment; once the foregone gain becomes larger, many participants switch to the determinately better outcome. Moreover, priority to need can dominate when need differences are substantial and intra-personal. These patterns fit neither an unrestricted maximization rule nor a categorical “always lottery when claims are equal” judgment. They instead point toward thresholded or partially aggregative accounts on which (i) chances matter in their own right, (ii) some benefits can be “disabled” in the presence of weighty competing claims, but (iii) this disabling is not absolute and can be overcome by weighty countervailing considerations.

Bibliography

- Kamm, F. M. 1993. *Morality, Mortality: Volume 1: Death & Whom to Save from It: Volume I: Death and Whom to Save from It*. Reprint edition. New York: Oxford University Press.
- Kamm, F. M. 2008. *Intricate Ethics: Rights, Responsibilities, and Permissible Harm*. New York: Oxford University Press.
- Kneer, Markus, and Juri Viehoff. 2023. “Are There Irrelevant Utilities? What the Folk Think (and Why This Is Relevant).” *Proceedings of the Annual Meeting of the Cognitive Science Society* 45(45).
- Kneer, Markus, and Juri Viehoff. 2025. “Partial Aggregation: What the People Think.” *Australasian Journal of Philosophy* Online First:1–22.
- Luptakova, Veronika, and Alex Voorhoeve. 2023. “How Do People Balance Death against Lesser Burdens?” in *Advances in Experimental Political Philosophy*, edited by M. Landauer. London: Bloomsbury.
- Mann, Kirsten. 2022. “Relevance and Nonbinary Choices.” *Ethics* 132(2):382–413. doi:[10.1086/716873](https://doi.org/10.1086/716873).
- Tomlin, Patrick. 2017. “On Limited Aggregation.” *Philosophy & Public Affairs* 45(3):232–60.
- Voorhoeve, Alex. 2014. “How Should We Aggregate Competing Claims?” *Ethics* 125(1):64–87.
- Voorhoeve, Alex. 2016. “Why One Should Count Only Claims with Which One Can Sympathize.” *Public Health Ethics* phw006.