Received: 14 September 2022 | Revised: 10 March 2023 | Accepted: 17 May 2023

DOI: 10.1111/bioe.13190

ORIGINAL ARTICLE

bioethics WILEY

ICU triage decisions and biases about time and identity

Joona Räsänen

Department of Political Science, School of Business and Social Sciences, Centre for the Experimental-Philosophical Study of Discrimination (CEPDISC), Aarhus University, Aarhus, Denmark

Correspondence
Joona Räsänen, Department of Political Science, School of Business and Social Sciences, Centre for the Experimental-Philosophical Study of Discrimination (CEPDISC), Aarhus University, Aarhus, Denmark.
Email: joona.rasanen@ps.au.dk

Funding information
Danish National Research Foundation, Grant/Award Number: DNRF: 144

Abstract

We often show a greater inclination to assist and avoid harming people identified as those at high risk of great harm than to assist and avoid harming people who will suffer similar harm but are not identified (as yet). Call this the identified person bias. Some ethicists think such bias is justified; others disagree and claim that the bias is discriminatory against statistical people. While the issue is present in public policy and politics, perhaps the most notable examples can be found in medical ethics such as in ICU triage decisions made during the COVID-19 pandemic. The Rule of Rescue, as the application of the identified person bias is sometimes called, states that spending large amounts of resources rescuing identifiable individuals who are in imminent danger is justified. In this paper, I show that our distorted attitudes toward time play a role in identified person bias. I claim that ICU triage decisions are better explained by a preference to treat people sooner rather than later, which is at least partially due to near bias (positive events are to be preferred to be near rather than distant), than by a preference to treat identified lives over statistical lives. Thus, another bias, near bias, is involved in the reasoning behind the identified person bias and the Rule of Rescue.

KEYWORDS
medical ethics, public policy, resource allocation, Rule of Rescue, time-bias

1 | INTRODUCTION

Imagine that you know all of the following and nothing more (of any relevance to your decision).

> **Ventilator$_1$**. There is a horrible pandemic raging and a hospital where incidences are high has just one ventilator. There is no way to get a ventilator from other hospitals and there is no possibility of patients being transferred to another hospital. At the moment, there is one patient, Ann, who needs the ventilator and will die for sure without it but with the help of the ventilator she will recover and return to normal within a few weeks. However, we know that during those few weeks, some of the people who are currently infected with the virus, but not yet severely ill, will need the ventilator. Without the ventilator, whoever falls ill within the next few weeks will die, but they will fully recover with the help of the ventilator. You know some information about Ann's personal life. You know nothing about the people who need the ventilator in a few weeks.[1]

[1]H. Orri Stefánsson used a similar example when presenting his work Stefánsson, H.O. (forthcoming). Identified person "bias" as decreasing marginal value of chances. *Noûs*. https://doi.org/10.1111/nous.12470. Stefánsson revised the example based on my criticism. I thank him for permission to use his earlier example here.

What should we do in this case? Who should get the ventilator? Should we connect the currently ill patient, Ann, to the ventilator or should we wait so that someone else can use it during the next few weeks?

Perhaps it seems obvious that we should give the ventilator to the person who is currently ill (in this case Ann) rather than let her die while we wait for the next patient. But this judgment is puzzling since it is certain that someone else will need the ventilator within the next few weeks, we just don't know who that person will be. So why do we think that we should give the ventilator to the patient who is currently ill; why should we prioritize this identified person's life over statistical lives? Call this the identified person bias.

A quote from Daniels illustrates the difference between identified and statistical victims.

> By an "identified victim," I mean Terry Q., lying injured in the passenger seat of a wrecked automobile on the corner of Main Street and Broadway, or Jessica McClure, the child who fell into the Texas a well in 1987 and whose family was sent $700.000 in donations for her. [...] By a "statistical victim," I mean the person who, extrapolating from traffic records, will be in a similar, serious car accident tomorrow (and may then be identified), or the children who will fall into wells next year if we do not cap them better than we did the well that trapped Jessica.[2]

Cohen, Daniels, and Eyal define identified person bias as follows:

> A greater inclination to assist (and avoid harming) persons and groups identified as those at high risk of great harm than to assist (and avoid harming) persons and groups who will suffer (or already suffer) similar harm but are not identified (as yet).[3]

So, according to the bias, people are more inclined to save Ann in Ventilator$_1$, than to save the statistical person. Likewise, people are more inclined to save Terry Q. than a statistical person who will be in a similar car accident tomorrow, and to save Jessica McClure than a statistical person who will fall into a well next year.

A version of the identified person bias is the Rule of Rescue. Very roughly, the Rule of Rescue states that it is permissible to spend large amounts of resources rescuing identified individuals who are in imminent peril.[4] The term Rule of Rescue was first used by Jonsen.[5]

The tension between saving identified individuals and the injunction to do as much good as possible can generate serious ethical difficulties.[6] As Cookson et al. note, policymakers are faced with difficult decisions, for instance, when funding controversial health technologies, such as costly new drugs.[7]

For example, in the 1990s, in Ontario, Canada, officials decided to fund medical care for twenty patients suffering from severely disabling, if not life-threatening complications (Gaucher disease) at a cost of 3.5 million dollars annually after a heated public debate.[8] National Gaucher Foundation of Canada appealed the Rule of Rescue, stating that the minister of health could not stand by and watch a very sick man suffer terribly when the means to relieve his suffering were at her fingertips.[9]

A more recent example can be found in the United Kingdom where Edward Willis-Hall, a baby from Colchester, with fatal spinal muscular atrophy, received the world's most expensive drug in the world, which costs £1.79 m to the NHS.[10] While the parents of Edward raised some of the money privately and the NHS managed to negotiate an undisclosed discount on its list price, it remains obvious that many statistical lives could have been saved with the money that now was used to save the life of Edward. Yet, according to the Rule of Rescue, it was morally permissible—perhaps even obligatory—to save Edward and let the statistical people die (even though an argument can be made that Edward, because he was just a young baby, would not have been harmed as much as older people are harmed by death[11]).

Many people think we cannot simply stand idle when an identified person's life—such as Edward's—is visibly threatened if rescue measures are available. Yet, some have claimed that saving identified lives is discriminatory against statistical lives because being identifiable is not a morally relevant feature.[12] This tension created by the opposite intuitions and opinions is the source of ethical and political difficulties.[13]

Some say that identified person bias is unjust or that it is not rational: we should *not* prioritize the identified victim or a patient over the unknown statistical victim or a patient.[14] Others say that we

[2]Daniels, N. (2012). Reasonable disagreement about identified vs. statistical victims. *Hastings Center Report, 42*, 35–45.

[3]Cohen, I. G., Daniels, N., & Eyal, N. (2015). Statistical versus identifed persons: An introduction. In I. G. Cohen, N. Daniels, & N. Eyal (Eds.), *Identified versus statistical lives: An interdisciplinary perspective*. Oxford University Press.

[4]For other, more detailed definitions of the Rule of Rescue see. Sheehan, M. (2007). Resources and the Rule of Rescue. *Journal of Applied Philosophy, 24*(4), 352–366.

[5]Jonsen, A. R. (1986). Bentham in a box: Technology assessment and health care allocation. *Law, Medicine and Health Care, 14*, 172–174

[6]Hadorn, D. C. (1991). Setting health care priorities in Oregon: Cost-effectiveness meets the Rule of Rescue. *JAMA, 265*(17), 2218–2225.

[7]Cookson, R., McCabe, C., & Tsuchiya A. (2008). Public healthcare resource allocation and the Rule of Rescue. *Journal of Medical Ethics, 34(7)*, 540–544.

[8]Gross, M. L. (2002). Ethics, policy, and rare genetic disorders: The case of Gaucher disease in Israel. *Theoretical Medicine and Bioethics, 23*(2), 151–170.

[9]Clarke, J. T. R., Amato, D., & Deber, R. B. (2001). Managing public payment for high-cost, high-benefit treatment: enzyme replacement therapy for Gaucher's disease in Ontario. *Canadian Medical Association Journal, 165*(5), 595–596.

[10]Jones, C. (2021, December 31). Essex baby given £1.79 m drug has new lease of life. *BBC News*. Retrieved July 1, 2022, from https://www.bbc.com/news/uk-england-essex-59565966

[11]McMahan, J. (2002). *The ethics of killing*. Oxford University Press; Räsänen J. (2020). Saving the babies or the elderly in a time of crisis? *The American Journal of Bioethics, 20*(7), 180–182.

[12]McKie, J., & Richardson, J. (2003). The Rule of Rescue. *Social Science & Medicine, 56*(12), 2407–2419. For critical discussion on discrimination and the Rule of Rescue see Lübbe, W. (2019). Appeal to the Rule of Rescue in health care: Discriminating and not benevolent? *Medicine, Health Care & Philosophy, 22*, 53–58.

[13]Cookson, R. et al., op. cit. note 7.

[14]Brock, D. (2015). Identified versus statistical lives: Some introductory issues and arguments. In I. G. Cohen, N. Daniels, & N. Eyal (Eds.), *Identified versus statistical lives: An interdisciplinary perspective*. Oxford University Press; Adler, M. (2015). Welfarism, equity, and the choice between statistical and identified victims. In I. G. Cohen, N. Daniels, & N. Eyal

should save the identified life[15]—sometimes even at the expense of saving a somewhat greater number of statistical lives.[16] I am not going to take a side on the issue.

Instead, in this paper, I argue that there is a natural explanation for why we often (but perhaps not always) think we should save the identified life over statistical life. The explanation has nothing to do with the distinction between identified and statistical lives, however. The explanation for the intuitive response to save the identified life rather than statistical lives comes from another bias: the near bias which is one instance of the time bias: our distorted attitude toward time.

2 | THE NEAR BIAS PLAYS A ROLE IN EXPLAINING THE IDENTIFIED PERSON'S BIAS

People are biased regarding time. If a genie tells you that you can have three wishes now, or in two months, you—like most people—will likely choose to have the wishes now rather than later. On the other hand, if a doctor tells you that you will need to go for a painful medical operation and that you can decide to have it performed in 2 weeks or in 2 years, you, as most people, will likely choose to go for the operation in 2 years (assuming other things are equal and postponing the operation does not have additional costs to you).[17]

People thus prefer good things happening in the present (or soon in the future) rather than far away in the future. People likewise prefer bad things happening far away in the future, rather than now or soon in the future. This is one kind of time bias. Call it the *near bias*.[18] There is an even stronger time bias where people prefer good things being in your future rather than the past and bad things being in your past rather than in your future. Call this the *future bias*.

There is a disagreement among philosophers and ethicists about whether it is rational to be future- or near-biased. Some are inclined to say future bias is irrational,[19] others think at least some kind of future biases are rational,[20] and some claim if you are near-biased you should be future-biased as well.[21] Yet some claim the arguments against the rationality of near bias are weak and bias toward the near does not necessarily violate the requirements of rationality.[22] Things get even more complicated when we turn our attention to other people and ask if we are—or should we be—time-biased on their behalf.[23]

Again, I am not going to take a side on the issue. I simply claim that the fact that many people are near-biased often explains our intuitive thinking in cases where we have to decide whether save identified lives or statistical lives. To put that another way, it is the time bias, in this case, the near bias, that explains why we prefer to save Ann in our example over a statistical life, why we prefer to save Terry Q. over a statistical car crash victim, and why we prefer saving Jessica from the well over a statistical child that will fall into well later.

To understand the explanation, consider the following revised version of the ventilator case. Keep in mind that you know all of the following and nothing more (of any relevance to your decision).

> **Ventilator$_2$**. There is a horrible pandemic raging and a hospital where incidences are high has just one ventilator. There is no option of getting a ventilator from other hospitals and there is no possibility of patients being transferred to another hospital. At the moment, there are a hundred patients who need the ventilator and will die for sure without it but with the help of the ventilator, the person who will get it will recover and return to normal within a few weeks. However, we know that during those few weeks, Ann, who is currently infected with the virus, but not yet severely ill, will need the ventilator. Without the ventilator, Ann will die for sure, but she will fully recover with the help of the ventilator. You know some information about Ann's personal life. You know nothing about the hundred people who need the ventilator now.

What should we do in this case? Who should get the ventilator? Should we wait, let the hundred currently ill patients die, and give the ventilator to Ann when she eventually needs it? Or should we randomly choose a person among the currently ill patients, save that person and let the 99 patients and Ann die?

(Eds.), *Identified versus statistical lives: An interdisciplinary perspective*. Oxford University Press; Otsuka, M. (2015). Risking life and limb: How to discount harms by their improbability. In I. G. Cohen, N. Daniels, & N. Eyal (Eds.), *Identified versus statistical lives: An interdisciplinary perspective*. Oxford University Press; Eyal, N. Concentrated risk, the coventry blitz, Chamberlain's cancer. In I. G. Cohen, N. Daniels, & N. Eyal (Eds.), *Identified versus statistical lives: An interdisciplinary perspective*. Oxford University Press.

[15]Zameska, J. (2022). An uncertainty argument for the identified victim bias. *Journal of Applied Philosophy*, *39*(3), 504–518.

[16]Daniels, N. (2015). Can there be moral force to favoring an identified over a statistical life? In I. G. Cohen, N. Daniels, & N. Eyal (Eds.), *Identified versus statistical lives: An interdisciplinary perspective*. Oxford University Press; Hare, C. (2012). Obligations to merely statistical people. *Journal of Philosophy*, *109* (5/6), 378–390; Verweij, M. (2015). How (not) to argue for the rule of rescue: Claims of individuals versus group solidarity. In I. G. Cohen, N. Daniels, & N. Eyal (Eds.), *Identified versus statistical lives: An interdisciplinary perspective*. Oxford University Press; Slote, M. (2015). Why not empathy? In I. G. Cohen, N. Daniels, & N. Eyal (Eds.), *Identified versus statistical lives: An interdisciplinary perspective*. Oxford University Press; Frick, J. (2015). Treatment versus prevention in the fight against HIV/AIDS and the problem of identified versus statistical lives. In I. G. Cohen, N. Daniels, & N. Eyal (Eds.), *Identified versus statistical lives: An interdisciplinary perspective*. Oxford University Press.

[17]Some people might say they would like to have the medical operation done sooner rather than later, perhaps so that they do not have to anticipate it. Maybe some will think this way. However, for the purpose of the paper, whether that is the case is irrelevant. More important is that people prefer that good things happen sooner rather than later.

[18]Sullivan, M. (2018). *Time biases: A theory of rational planning and personal persistence*. Oxford University Press.

[19]Dougherty, T. (2011). On whether to prefer pain to pass. *Ethics*, *121*(3), 521–537.

[20]Lowry, R., & Peterson M. (2011). Pure time preference. *Pacific Philosophical Quarterly*, *92*(4), 490–508; Kauppinen, A. (2018). Agency, experience, and future bias. *Thought*, *7*(4), 237–245.

[21]Greene, P., & Sullivan, M. (2015). Against time bias. *Ethics*, *125*(4), 947–970.

[22]Żuradzki, T. (2016). Time-biases and rationality: The philosophical perspectives on empirical research about time preferences. In J. Stelmach, B. Brożek, & L. Kurek (Eds.), *The emergence of normative orders*, Copernicus Center Press.

[23]Hare, C. (2008). A puzzle about other-directed time-bias. *Australasian Journal of Philosophy*, *86*(2), 269–277.

There are obvious answers to these questions. We should *not* wait and give the ventilator to Ann. Instead, we should give the ventilator to one of those hundred patients who need the ventilator now—we should do that even though they are not identified lives, after all, they are just some random people who happened to be sick and need the ventilator. While they have names and personalities, they are unidentified to *you* who in this case is making the decision. While perhaps after a careful examination we conclude that we should save Ann after all, our intuitive response to the case is clear: we should save the life we can save right now.

It thus seems that it is our time bias that explains our thinking in these cases. We should save lives now rather than later— it does not matter whether the lives we save are identified or statistical lives.

3 | OBJECTIONS AND THE REPLIES

Perhaps one objects to my example. One might, for instance, claim that in the revised ventilator case, those hundred patients whom we can save now are not statistical patients because they have identities that we *could* know, we just do not know them. To put it another way, the hundred people who are currently sick are determined: but the problem really is not about (not) *knowing*; it is about (not) *being*. Thus, one might claim that in a genuine case of the problem of statistical lives, those people who will get sick should be metaphysically undetermined. It is not enough that we do not know who is sick, it should be the case that it is actually undetermined who will get sick. An objector could thus claim that my example cannot show why we prefer saving identified lives over statistical lives because I am not comparing saving identified lives versus statistical lives, since I am not comparing saving an identified person versus saving metaphysically undetermined lives.

To see if the objection has force, we can revise the ventilator case based on the objection. Consider the following case while again keeping in mind that you know all of the following and nothing more (of any relevance to your decision).

> **Ventilator$_3$**. In two months, I am about to hit Ann in the head with a baseball bat. Suppose further that Ann will be severely injured because of my attack and she will need a ventilator for a few weeks. If she gets the ventilator she will fully recover, if she doesn't, however, she will die. In addition to hitting Ann in two weeks, I am about to hit someone in the head with a baseball bat tonight. At the moment, I do not know whom I will hit but I am sure that I will hit someone. Suppose for instance, that the person I will hit is whoever happens to be the first person I see wearing a red shirt when I walk out of my apartment tonight. Suppose further that the person will be severely hurt because of my attack, and they will need a ventilator for a few weeks. If they get the ventilator, they will fully recover but if they won't they will die. You know some information about Ann's personal life. Obviously, you know nothing about the person who I will attack later tonight (other than that they are wearing a red shirt, which is morally irrelevant).

As with the earlier thought experiments, let us assume that the hospital has only one ventilator and there is no way to get more ventilators or transfer patients to other hospitals. If these are the details of the case, whom should we save? Who should get the ventilator? Should we give the ventilator to the metaphysically undetermined statistical patient tonight, or should we let them die and give the ventilator to Ann, who we know will need the ventilator in 2 weeks?

If your moral intuition aligns with mine (and I assume it does) then we should save the person who needs the ventilator tonight—even though it is yet undetermined who that person is, and even though Ann will then die.

The example shows that even in the cases where the lives are metaphysically undetermined, we intuitively prefer saving lives now versus saving lives later—even though it means that we will save statistical lives over identified lives. Thus, the identified person bias is not about identity or statistics after all—the bias is about when we save the people.

However, one might object at this point based on the following case.

> You can either save the person in front of you or you can press a button that you know will save two other unidentified people.

There is no difference in time here. Yet, arguably many people would save the one in front of them, perhaps due to the fact that the person is "identified" to you—even when not saving them would save two people that are not in front of you. Thus, while the near bias is part of the explanation in ICU triage decisions, it might not always explain the identified person bias.

To see this in detail, consider the following case.

> **Ventilator$_4$**. Tonight, I am about to hit Ann in the head with a baseball bat. Suppose further that Ann will be severely injured because of my attack and she will need a ventilator for a few weeks. If she gets the ventilator she will fully recover, if she doesn't, however, she will die. In addition to hitting Ann in the head tonight, I am about to hit someone else in the head with a baseball bat tonight also. At the moment, I do not know whom I will hit but I am sure that I will hit someone. Suppose for instance, that the person I will hit is whoever happens to be the first person I see wearing a red shirt when I walk out of my apartment tonight. Suppose further that the person will be severely hurt because of my attack, and they will need a ventilator for a few weeks. If they get the

ventilator they will fully recover but if they won't they will die. You know some information about Ann's personal life. Obviously, you know nothing about the person whom I will attack later tonight (other than that they are wearing a red shirt, which is morally irrelevant). There is only one ventilator available.

Should we save Ann or the unknown, undetermined, statistical victim? Perhaps in this case, people have the intuition that we should save Ann. If that is the case, then the identified person bias is present after all. The identified person bias is present in cases where we can save either identified life now or statistical life now. The identified lives bias also *explains* why we think we should save the life we think we should save: we should save identified life because (rightly, or wrongly) we think that it is better to save an identified rather than a statistical life. Whether it is truly morally just, is another thing.

But as my argumentation showed, in at least in some cases where the lives saved are saved at different times, the explanation behind our intuition changes. Then the time bias explains our intuitive reactions to the cases: we should save a life now because (rightly, or wrongly) we think it is better to save a life now rather than later.

However, one might also claim that it is because of differences in certainty that we prefer saving lives now rather than saving lives later. If we use the resources now and save Ann now rather than wait and save someone else later, we are using the resources in the most efficient way. If we do not save Ann but wait, it could be that some unexpected and unlikely factor contributes to the case so that the people we are planning to save in the future cannot be saved (maybe they die for other causes or maybe they will be cured because of some other reason or maybe for some reason they do not who at the hospital next week). Also one could claim that using the ventilator sooner could result in its being free to use on someone else sooner—if Ann for instance gets better sooner than was expected.[24]

The objection is not without merit. Certainty surely plays a role in whom we should save. However, I do not think it refutes the main claim of this paper. One reason for it is that uncertainty may sometimes be greater with identified lives at risk.

It may be, for instance, highly uncertain that acute care interventions on identified people will succeed now, but that prevention programs that will save statistical lives later may be highly certain.[25] For instance, consider the following case.

> **Ventilator$_5$.** There is a horrible pandemic raging and a hospital where incidences are high has just one ventilator. There is no way to get a ventilator from other hospitals and there is no possibility of patients being transferred to another hospital. At the moment, there is one patient, Ann, who needs the ventilator and will die for sure without it but with the help of the ventilator she could recover and return to normal within a few weeks. However, it is uncertain whether she will recover. In fact, the doctors have estimated that she has a 75% chance of recovering if getting the ventilator. We also know that during those few weeks, some of the people who are currently infected with the virus, but not yet severely ill, will need the ventilator. Without the ventilator, whoever falls ill within the next few weeks will die, but they will fully recover with the help of the ventilator. We know from previous cases that these statistical victims have a recovery rate of 99%. You know some information about Ann's personal life. You know nothing about the people who need the ventilator in a few weeks.

I am inclined to think that there is still the intuition of saving Ann now rather than the statistical victims later—despite Ann's likelihood of surviving is less than that of the statistical victims. But if the explanation of why we prefer saving Ann rather than the statistical victim is about the probabilities, we should be inclined to think that statistical victims should be saved in this case. However, we do not think so, it remains very plausible that it is the time bias, as argued in this paper, that plays a role in explaining the intuition on whom to save in such cases.

One could also point out that generally, waiting, rather than admitting the patient who is currently at the ICU, will leave resources unused for longer and therefore be less efficient.[26] This does not seem to be because of near bias, but rather, because of a reasonable strategy aiming to optimize health outcomes. This could very well be a part of the explanation of why identified individuals are saved rather than statistical lives. It does not mean, however, that near biases, for instance, do not have any role here. I am inclined to think that efficiency, risk aversion, near bias, and sometimes identified-person bias itself all play a role in these decisions.

4 | CONCLUSION

In this paper, I have argued that the identified person bias is often actually about time bias. I argued that while we initially thought that identifying the victim matters (Ventilator$_1$), when faced with revised cases (Ventilator$_2$ and Ventilator$_3$), we realized that it is not our ability to identify the victim that matters after all. What matters is *when* we save the people (now or in the future). Thus, our distorted attitude toward the future explains why we prefer to save people now rather than to save different people later. Of course, it could be that other factors are present also. Factors such as risk aversion and efficiency, for instance. Be that as it

[24] I thank two anonymous referees for pressing me on these issues.
[25] Brock, op. cit. note 14, p. 48.
[26] I thank the anonymous reviewer at *Bioethics* for this point.

may, time bias seems to play a role in explaining identified person bias.

ACKNOWLEDGMENTS

I would like to thank Jay Zameska and H. Orri Stefánsson for commenting on earlier versions of this paper. Thanks also to colleagues at CEPDISC for the discussion on the topic. I would also like to thank two anonymous reviewers at Bioethics for commenting on an earlier version of this paper.

CONFLICT OF INTEREST STATEMENT

The author declares no conflict of interest.

ORCID

Joona Räsänen http://orcid.org/0000-0002-7383-6138

AUTHOR BIOGRAPHY

Joona Räsänen is a postdoctoral researcher at the Centre for the Experimental-Philosophical Study of Discrimination at Aarhus University School of Business and Social Sciences. His research interests are in bioethics and applied ethics and he is a frequent contributor to bioethical journals such as the *Journal of Medical Ethics*, *The American Journal of Bioethics*, and *Bioethics*.

How to cite this article: Räsänen, J. (2023). ICU triage decisions and biases about time and identity. *Bioethics*, 1–6. https://doi.org/10.1111/bioe.13190