



The normative significance of identifiability

Tomasz Żuradzki¹

© The Author(s) 2018

Abstract

According to psychological research, people are more eager to help identified individuals than unidentified ones. This phenomenon significantly influences many important decisions, both individual and public, regarding, for example, vaccinations or the distribution of healthcare resources. This paper aims at presenting definitions of various levels of identifiability as well as a critical analysis of the main philosophical arguments regarding the normative significance of the identifiability effect, which refer to: (1) *ex ante* contractualism; (2) fair distribution of chances and risks; (3) anti-aggregationist principles that recommend the distribution of bad effects and the concentration of good ones. I will show that these arguments, although connected with interesting philosophical problems regarding e.g. counterfactuals, aggregation, or probability, are unconvincing.

Keywords Identifiability effect · Identified victim effect · Statistical lives · Uncertainty · Risk · Healthcare · Rule of rescue

Introduction¹

The concepts ‘identified individual’s life’ and ‘statistical life’ were introduced in 1968 by the American economist (and later Nobel laureate) Schelling (1968). In his article ‘The Life You Save May Be Your Own’ (which adopted the title of a previously published short story by the American writer Flannery O’Connor), Schelling noted that people are more eager to contribute to saving an ‘identified individual’s life’ than to saving a ‘statistical life’. He understood an ‘identified individual’s life’ as the life of a known person who needs help at a given moment, a ‘statistical life’ as the life of a person whose identity is not known to the decision-maker at the moment of deciding (I will discuss more detailed definitions in “[Defining levels of identifiability](#)” section). Schelling pointed out that such a preference may have a significant impact on the way public or private resources are spent. Here is a frequently quoted example from Schelling’s text regarding healthcare:

Let a 6-year-old girl with brown hair need thousands of dollars for an operation that will prolong her life until Christmas, and the post office will be swamped with nickels and dimes to save her. But let it be reported that

without a sales tax the hospital facilities of Massachusetts will deteriorate and cause a barely perceptible increase in preventable deaths—not many will drop a tear or reach for their checkbooks. (115)

This story is one of many examples—either real or hypothetical—which are still used in empirical and theoretical studies on the problem of identifiable individuals. This anecdotal evidence has been partly confirmed by empirical studies (for reviews, see: Kogut and Ritov 2015; Small 2015; Lewinsohn-Zamir et al. 2017). This is one way in which, for example, one can interpret studies whose participants were more eager to compensate a loss to others, or would give more money to a charity, if the potential beneficiaries had already been chosen at the moment of decision-making than in a situation in which the beneficiaries were supposed to be chosen after the participants made the decision. The preference was visible in a situation where the participants themselves did not know—and were aware of the fact that they would never know—any other meaningful personalising information concerning the individuals whose loss they were compensating or whom they were helping (Small and Loewenstein 2003). However, studies on this topic do not prove clearly which psychological mechanisms this preference is based on: the fact that a given person was identified *ex ante* (by a number, name, etc.), even if the decision-maker knew none of the person’s particular features, or other

✉ Tomasz Żuradzki
t.zuradzki@uj.edu.pl

¹ Institute of Philosophy, Jagiellonian University, ul. Grodzka 52, 31-044 Kraków, Poland

¹ This article develops arguments previously presented in Polish in Żuradzki (2017).

characteristics of the situation, e.g. singularity (Kogut and Ritov 2005).

Another assumption is that more people prefer to harm innocent unidentified people at the moment of decision-making than to harm identified individuals. Although the effect in this case is not as well-documented in behavioural research, it is apparently one way in which one can interpret studies that show, for example, the reduced level of acceptance of a policy requiring trade-offs between different people (e.g. affirmative action) when the cost of such a policy is borne by identified people (Ritov and Zamir 2014).²

Therefore, I assume that the identifiability effect has two main forms: (1) greater readiness to help identified people than unidentified ones; (2) greater readiness to take actions harmful to (innocent) unidentified people than to take actions harmful to (innocent) identified ones.

In this article I analyse the normative status of decision-makers' preferences of this type and consider whether they are morally justifiable or, rather, constitute certain kinds of preferential anomalies. Answering this question about the normative significance of the identifiability effect is important in the context of moral enhancement or moral technologies that 'attempts to bridge the gap between moral psychology and normative theory by recommending ways in which we, as moral psychology describes us, can become more as we should be, as normative theory prescribes' (Alfano 2013). Obviously, before we start using any moral technology we must decide in what direction are we going to steer our own or others' behaviour; the aim of this paper is to describe situations in which this task is particularly difficult, because it depends on the content of normative doctrines or accepted values. I also show how difficult it is to use knowledge of the psychological foundations of human moral behaviour to 'debias' behaviour or 'steer it toward a more rational path'—for example, in healthcare, since this process depends on what constitutes 'a rational path' as well as on the distinction between biased processes and perfectly rational reactions. Thus, before we can ask whether it is even possible or feasible to debias human behaviour in the case of the identifiability effect, we should establish its normative significance.

The identifiability effect is visible in many individual and public decisions. In the long run it may lead to inefficient use of scarce resources: decision-makers with these preferences have a lesser chance of achieving the goals they consider the most valuable, e.g. maximising the number of people saved.

This may be visible in many healthcare decisions: the rule of rescue recommends rescuing identified people whose life is in imminent danger, (Hope 2001; McKie and Richardson 2003). The significance of the identifiability effect was analysed recently for example in discussions over dividing limited resources between the prevention and treatment of AIDS in Africa. Some authors argued that it is more effective to spend a limited budget on prevention (i.e. saving unidentified people) than on the treatment of the actually ill, and that identifiability is not a morally significant factor: 'statistical lives saved are just as real as identified lives saved; all have the same equal worth' (Brock and Wikler 2009: 1671). Others answered that even if such an allocation produced the greatest good (i.e. number of lives saved), it would not be fair: concentrating on prevention could save more lives overall, but there are important moral reasons to give priority to identified people who are sick now, even if we know that this policy will save fewer lives overall. Some argued (Frick 2015a) that one of these reasons is the fact that sick people are identified (this claim does not exclude the possible existence of strong moral reasons other than identifiability to give priority to those who are sick now). Similar arguments were also put forward in a discussion of the alleged obligation to reveal incidental findings noted in the course of scientific research to interested (and identified) people, even when it interrupts the process of conducting the studies (Żuradzki 2015).

The second form of the identifiability effect is significant in the case of public undertakings which can threaten human lives or interests. For example, the economist and philosopher John Broome, discussing the problem of valuing human life in economics, noted that evaluation of the acceptability of risky public projects (e.g. large construction projects) depends on whether there is a risk to identified or unidentified (at the moment of making the decision) individuals (Broome 1978). The former case meets much stronger objections than the latter, even if it is nearly certain that someone will be harmed, or even die, during the execution of a given project. A similar effect is visible in disputes between retributivism and consequentialism in the context of the risk of unintentionally punishing (unidentified) innocents (Bystranowski 2017).

In recent years the topic has been mentioned in the context of discussions of rational and emotional components of normative judgments and the relationship between these components and normative theories of a consequentialist character (i.e. the rightness of a given action depends on actual or expected results only) or those of non-consequentialist character (i.e. rightness depends on other factors as well, e.g. intentions, consent, etc.). Some researchers assume that the differentiation of decisions or evaluations based only on the existence in one case of an identified and in another of an unidentified person has no 'rational basis', since it is an

² Interestingly, this preference is reversed when people believe in someone's guilt: experiments showed that people are ready to punish criminals more severely when they are identified in comparison to unidentified wrongdoers who have committed an analogous crime (Small and Loewenstein 2005).

effect of the automatic manifestation of our emotional preoccupation with those ‘up-close and personal’ (e.g. Greene 2007: 49). However, discussions are being conducted among philosophers on possible ways to justify the normative significance of this preference; some scholars (Daniels 2012; Hare 2012; Frick 2015b; Zaluski forthcoming) claim that the preference towards identifiable individuals is not necessarily irrational or morally wrong. This article will explore these arguments and find that they are not convincing, although they refer to interesting philosophical problems connected with e.g. interpretations of counterfactuals, aggregation, or probability.

Defining levels of identifiability

The difference between identified and statistical individuals (victims or beneficiaries) is not always clearly stated. In particular, the meaning of ‘statistical’ is not clear, and the term is often understood differently in psychological studies and in philosophical papers. Let me present the following fictional example, inspired by situations described by Hare (2012) and, previously, by Kamm (2001, Chapter 11).

Example 1: Alice and her friends (rescuing an identified vs an unidentified individual)

Five people (Alice, Barbara, Cecilia, Dorothea, Eva) are in danger. If you do nothing, Alice and one of the remaining women (either Barbara, or Cecilia, or Dorothea, or Eva) will die. Fortunately, you can take one of two actions (unfortunately, you cannot take both simultaneously): (1) one, thanks to which you will rescue Alice, or (2) the other, thanks to which you will rescue one of the four remaining women.

Who should be rescued? Are there any reasons in favour of either of these two options? I assume—both in this example and in the following ones—that you have no special relationship with and no special obligations to any of them, and that, generally speaking, you do not know anything beyond the description in the example. At first glance it may seem that it does not matter which action you take, since in either case you rescue one person, while one person dies. The only difference is that in the first case (action 1) you know *ex ante* exactly whom you will rescue, whereas in the second (action 2), you do not know *ex ante* exactly whom you will rescue; in fact, you may never learn whose life you saved.

Let us assume that you choose action 1: in this situation you save Alice (who is an *ex ante* identified beneficiary) and after that you can find out which of the remaining individuals will die (*ex post* identifiable victim). This is a case of two identifiable individuals. But if you choose action 2, the situation is more complicated: Alice will die (*ex ante* identified victim) and you save one person, although you do not

know whom. Of course it is possible that after your decision you may learn the identity of the person you have rescued (*ex post* identifiable victim), but let us assume that by taking action 2 you have certainly rescued one person out of four, but you do not know who exactly the beneficiary of this action is, that is, which of the four women would have died had you taken action 1 (*ex post* unidentified beneficiary).

An analogous situation may occur in a case of harm. Let us examine the following example inspired by a hypothetical situation described by Michael Otsuka (2015).³

Example 2: Adam and his friends (sacrificing an identifiable vs an identified individual)

Five people (Adam, Bernard, Caspar, Daniel, Emil) are in danger. If you do nothing, all of them will die. Fortunately, you can take one of two actions (unfortunately, you cannot take both simultaneously): (1) one, thanks to which you will certainly rescue Adam and three other men; however, through this action you will unintentionally contribute to the death of one of the remaining four (either Bernard, or Caspar, or Daniel, or Emil); if you choose this action, you will learn the identity of the victim after making the decision; (2) the other, thanks to which you will rescue the four remaining men (Bernard, Caspar, Daniel, Emil); however, it can be predicted that because of this action you will unintentionally contribute to Adam’s death.

Which action should be taken in this situation? Again, it may seem that it does not matter which action you take, since in either case you rescue four persons, while one person dies. Let us assume that you choose action 1: in this situation you save four people in total, but at the cost of the death of one. In fact, you do not know the identity of the victim before making the decision, but you can learn it afterwards (*ex post* identifiable victim). When you choose action 2, Adam dies (*ex ante* identified victim), but the remaining people survive and (as in the previous example) after your decision you may not learn the identity of the person you have rescued, that is, which of the four men would have died had you taken action 1. Similarly to the previous example, let us assume that after taking action 2, you are not able to determine exactly who the beneficiary of this action is (*ex post* unidentified beneficiary).

In both of the examples presented above, lack of awareness of the identity of unidentified beneficiary can be interpreted in two ways. First, it may be the result of the contingent cognitive limits of the decision-maker (including a lack

³ In Otsuka’s example the number of people in danger is much greater; the decision-maker is a public official, and the cost of rescuing this great number of people is borne by a person from outside the group.

of technological accessibility that may be overcome in the future). Given this interpretation, in Example 1, after deciding to take action 2, the result would be, for example, that if you had not chosen the action, one woman in particular (either Barbara or Cecilia or Dorothea or Eva) would have died—exactly that one, with specific DNA, although you are unable to confirm this. An analogous situation may happen in Example 2. I will call people of this kind ‘unidentifiable individuals’.

Secondly, lack of awareness of a person’s identity can result from non-contingent limitations, that is, from the indeterminacy of conditions of certain natural processes. In this case, in Example 1, after deciding to choose action 2, it would be no fact of the matter whom you rescued—in order to establish this information, you would have had to make a different decision previously, choosing action 1. Similarly, in Example 2, after deciding to choose action 2, it would be no fact of the matter who would have died had you chosen action 1. I assume that we are dealing with the statistical individual *sensu stricto* in the case of this second interpretation only, that is, when lack of awareness of the identity of the beneficiary does not result from the contingent inability of an agent or a technology to check someone’s identity, but from the very fact that after choosing one action, there is no such fact of the matter that would correspond to statements about what would have happened had we chosen differently. It is worth noting that this second interpretation is not only a hypothetical scenario discussed by philosophers, but also seems important in many real-life situations related to healthcare. For example, it is reasonable to assume that, following vaccination of a large population, there is no fact of the matter who exactly would have died had we not vaccinated this population. This may also be true of many preventive actions in healthcare.⁴

It is worth summing up our terminological distinctions, since all of these terms have been used inconsistently in the literature (both philosophical and psychological). In my paper *an identified person* is an individual who has been in some way identified *ex ante* by a decision-maker. In contrast, a person whose identity cannot be established *ex ante* is an *ex ante unidentified* individual. This class may include three types. A *ex post identifiable person* is an individual who may be in some way identified *ex post* by a decision maker. A person who may not be identified even *ex post* may be either *an unidentifiable* or *a statistical individual*. The former means that the person’s identity cannot be established only because of the cognitive or technical limitations of a decision-maker;

the latter means that an individual in a group cannot be identified for certain genuine non-contingent reasons: there is no fact of the matter statements concerning this person could refer to, thus the identity of this person could have been established only if a different decision had been made in the past.

Although the terms ‘statistical lives’, ‘statistical individuals’, ‘statistical victims’, etc. often appear in the psychological literature (Small and Loewenstein 2003, 2005; Small 2015; Kogut and Ritov 2015), it is worth noting that none of the empirical studies of which I am aware has distinguished these levels of identifiability, and none has traced the differences in terms of human decisions, evaluations, emotions, or cognitions between them. In particular, psychological scenarios do not distinguish between a lack of awareness of identity resulting from the contingent limitations of a decision-maker and a lack of awareness resulting from the indeterminacy of conditions; that is, they fail to distinguish between statistical and unidentifiable individuals.

In the next chapter I will show that the difference between identified individuals and those who cannot be identified *ex ante* (because they are either identifiable *ex post*, or unidentifiable, or statistical individuals) is treated as normatively significant by some philosophers, i.e. how they differentiate the strength of reasons for action on the basis on the applicable type of identifiability.

Justifications of the normative significance of identifiability

Are there any reasons to prefer to rescue an identified (Alice) rather than an unidentified (one of the four remaining women) person in the above-mentioned situation (Example 1)? Or, in Example 2, to prefer to unintentionally harm an *ex ante* unidentified (but *ex post* identifiable) person rather than to harm an *ex ante* identified person (Adam)? Do moral reasons depend on the level of identifiability (*ex post* identifiable, unidentifiable, statistical)? In recent years some philosophers have defended the normative significance of identifiability and have tried to show that in situations similar to Examples 1 and 2 there are specific reasons to choose action 1, thereby rescuing an identified person (Example 1), or harming a person that is unidentified at the very moment of decision-making (Example 2).

Ex ante contractualism

Let us leave aside the obvious case of act consequentialism, in which differences in identifiability do not matter (what we ought to do depends only on the actual or expected state of the world) and start with contractualism. The already classic

⁴ E.g.: ‘This is the case, for instance, with the treatments for coronary heart disease and hypertension mentioned in the outset: it is impossible to say, both before and after treatment, whom in the population that actually has benefited from the treatment, i.e., who would not have survived was it not for the treatment’ (Juth 2017: 5).

statements of Scanlon include the following definition of ethical contractualism:

[T]hinking about right and wrong is, at the most basic level, thinking about what could be justified to others on grounds that they, if appropriately motivated, could not reasonably reject. (Scanlon 1998: 5)

In considering whether a principle could reasonably be rejected we should consider the weightiness of the burdens it involves for those on whom they fall, and the importance of the benefits it offers for those who enjoy them, leaving aside the likelihood of one's actually falling into either of these two classes. (Scanlon 1998: 208)

In the popular interpretation, Scanlon refers in these fragments to people who will actually suffer or benefit from a given action (Reibetanz 1998). It means that the obligation to justify, to every person, the rule which was the basis of a given action is not related to the likelihood of gain or loss before the action was taken (*ex ante*), but to a hypothetical situation in which we would know what state of the world result from a given action. This interpretation suggests that, in my first two examples, contractualism of this type would not prefer either of the possible actions: one would deal with Alice's claim versus the claim of one of the remaining women (or Adam's claim versus the claim of one of the four men); whether the decision-maker knew or didn't know the victim's identity before making the decision would not matter at all.

One of the most controversial elements of Scanlon's *ex post* contractualism is its demand for a comparison of the claims of people in different possible worlds that may exist as the result of our actions, depending on who is the beneficiary and who is the victim in a given possible world, and therefore this view is inconclusive in many situations that are more complicated than my two initial examples (Frick 2015b: 196). This inconclusiveness stems from the fact that almost all our actions may lead to many, or even an infinite number of, possible scenarios. *Ex post* contractualism requires taking into account claims from all of these possible worlds, even from those that are extremely improbable (this view does not discount claims due to their improbability), and fulfilling the strongest individual claim (it also prohibits aggregating individual's claims). Thus, the *ex post* perspective renders contractualist deliberations extremely risk-averse and contradicts many of our ordinary moral convictions, since we often accept risky actions in large populations that certainly lead to death (e.g. when we vaccinate newborns against certain serious but non-fatal childhood

diseases, accepting that there is a remote chance of fatal side effects from the vaccination itself). Therefore, some authors have criticised Scanlon's views, without rejecting contractualism itself, suggesting comparing the *ex ante* rather than the *ex post* claims of the interested parties (John 2014; Frick 2015b). Below, I present two examples inspired by Frick (2015b) which may be helpful in distinguishing a different type of contractualism.⁵

Example 3: Mass vaccination (identified vs statistical individuals)

A million children are threatened by a viral disease that will kill them all if you do nothing. You have to choose between one of two vaccinations (you cannot have both).

- Vaccination 1 is a sure thing: it will certainly save the life of every child. However, it does not completely protect children from a virus that results in incurable paralysis of one leg of every child. After choosing this vaccination it is impossible to determine who would have died if Vaccination 2 had been chosen instead (i.e. it is statistically certain that, with Vaccination 1, some children, more or less 1000, are statistical beneficiaries).
- Vaccination 2 is risky: on the basis of previous cases, doctors estimate that thanks to it every child will have a 0.999 chance of being completely healed. However, every child runs a 0.001 risk that this vaccination will not work and that the virus will cause the child's death; I assume that probabilities are independent⁶ in the case of each child (victims are unidentified *ex ante*, but they will be identifiable *ex post*).

Example 4: Mass vaccination (identified vs identified individuals)

The situation is identical to Example 3, except that:

⁵ However, Frick (2015b) does not differentiate between what I have called statistical and unidentifiable individuals, and he uses the term 'statistical lives' to refer to both types.

⁶ There is a significant difference between Examples 1 or 2 and Example 3 regarding the dependence or independence of probabilities: In the two initial examples the probabilities were dependent (it was sure that one and only one person would die); here, I assume that they are independent (it may happen that if we choose Vaccination 2, more or less than 1000 children will die, or even—although this would be extremely improbable—all of them, or none of them). Does this difference matter? I assume that this is not normatively important because the difference between the dependence or independence of probabilities is not as clear as it may seem at first sight (see below and see also: Otsuka 2015: 90).

- Vaccination 2a will save 999,000 children out of a million. However, due to a particular genotype it will certainly not work for a thousand children, those with a certain genotype, who will die as a result of this viral disease and who can be identified *ex ante* by the decision-maker.

What should we do in these situations? In both of them, if we use Vaccination 2 (or 2a), a similar number of children will most probably die. In Example 3, the probability that 1000 children (± 100) will die is very high (0.99), and the probability that at least 1 child will die is extremely high or even statistically certain ($1 - (999/1000)^{1,000,000}$). In Example 4, it is known for certain that 1000 children with a certain genotype will not survive the disease (identified victims). These examples can easily illustrate the difference between the three types of views, two of which do not recognise the normative significance of the preference towards identified individuals, one of which (the third) does. Below, I enumerate these two views and discuss the third one.

The first is the *ex post* view referring to the principle of full information. Here identifiability is not itself normatively significant, so in both the above Examples (3 and 4), Vaccination 1 should be used. This view recommends acting in a way that substantially improves the outcomes of those who otherwise would end up as worst off, no matter if they are identified or not for a decision maker (Fleurbaey and Voorhoeve 2013, their slogan is: ‘decide as you would with full information’).

Second, contractualism behind the hypothetical veil of ignorance along with the expected utility rule. This view is also agnostic about the normative significance of identifiability, since it claims that because of fairness a (public) decision maker should always act as if he didn’t know people’s identities (even if in fact he knows them), therefore in both the above Examples (3 and 4), the second Vaccination should be used (2 or 2a). This is so because this view also assumes the expected utility rule: in this particular case it is reasonable to assume that a choice that involves the tiny risk of death (0.001) would have a higher expected value for every child than a choice that involves the certain and incurable paralysis of one leg. Thus a (public) decision maker should maximise the expected utility on behalf of each child separately because it would be in every child’s interest (Harsanyi 1977).

Since the first two views do not consider identifiability to be practically significant, I want to concentrate on another view that finds it important.

This third view is *ex ante* contractualism behind the natural veil of ignorance. In contrast to the second view, *ex ante* contractualism ascribes an important normative role only to the lack of knowledge of a decision maker that is the result of ‘natural’ processes, and not to the result of a hypothetical (or forced) veil of ignorance in the sense proposed by Rawls

(1999) and used by Harsanyi (1977). Similarly to the second view, it appeals to an argument from a single person case: ‘the rightness of an action is a function of each individual’s personal reasons for rejecting a principle that licenses the action. An act is wrong if and only if there is someone who can complain that we failed to treat her in a way that was justifiable to her, not because its consequences were impersonally bad.’ (Frick 2015b: 187). Therefore, this view assumes the following rule:

Rule 1: Satisfy the strongest individual claim.

Satisfying the strongest individual claim known *ex ante* by a decision-maker is a *pro tanto* moral reason that contributes to making an action right or wrong, all things considered. (Frick 2015b: 216, 220; cf.; John 2014: 35)

This rule implies that identifiability is normatively significant, therefore in Example 3, Vaccination 2 should be used; whereas in Example 4—Vaccination 1. It is so because both these Examples might be decomposed into one million single-person choices occurring in parallel. Due to the lack of knowledge regarding the identity of potential victims, the decision-maker in Example 3 might compare the claims of someone who would certainly be afflicted with a paralysed leg (Vaccination 1) with the claims of someone who would be exposed to a minor risk of death (0.001) during treatment (Vaccination 2). In this case, a decision maker could assume that every reasonable agent would prefer to be exposed to a minor risk of death and thus Vaccination 2 is morally preferable, because, as the definition says, no one ‘could reasonably reject’ this decision *ex ante*.⁷

Nevertheless, the situation would be different in Example 4 because a decision-maker would know (or could easily learn) the identity of carriers of a certain genotype before making the decision. In this case, a decision maker should compare the claims of someone who would certainly be afflicted with a paralysed leg (Vaccination 1) with the claims of a representative of the identified (or easily identifiable) 1000 children who would certainly die after taking Vaccination 2. Therefore, the strongest claim in Example 4 would obviously be made by a representative of potential victims and thus we should (morally) prefer Vaccination 1.

The main problem with *ex ante* contractualism is related to the ambiguity of the concept of the ‘natural’ veil of

⁷ However, Frick emphasises that *ex ante* contractualism should not be treated as the sole rule determining what is right and what is wrong, but one of many right/wrong-making properties. This means that sometimes other properties, in particular the number of saved people, can prevail over the normative significance of identifiability (meaning that, in examples such as ‘one identified vs x statistical persons’, Frick may assume that rescuing a greater number of *ex ante* unidentified persons is preferable, if only x is a sufficiently large number).

ignorance. As I demonstrate, the expected ‘naturalness’ of the veil of ignorance is problematic in cases that are intermediate between Examples 3 and 4, in which we can be sure that some as yet unidentified individuals exist and will be harmed because of our actions, but in which, due to cognitive limitations, it is impossible to obtain this information before a decision. Let us view this problem by means of another example.

Example 5: Mass vaccination (identified vs unidentifiable individuals)

- Vaccination 1, as in Example 3, is a sure thing, but after choosing this vaccination it is impossible due to technical reasons alone to determine who would have died if Vaccination 2b had been chosen (if Vaccination 1 is chosen, they are unidentifiable beneficiaries).
- Vaccination 2b is risky but will save most children. However, it will certainly not work for carriers of the very rare gene X, who will die as a result of this viral disease. In large populations, usually about 1000 children in one million are carriers of X, but due to certain technical limitations, the decision-maker has not identified *ex ante* carriers of X in this population, neither have the carriers of X identified themselves as potential victims (they, though unidentified *ex ante*, will be identifiable *ex post*).

In this example, a decision maker initial lack of knowledge of their identity results from technical limitations alone, so there are no statistical people here, only unidentifiable ones. In this case, *ex ante* contractualism would assume that the question of which vaccination should be used depends on the very nature of these technical limitations. If a genetic test for the presence of gene X in Example 5 were ‘simple and costless’ (Frick 2015b: 193), we should use Vaccination 1, as this last example would be similar to Example 4. Contrastingly, if the test were too complicated or expensive to carry out, we should use Vaccination 2b, as this example in such a case would be much closer to Example 3.⁸ Why? Because *ex ante* contractualism differentiates decisions not on the basis of a distinction between statistical and unidentifiable people (as defined in “Defining levels of identifiability” section), but based on the ‘natural’ availability of information to a decision-maker about the identity of the people involved. In this example, even if we failed to identify carriers of X and we used Vaccination 1, there still would be a fact of the matter of who is a carrier of X and who is not. Therefore, in Example 5, the reasons for choosing Vaccination 2b are

based not only on a specific moment of comparing claims (*ex ante*), but also on the state of knowledge, technical expertise, and subjective perspective of the decision-maker. For example, if determining the genotype of the people in Example 5 were made easier as a result of some new scientific discovery, this could have normative implications for a decision-maker. Moreover, according to this view, in many cases evaluation of the rightness of a decision in this type of situation may depend on the inclination (or aversion) of a given decision-maker to acquire new information on the identity of the victims.

Another problematic aspect of this view is how to interpret what it means ‘to identify’ an individual. In the above quoted psychological research by Small and Loewenstein (2003), the mere awareness of a decision-maker that someone has been already been chosen (with no further information expected) was enough for many people to favour the person identified in such a way in comparison with a person who was expected to be chosen from a group. But ‘knowing’ someone’s identity in such a way is meaningless. Similarly meaningless is knowing someone only by her surname, ID or genotype (as in Example 4). Is there any genuine moral reason to prefer helping some people over others solely on the basis of such meaningless information? I do not think so. So maybe ‘to identify’ someone it is necessary to establish some relation, for example to have a direct contact with her, e.g. an eye contact that enable to see her body or its particular fragment (e.g. is seeing someone’s face more meaningful than seeing someone’s finger?). Or, maybe we must know something about her e.g. something about her preferences, views, personal history, career. In contrast with Hare (2016) who writes that ‘it [morally] matters whether you know who is who’ I do not find such information morally relevant, in particular in the case of a public decision maker.

The next significant problem with this view is related to the violation of the seemingly obvious rule which states that a decision-maker (in particular a public decision-maker) should rescue a larger rather than a smaller number of people, in particular in a context similar to the above examples, which refer to public health decisions (but see Taurek 1977). However, according to *ex ante* contractualism, there are situations in which rescuing x people identified by a decision-maker is more important than reducing individual risks of death in a group of people, even if it is known in advance that this second policy will lead to a greater number (i.e. more than x) of deaths of *ex ante* unidentified people. Frick surprisingly claims that this view ‘is consistent with many people’s intuition’ (p. 219). Accepting such a rule in the case of public decision-makers would mean that their actions would not be required to maximise the number of saved people, solely because the decision-makers would have some irrelevant or meaningless information about some of them (such as ID numbers, names, etc.) at the moment of decision.

⁸ Frick (2015b: 193) refers to procedures that are one the one hand ‘simple and costless’ but on the other ‘extremely costly’, so his distinction seems not to be exhaustive.

This, in contrast to Frick's opinion, is obviously a controversial claim. Although knowledge of identity is certainly an important factor in personal relations and may have an influence on the psychological aspects of the decision-making process, *ex ante* contractualism does not vindicate its normative significance in the case of public decisions.

Distributive justice and the concentration of risks and chances

Another interesting explanation of the normative implications of the effect of identified individuals was presented by Daniels (2012; see also Frick 2015a). His starting point was a story similar to the one in Example 1. He argued that in this kind of situation there is an additional reason to rescue Alice, only because 'the concentration of risk matters morally' (Daniels 2012: 41). At the moment of decision-making, Alice's situation is worse than that of any of her four friends: she will die without medication, but if we decide to rescue her, each of her four friends still has a 0.75 probability for survival.⁹ Daniels also suggested that the stronger the reasons for rescuing Alice, the wider the distribution of risk within the group of unidentified victims. For example, if there were 100 women instead of four, and the risk of dying was 0.01, the reasons to rescue Alice rather than reducing the risk of death for each of those 100 women would be even stronger. It is worth noting that the problem does not lie in awareness of the probability of death or in possible psychological costs: his view does not depend on whether the interested parties are aware of the danger. In this type of situation, he seems to accept the following rule of distributive fairness:

Rule 2: Distribute chances and risks fairly.

Helping those who have a smaller chance of survival in a given situation is a *pro tanto* moral reason that contributes to making an action right or wrong, all things considered.

Philosophers who seemingly accept such a rule (Broome 1984; Saunders 2008; James 2013) perceive the chance to survive or the risk of death as separate values that are subject to distribution, similarly to other goods, due to the requirements of fairness. For example, John Broome argued in favour of the requirement of a random method of the allocation of indivisible good in situations in which more than one person submits a claim to this good and there are no other reasons for ascribing it to one given person rather than another. Another follower of this approach summarised this view: 'I believe the chance of a benefit is itself a benefit;

after all, almost everyone prefers to have any chance rather than none' (Saunders 2008: 367). However, such a view is controversial and the following easy counterargument may be put forward against it: 'Why isn't living itself the relevant good here, rather than chances of living? If I die then I am not significantly better off for having had a high chance of living. If I live then I am not significantly worse off for having had a low chance of living' (Hare 2012: 385; cf.; Henning 2015).

It is, however, easy to prove that there are situations in which Rule 2, along with some non-controversial premises, would lead a decision-maker to favour *ex ante* unidentified people and not *ex ante* identified people; thus, apart from other problems related to the perception of a chance to survive as an independent good in itself, whose distribution is subject to evaluation, this solution would not even accomplish the goal it was designed to achieve. Let us examine this on the following example, inspired by a situation described by Hare (2012).

Example 1a: Anthony and his friends (rescuing an identified vs an unidentified individual)

Five people are in danger (Anthony, Bernard, Caspar, Daniel, Emil). If you do nothing, everyone will die. Fortunately, you can take one of two actions (but, unfortunately, you can't take both at the same time): (1) one, thanks to which you will certainly rescue Anthony (the remaining four men will die); (2) the other, thanks to which you will rescue one of the four remaining men; however, at the very moment of decision-making you do not know which one; you will learn his identity after the event (Anthony and the three remaining men will die).

In this example, all five people are in the same situation: if you choose action 1, none of the four men (Bernard, Caspar, Daniel, Emil) has any chance of survival. If you choose action 2, Anthony is the one who has no chance to survive, but each of the remaining four men has a 0.25 probability of staying alive. Accepting Daniels's rule that the distribution of chances is of moral importance and adding the non-controversial rule of aggregating identical claims (in a situation when every person from a larger group has a claim identical to a claim of every person from a smaller group, the size of the group is decisive and one should help a larger rather than a smaller group of people), in this case the *ex ante* unidentified person should be rescued (one of the four men) rather than Anthony, an *ex ante* identified person. Why? If you take action 1, it is obvious that you will rescue Anthony (an identified beneficiary), but it is obvious that the other four men, having no chances to survive, will die (*ex ante* identified victims). However, if you take action 2, it is obvious that Anthony will die (*ex ante* identified victim), but simultaneously, you are not only saving one of the four

⁹ The main difference is that Daniels assumes that the probabilities are independent but see footnote 6.

remaining men, but you are distributing a chance of survival amounting to 0.25 among four people. Since Daniels assumed that the chances of survival constitute an independent good subject to fair distribution, in this case it would be much fairer, in accordance with his view, to distribute the chances to survive among four people than to ‘give’ all chances to Anthony alone. Therefore, the problem in this case is that in some situations Daniels’ proposal would lead to favouring unidentified people, which is contrary to the goal it was created to achieve.

Concentrating good effects and distributing bad ones

Another way of defending the normative significance of the identifiability effect was presented by Hare (2012). Contrary to the previous two views, he assumes that being a statistical individual, in the strictest sense, is of crucial significance. In contrast, being classified as *ex post* identifiable or as unidentifiable is not morally significant. Hare developed two rules that, in his opinion, lie behind the normative significance of preference for identified victims.

Rule 3: Concentrate good effects.

‘Other things being equal, given a choice between doing something very good for one person but very bad for one person, and doing something quite good for four people but very bad for one person, you ought, other things being equal, to do the former’.

Rule 4: Distribute bad effects.

‘Given a choice between doing something very good for one person but very bad for one person, and doing something very good for one person but quite bad for each of four people, you ought, other things being equal, to do the latter’ (Hare 2012: 387).

These rules would work in the two initial Examples 1 and 2 as follows (as long as they are interpreted as concerning statistical individuals *sensu stricto*). The result of action 1 is very good for one person (Alice/Adam), very bad for one person (one of four remaining persons, i.e. the one who will die) and neutral for the three remaining people (they will survive regardless of the action we choose). The result of action 2 is only moderately good for each of four people (because it only increases the probability of their survival from 0.75 to 1) and very bad for one person (Alice/Adam). The key question is: why is action 2 only moderately good for four people but not very good for one person in Examples 1 and 2? This is because statistical beneficiaries *sensu stricto* are assumed here: after taking action 2, there is no such fact of the matter to which the statements regarding the people we saved would refer. Therefore we cannot claim that we rescued one identified person, but only that we increased the probability of survival of each of four identified people

from 0.75 to 1¹⁰, which means that we certainly rescued one statistical person, but it is impossible to know who exactly was rescued. Therefore, in this case, if one accepts Rules 3 and 4, one should choose action 1 in Examples 1 and 2. In both these cases our action is very bad for one person; the actions in these examples do not differ in this respect. However, according to Rule 3, it is better to do something very good for one person than something moderately good for four people.

In Example 1a, discussed in the previous subsection (provided that it is interpreted as concerning statistical individuals *sensu stricto*), the result of action 1 would be very good for one person (Anthony) and moderately bad for each of the remaining four people (the probability of survival is reduced from 0.25 to 0 for each of them). The result of action 2 is very good for one person (the one that will survive), very bad for one person (Anthony), and neutral for the three remaining people (they will die regardless of which action we choose). In this case also, one should choose action 1, because Rule 4 says it is better to do something moderately bad for each of four people than something very bad for one person (moreover, in both cases we do something very good for one person, so the actions do not differ in this respect).

Those two rules are also interesting in terms of the dilemmas from Examples 3–5. In Example 3 they provide additional reasons (different from *ex ante* contractualism) to favour Vaccination 2. After choosing vaccination 1 there is no such fact of the matter that would refer to a statement concerning which of the one million children we rescued (as previously, the point of reference is another possible action rather than the state of the world, in which we do not intervene). Therefore choosing vaccination 1 we are not doing anything very good for any group of children, but something minimally good for each of one million children. However, no argument of this kind exists in either of Examples 4 and 5, since, after choosing Vaccination 1, in both of these examples there would be such a fact of the matter that somebody was rescued (although in Example 4 we would not know this person, but in Example 5 we would). It means that in both of these cases (Examples 4 and 5), in choosing vaccination 1 we would do something very good for 1000 existing children.

The most important problem regarding this view refers to the distinction between statistical beneficiaries *sensu stricto* and the impossibility of their identification by a decision-maker. The distinction, though theoretically very precise, may seem morally arbitrary. The way the vaccination works in the above-mentioned examples provides an additional moral reason: if it works in an indeterminate way, one should

¹⁰ Let me remind the reader that, for simplicity, the probabilities in these two initial examples are dependent.

choose Vaccination 2; if it works only for people with particular genetic material, one should choose Vaccination 1. A similar problem is connected with Examples 1 and 2: if we interpret them so that they refer to unidentifiable individuals whose identity is not known to the decision-maker, those two rules give us no additional reason to choose action 1 over action 2 in either case.

Another doubt regarding this view is related to the normative weight of the two rules presented above. Hare writes clearly: ‘Other things being equal, our obligation to aid two “merely statistical” people is stronger than our obligation to aid one “identified” one’ (Hare 2012: 389), suggesting that he limits the normative meaning of the identified victim’s effect to a very narrow group of hypothetical situations in which we would have a choice between rescuing either one statistical or one identified person. This means that the practical significance of his rules would be rather negligible and they would have no application to the real problems connected with medical care, vaccinations, or research ethics.

Conclusion

In this paper I have discussed three main propositions that attempt to vindicate the normative significance of identifiability. In each case, I showed that the arguments in favour of the normative implications of the effect are problematic. However, the expectation that decision-makers will treat unidentified individuals equally to identified ones—which is the conclusion of this paper—is problematic as well. First, it was argued that informing the decision-maker about this psychological phenomenon reduces the number of decisions to make helpful donations, probably because the awareness of the identifiability effect undermines feelings of sympathy, and, as a result, reduces the total sum of donations (Small et al. 2007). Secondly, many people treat unidentified individuals (and in particular statistical individuals) as fictional characters. If we want to use moral technology to enhance decisions about unidentified people, we cannot require morally sensitive agents to be akin to the character played by Mia Farrow in Woody Allen’s film *The Purple Rose of Cairo*, who was undisturbed by the fact that her beloved was fictional and who excitedly told a friend: ‘I just met the most wonderful man (...) He’s fictional, but you can’t have everything!’

Acknowledgements This research was funded by the Ministry of Science and Higher Education in Poland, National Programme for the Development of Humanities, No. 0068/NPRH4/H2b/83/2016, and benefited from a research stay at the Fondation Brocher (<http://www.brocher.ch>). Earlier versions were presented at the conference Designing Moral Technologies: Theoretical, Practical and Ethical Issues in Ascona; Society for Applied Philosophy Annual Conference in Edinburgh; the Tilburg Centre for Logic, Ethics, and Philosophy of Science; the Joint Session of the Aristotelian Society & the Mind Association in Oxford. I am grateful to the audiences for their comments, in particular this work benefits from discussions with my colleagues at the

Jagiellonian University and anonymous reviewers of *Ethics and Information Technology*.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Alfano, M. (2013). *Character as moral fiction*. Cambridge: Cambridge University Press.
- Brock, D. W., & Wikler, D. (2009). Ethical challenges in long-term funding for HIV/AIDS. *Health Affairs*, 28(6), 1666–1676.
- Broome, J. (1978). Trying to value a life. *Journal of Public Economics*, 9(1), 91–100.
- Broome, J. (1984). Selecting people randomly. *Ethics*, 95(1), 38–55.
- Bystranowski, P. (2017). Retributivism, consequentialism, and the risk of punishing the innocent: The troublesome case of proxy crimes. *Diametros*, 53, 26–49.
- Daniels, N. (2012). Reasonable disagreement about identified vs. statistical victims. *Hastings Center Report*, 42(1), 35–45.
- Fleurbaey, M., & Voorhoeve, A. (2013). Decide as you would with full information! In N. Eyal, S. Hurst, O. F. Norheim & D. Wikler (Eds.), *Inequalities in health: Concepts, measures, and ethics* (pp. 113–128). Oxford: Oxford University Press.
- Frick, J. (2015a). 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 persons: An interdisciplinary perspective* (pp. 182–201). Oxford: Oxford University Press.
- Frick, J. (2015b). Contractualism and social risk. *Philosophy & Public Affairs*, 43(3), 175–223.
- Greene, J. D. (2007). The secret joke of Kant’s soul. In W. Sinnott-Armstrong (Ed.), *Moral psychology. The neuroscience of morality: Emotion, brain disorders, and development* (pp. 35–117). Cambridge (MA): MIT Press.
- Hare, C. (2012). Obligations to merely statistical people. *Journal of Philosophy*, 109(5/6), 378–90.
- Hare, C. (2016). Should we wish well to all? *Philosophical Review*, 125(4), 451–472.
- Harsanyi, J. C. (1977). Morality and the theory of rational behavior. *Social Research*, 44(4), 623–656.
- Henning, T. (2015). From choice to chance? Saving people, fairness, and lotteries. *Philosophical Review*, 124(2), 169–206.
- Hope, T. (2001). Rationing and life-saving treatments: should identified patients have higher priority? *Journal of Medical Ethics*, 27(3), 179–185.
- James, S. M. (2013). When helping the victim matters more than helping a victim. *Utilitas*, 25(1), 32–45.
- John, S. (2014). Risk, contractualism, and Rose’s ‘Prevention Paradox’. *Social Theory & Practice*, 40(11), 28–50.
- Juth, N. (2017). For the sake of justice, should we prioritize rare diseases? *Health Care Analysis*, 25(1), 1–20.
- Kamm, F. M. (2001). *Morality, mortality* (Vol. II). *Rights, duties, and status*. Oxford: Oxford University Press.
- Kogut, T., & Ritov, I. (2005). The “identified victim” effect: an identified group, or just a single individual? *Journal of Behavioral Decision Making*, 18(3), 157–167.

- Kogut, T., & Ritov, I. (2015). Target dependent ethics: discrepancies between ethical decisions toward specific and general targets. *Current Opinion in Psychology*, 6, 145–149.
- Lewinsohn-Zamir, D., Ritov, I., & Kogut, T. (2017). Law and identifiability. *Indiana Law Journal*, 92(2), 505–555.
- McKie, J., & Richardson, J. (2003). The rule of rescue. *Social Science & Medicine*, 56(12), 2407–2419.
- 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 persons: An interdisciplinary perspective* (pp. 77–93). Oxford: Oxford University Press.
- Rawls, J. (1999). *A theory of justice*. Oxford: Oxford University Press.
- Reibetanz, S. (1998). Contractualism and aggregation. *Ethics*, 108(2), 296–311.
- Ritov, I., & Zamir, E. (2014). Affirmative action and other group trade-off policies: Identifiability of those adversely affected. *Organizational Behavior and Human Decision Processes*, 125(1), 50–60.
- Saunders, B. (2008). The equality of lotteries. *Philosophy*, 83(3), 359–372.
- Scanlon, T. M. (1998). *What we owe to each other*. Cambridge (MA): Harvard University Press.
- Schelling, T. C. (1968). The life you save may be your own. In S. B. Chase (Ed.), *Problems in public expenditure analysis* (pp. 127–176). Washington: Brookings Institution.
- Small, D. A. (2015). On the psychology of the identified victim effect. In I. G. Cohen, N. Daniels & N. Eyal (Eds.), *Identified versus statistical persons: An interdisciplinary perspective* (pp. 13–23). Oxford: Oxford University Press.
- Small, D. A., & Loewenstein, G. (2003). Helping a victim or helping the victim: Altruism and identifiability. *Journal of Risk and Uncertainty*, 26(1), 5–16.
- Small, D. A., & Loewenstein, G. (2005). The devil you know: The effects of identifiability on punishment. *Journal of Behavioral Decision Making*, 18(5), 311–318.
- Small, D. A., Loewenstein, G., & Slovic, P. (2007). Sympathy and callousness: The impact of deliberative thought on donations to identifiable and statistical victims. *Organizational Behavior and Human Decision Processes*, 102(2), 143–153.
- Taurek, J. M. (1977). Should the numbers count? *Philosophy & Public Affairs*, 6(4), 293–316.
- Załoski, W. (forthcoming). The moral status of helping and the identified victim effect. *Principia*.
- Żuradzki, T. (2015). The preference toward identified victims and rescue duties. *American Journal of Bioethics*, 15(2), 25–27.
- Żuradzki, T. (2017). Normatywne implikacje preferencji wobec osób zidentyfikowanych. *Diametros*, 51, 113–136.