

**The Ethics of the Rule of Rescue: Guidelines for Use in the
Medical Setting**

Duke University

For Graduation with Distinction in Philosophy

By Spencer Flynn

TABLE OF CONTENTS

Chapter One: Background and Empirical Considerations	Pg. 3
Introduction, 3	
Describing Rescue, 7	
Psychology of Rescue, 16	
Chapter Two: Considerations from Ethical Theory	Pg. 25
Rescue Duties, 25	
Consequentialist Considerations, 35	
Deontological Considerations, 45	
Chapter Three: From Theory to Practice	Pg. 57
Procedural Justice, 57	
Checklist of Considerations, 60	
Case Studies, 61	

Chapter One: Background and Empirical Considerations

1. Introduction

Humans are naturally fascinated with emergencies and rescues. This fascination consumes our imaginations—some of our oldest, and most popular, narrative arcs focus on a hero narrowly rescuing a group from some great villain or disaster.¹ It defines how we characterize bravery—we laud and praise the firefighter leaping into flames, the soldier charging into battle to save their comrades. But even more, we are obsessed with rescues because of what they represent: a moment of great moral significance. The ancient Confucian, Mengzi, writes that any person will stop to help a child who has fallen in a well, for no other reason than a feeling of benevolence and a sense of moral obligation (Mengzi 2A6, Van Norden 2008, p. 46). There is an imminent, identifiable threat, so we feel a strong impulse to help.

Psychological research is clear that such an impulse exists and exerts a powerful influence over our behavior (Latane & Darley, 1970; Piliavin et. al, 1981; Jenni & Loewenstein, 1997). Real-life examples of the impulse to rescue are abundant and oftentimes, beautiful. There was the heart-wrenching evacuation of thousands of British soldiers from the beach of Dunkirk in the Second World War by hundreds of civilian boats. There were massive relief efforts to famine-stricken areas of Ethiopia in the 1980s. There are numerous heart-warming stories about children who need a medical procedure to survive, and communities who gathered together to fund their rescue operation. Beyond these dramatic examples, there are also countless little instances of rescue—many people feel compelled to save even insects from their dangers.²

¹ Potential examples are numerous: Beowulf slaying Grendel, Luke Skywalker saving the Rebellion from the Death Star, and of course almost every superhero movie or comic book.

² Each of these rescues described—population level rescues, individual rescues, and non-human rescues—all have their own relevant ethical considerations. My point in listing them here is not to conflate them as entirely similar, but simply to show the range of circumstances in which rescue impulses are elicited.

Yet acting on impulses to rescue does not always yield desirable—or moral—results. We often continue to operate on the terminally ill, even when they seem far past saving and the treatment causes excruciating pain. We organize expensive rescue efforts for plane or boat crashes with little hope of rescue, when the same amount of money could have saved dozens—if not hundreds—from malaria (Creadon, 1997). The United States spends exorbitantly on high-intensity treatments in place of more cost-effective preventive care (Squires, 2012). The stories that elicit the greatest sympathy often reflect our own biases – we are much more likely to donate and lobby to save an American citizen who looks like us than a Syrian refugee who does not. The impulse to rescue can be a powerful motivator to act altruistically, but it can also be an instinct to focus on a more immediate but lesser need at the expense of either equity or the greater good.

Striking the appropriate balance between how we utilize our impulse to rescue versus concerns about equity and rational cost-benefit analysis, then, is important. Nowhere is this balance more important than in medicine, where questions of who to help and when arise countless times each day. Bioethicists leverage accounts of duties to rescue for numerous purposes, including to establish the need for researchers to release urgent incidental findings (Miller et. al, 2008), provide ancillary care (Merritt et. al, 2010), argue for the duty to become an organ donor (Peters, 1986) and to critique placebo controls in clinical trials in developing countries (Hawkins, 2006). Given the practical and theoretical importance of duties to rescue, the bioethics community has engaged in a lively conversation on “the Rule of Rescue” (RR). The RR is defined as the human impulse to rescue an identifiable person(s) facing an imminent threat who presents themselves to a health service. These discussions have focused largely on whether the RR passes tests imposed by traditional ethical theories, and general remarks about the scope

and force of the duty it may (or may not) impose. Yet there is little consensus between accounts of the RR and no sufficiently comprehensive or systematic description of it to focus and guide decision-making processes, either at the individual or policy level. My goal is to provide such an account for the medical community.

In my first chapter,, I discuss the concept of rescue broadly, offering my own formulation of the RR and highlighting the primary controversies that it faces—the identifiable victim bias and claims that rescue is discriminatory. Next, I proceed into psychology: I briefly describe current knowledge of the *psychological* impulse to rescue, and I attempt to determine if descriptive psychology can be useful in identifying when the impulse to rescue might constitute an ethical or moral imperative.

This empirical evidence provides a grounding for the following chapter, in which I consider the ethical status of the RR. Here, I first examine various common-morality accounts of duties to rescue, including 1) a general duty to rescue for all people, 2) a medical professional duty to rescue (although this may better be called a “duty to care” due to medical professional’s financial obligations to care as opposed to solely moral obligations), and 3) an institutional duty to rescue. I conclude that the general duty to rescue is too vague for proper analysis and the medical professional duty to rescue is relatively clear and so does not require analysis, leaving the institutional setting as the focus of the remainder of the treatment. Next, I place myself in the position of what I take to be the prototypical institutional actor. That is, I consider welfare optimization (i.e. consequentialism) of the institution’s target population as the first basis of action and then consider “moral externalities” — deontological considerations like equity and justice — as side constraints. Through this analysis, I show that the RR can be (reasonably) consistent with the major ethical theories of consequentialism and deontology, and in doing so,

realize a more tightly defined and nuanced formulation. However, some objections raise lasting questions. They indicate areas of reasonable disagreement about whether the RR should hold, or to what extent it should, in certain circumstances.

In the third and final chapter, I remain largely agnostic both about the status and relative weights of these moral complaints that have staying power. Instead, I catalogue them in a checklist of questions for policymakers to consider for themselves within the specific context of their situation. Further, I argue that in order to ensure equity and moral sensitivity to the social value of rescue, deliberations on contentious issues of rescue should engage a variety of community members and follow the requirements of procedural justice (Daniels and Sabin, 1997); Daniels 2000). Lastly, I provide various examples of my formulation—and the checklist of guiding questions—in action by using historical case studies. In these case studies, we can place the RR where it invariably must dwell—in real contexts—to gain insight into how the checklist might be used, to nuance it with additional practical considerations, and to test its conclusions.

My method allows me to engage with the difficult and controversial issues around the RR, instead of casting my lot on one end of a theoretical issue or another and then stopping short of a complete treatment. Many theorists have decided that particular ethical objections to the RR are insoluble, and so have simply despaired of there being much to say about it (Cookson et. al, 2008; Rulli & Millum, 2016). In doing so, theorists have halted their investigations prior to offering a complete treatment. My method also avoids the twin errors common in discussions of the RR: either to focus solely on theorizing, and so neglect the significance of psychological impulses to rescue, or to focus too heavily on describing human impulses and fail to properly respect theoretical considerations like efficiency or fairness. I will add my own positions on

these complex issues (particularly through a novel argument about negating symbolic value concerns, and a novel rebuttal of the general discrimination argument against the RR), but my main additions to the conversation around the RR are the comprehensiveness of my account, the practicality of adopting the perspective of institutional actors, and my willingness to build uncertainty into my final answer.³

Ultimately, I conclude that if there is sufficient social value to following a rescue impulse, and negative cost-benefit calculus or discriminatory concerns are relatively low, then it is morally appropriate to follow that impulse. However, there is no silver bullet for precisely determining these occasions, but rather contextual considerations best accounted for in a procedurally fair deliberative process. Accordingly, I synthesize the relevant considerations for which no definite moral answer is apparent, and provide them to decision-makers for use in deriving final, context-dependent conclusions.

2. Describing Rescue, the Rule of Rescue, and Controversies around the Rule of Rescue

Introductory remarks are in order. Here, I outline rescue, the Rule of Rescue, and the source of controversy around the Rule of Rescue. These definitions and delineations will be leveraged extensively in the later theoretical discussion.

Defining Individual Rescue and Institutional Rescue

Rescue is a specific category of the general case of actions that constitute helping another person. As I will use the term, it requires several contextual features: 1) that there is a very high, imminent risk to the person being rescued, 2) that the rescuer is helping a stranger rather than a

³ My treatment is similar to Daniels 2012 article in regard to accepting uncertainty, but Daniels stops short of clearly identifying the points of moral uncertainty that merit discussion (Daniels, 2012).

family member or friend,⁴ and 3) that the person is identifiable rather than seen as a statistical person (this delineation between identifiable and statistical will be explained in greater detail in the following section). By “high risk” I mean that life or permanent injury are in danger of occurring and will come about certainly or near-certainly without intervention. By “imminent” I mean that the risk will be realized speedily; the exact cutoff is unimportant, but perhaps less than a few hours or few days.

Consider a passerby who helps a child drowning in a well. This is an act of rescue, for the child 1) faced imminent, time sensitive peril 2) is rescued by a stranger 3) is identified, rather than statistical. Further, a child who receives a flu vaccine, and otherwise would have contracted a lethal case of the flu, is not rescued by my definition. This is because conditions 1) and 3) are not satisfied. The child was not in imminent danger when they received the flu vaccine, and, while the nurse administering the vaccine certainly identified the child, the attempt to help the child via the vaccine was an attempt to help a *statistical* child, rather than an attempt to help that child in particular. Since the child only had a statistical risk of contracting the flu, rather than a risk that was currently being realized as in the case of the drowning child, they would be considered a statistical victim.

Beyond general individual rescue, the primary focus of this paper will ultimately center on institutional rescue. An institution can promote rescue in two ways. First, the institution can rescuing a specific individual. For example, a child with extreme life-threatening traumatic injuries could be referred from a community hospital to a level 1 trauma center for

⁴ I include this condition in order to disentangle notions of a duty to rescue from notions of duty to family members, friends, or others who occupy special relationships of duty to the would-be rescuer. Notice, however, that this condition violates common-language use of “rescue”: we would still say that a father rescues his daughter from drowning in a well, even if my condition 2) here would not deem that to be rescue.

extracorporeal membrane oxygenation (in which the child's heart and lungs are bypassed via mechanical oxygenation and pumping of blood) and mass blood transfusion, with very low chance of survival. The Chief Medical Officer of the level 1 trauma center, which might or might not ordinarily not grant such intensive use of its resources for a single patient with a low chance of survival, could hear about the case and decide to accept the patient. This would constitute the institution rescuing the individual.

Second, the institution could engage in what I will call "systematized rescue." In systematized rescue, the institution establishes a policy that promotes rescue in some general sense. For example, the Chief Medical Officer and hospital decision-makers could broadly decide that all children with such and such extensive injuries will be accepted for transfer.⁵ Both institutional systematized rescue and institutional individual rescue become controversial when they either violate cost-benefit calculus or appear discriminatory. These sources of controversy will be discussed in depth throughout this paper.

Defining the Rule of Rescue

Ever since A.R. Jonsen first coined the phrase "Rule of Rescue"⁶ in 1986, there have been various competing formulations of the rule (Jonsen, 1986). These primarily differ in: 1) ethical assumptions, or whether to define the impulse to rescue as an ethical imperative, a psychological impulse, or an apparent duty; 2) context relevance, or whether to explicitly tie the

⁵ It is worth noting that the Emergency Medical Treatment and Active Labor Act requires that a higher-level hospital accept transfers from lower level hospitals, so long as there is space in the receiving facility and the transferring physician determines that it is in the best interest of the patient to be transferred. In practice, however, transfer decisions are often the result of a conversation between the transferring physician and the receiving physician/facility, so general institutional recommendations on what types of patients should be accepted for transfer would likely impact transfer decisions.

⁶ Jonsen intentionally picked the word "rule" to describe the sort of psychological impulse the human instinct to rescue is; it is somehow binding, more forceful than a typical impulse. Further, the Rule of Rescue is commonly seen as a moral rule, as constituting a duty.

rule to the medical context; and 3) stated implications, or whether to explicitly link the RR to the problem of the identified victim bias (to be discussed shortly). Take, for example, Bochner's very broad definition of the RR: "a perceived duty to save endangered life where possible" (Bochner et al., 1994, p. 901). Bochner leaves it an open question whether the RR is an ethical or merely psychological imperative, leaves the context relevance broad, and does not tie the RR to any implications. Compare Bochner's formulation to Osborne and Evans' far more specific definition: "the powerful human proclivity to rescue a single identified endangered life, regardless of cost, at the expense of any nameless faces who will therefore be denied health care" (Osborne and Evans, 1994, p.779). Here we have a definition that is again agnostic about the ethicality of the RR but has a specific context relevance (healthcare setting) and stated implications (discrimination against statistical lives).

For my purposes, I will use the following definition of the RR: the human impulse to rescue an identifiable person facing an imminent threat who presents themselves to a health service.⁷ I will call this formulation the healthcare Rule of Rescue (hRR). This definition is most amenable to my purposes because it makes no ethical assumptions nor assumptions about implications, but specifies a clear context: an identifiable person in need of urgent assistance from a health service. Essentially, I have selected the broadest formulation of the RR with respect to the medical establishment—although it should be noted that much specificity is still lacking, particularly in defining "an imminent threat" as well "identifiable". It is this version of

⁷ My definition here was influenced by a definition proposed by Nord et al. in 1995, in which the rule of rescue was considered "the sense of immediate duty that people feel towards those who present themselves to a health service with a serious condition"(pp. 90). My primary change from Nord's explication is the omission of the term "duty" to clearly keep any language of ethics or morality clear of the starting definition of the RR.

the RR, then, that will be the focus of this ethical treatment, and substantiating its currently vague implications will be the end goal.

Controversies Around the Rule of Rescue: The Identifiable Victim Bias and Discrimination

The Rule of Rescue was first articulated by Jonsen as a psychological fact, as a description of how people generally operate in situations that token a rescue response. People, Jonsen argued and psychologists have demonstrated, have a predisposition to help *identifiable* victims in immediate peril. The question for the moral theorist is to determine when this instinctual response is moral, and when it is immoral. The primary battle grounds for this debate for the Rule of Rescue center around what the “the identifiable victim bias” or, alternately, “the problem of statistical versus identifiable lives”, and on arguments that claim that rescue impulses are inherently discriminatory. The discussion of the identifiable victim bias is primarily directed towards the consequentialist, while the arguments on discrimination are aimed at the deontologist.

The Identifiable Victim Bias and Consequentialism

The identifiable victim bias⁸ is the bias to prefer to save identifiable lives over an equal or greater number of statistical lives. That is, the identifiable victim bias describes an apparently counter-consequentialist predisposition that arises from the hRR: most people instinctually allot greater resources and attention to individuals in immediate, identifiable peril than to an equal number of statistical lives. The strong desire to rescue an identifiable individual can be in conflict with a more cost-effective choice that would optimize overall lives saved. As Norman Daniels describes it, the identifiable victim bias is “the psychological fact that people are more

⁸ While A.R. Jonsen was the first to use the phrase “rule of rescue” in 1986, the identifiable victim bias was first articulated by Thomas Schelling in his famous 1968 paper (Schelling, 1968). Many thinkers on the topic of the rule of rescue and the identifiable victim effect therefore trace the start of the conversation on health allocation rescue morality not to Jonsen, but to Schelling.

likely to rescue an identified person we know is being harmed than a ‘statistical’ victim: the person who might be harmed if we don’t take steps to prevent it” (Daniels, 2012, p. 35). Take note that, according to definition I have laid out, the identifiable victim bias always entails a situation in which there is an equal or greater number of statistical lives that would have been saved without the bias for identified lives.

By identified person, I mean that they are seen as an individual rather than as one person among many similarly situated people (i.e. a statistic). This does not require that the identified person’s name is known, but just that, in the case in question, they are being identified as an individual in a specific danger rather than as one of many facing a probabilistic risk. Most thinkers illustrate the distinction between identified and statistical lives through examples:

By an “identified victim,” I mean [a woman], lying injured in the passenger seat of the wrecked automobile on the corner of Main Street and Broadway...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) (Daniels 2012 p. 35).

The identified victim bias is by far the most frequently cited complaint against the hRR (Daniels, 2012; McKie, 2003; Rulli and Millum, 2016), and for good reason. Critics of the hRR argue that the human bias towards rescue leads the healthcare system to direct disproportionate resources towards the identified victims, at the cost of overall worse outcomes for a greater number of statistical victims. Critics make this claim at both the individual level (i.e. physicians making poor allocation decisions among their patients) but are also frequently making a system-level claim about established health policy. In the sense of “systematized rescue” discussed previously, the hRR could be interpreted as the mandate to design systems that are adept at rescuing individuals once they have become identifiable (i.e. after they are personally

endangered), rather than designing systems to stop statistical individuals from ever becoming identifiably endangered. To draw from Daniels' example of a motor vehicle crash above, critics might say that we prioritize developing systems skilled at saving the most recent motor vehicle crash victims (like our EMS system), rather than funding interventions to improve driving safety to save potentially greater number of individuals who have not yet crashed (like driving safety classes—disregarding controversies around the efficacy of those classes). Thus, at the institutional level, the identifiable victim bias leads to investment in treatment services instead of in public health.

The identifiable victim bias is not identical with the hRR, because acting on the hRR does not always entail acting on a problematic bias in favor of the identified over the statistical—there may be times where the option that leads to the most net lives saved is to act on the hRR, like with a choking child that could very easily be rescued.⁹ Instances of easy rescue are largely uncontroversial, because there is no tradeoff involved. In them, there is clearly a duty to act on the hRR.

Rather, the identified victim bias¹⁰ is the major theoretical bound on the hRR from a consequentialist perspective. So long as rescue has clear benefits and does not lead to greater harms imposed on statistical victims, it seems clear that we should rescue (holding concerns about racial discrimination or justice equal). But there are many situations where the cost of

⁹ The RR is also often conflated with additional topics. Two of the most significant examples are 1) the “treatment bias”—that healthcare systems prioritize treatment over prevention—and 2) the notion that the RR is really a derivative of the moral imperative to help the worst off. Both characterizations of the RR are flawed. It is true that the RR can lead to prioritization of treatment rather than preventive services, but the RR is broader than this—the RR is also in action in cases where rescue is the right choice over preventive care. The RR is also separate from a Rawlsian imperative to help the worst off (the “difference principle”). The RR applies only to identified victims, whereas the moral imperative to help the worst off applies to anonymous individuals as well. Indeed, Rawls' underlying theory for this imperative relies on a rational egoist behind a veil of ignorance. This setup precludes identifiability as a consideration of the difference principle (McKie and Richardson, 2003).

¹⁰ From this point forward, I will always be taking the identifiable victim bias to mean the bias to prefer identifiable over statistical lives *in the case of equal to or greater expected statistical lives lost*.

rescue may come at the expense of greater downstream life-saving. It is in these situations where the theorist, physician, or policy-maker must pause and inquire whether they should act on or control their impulse to follow the hRR. They must ask: should we act on the hRR to save an individual at the expense of a greater number of statistical lives? What is the acceptable trade-off between identifiable and statistical lives? Is it 1:1, 1:2, 1:3? What rationale could be provided for such a decision, and where are its bounds? These questions will be the focus of the later consequentialist section.

Discrimination and Deontology

Beyond the consequentialist aspects of the identifiable victim bias, there is also another thread of attack commonly levied at the bias. This attack is not concerned about overall utility; rather, it focuses on fairness and justice. By far the most influential deontological¹¹ critique—that is, a critique based on the viewpoint that acts are moral or immoral according to rules, motives, or duties rather than due to consequences—takes the form of the discrimination argument.

The discrimination argument, in turn, has two forms. First, and far broader in its implications, is the version of the discrimination argument that holds that it is discriminatory—unjust, unfair—to act on the identifiable victim bias. In other words, that it is unfair to favor identified over statistical lives (Daniels, 2012). One of the most common deontological duties is to equally respect people due to their status as persons. The thrust of this first discrimination argument is to cast the identifiable victim bias as a case of unequally respecting human lives—in particular, of favoring the identified over the statistical. If this argument were a slogan, it would

¹¹ Throughout this paper, I use the term “deontological” rather broadly. I am not tying it to any particular deontological theory, but rather referencing any choice-based ethical theory that incorporates more than consequences into its ethical framework. It is therefore suitable to replace “deontology” throughout this paper with “non-consequentialist”.

read “statistical persons are persons too”. Accordingly, the deontological critic of the hRR would say that any rescue that entails acting on the identifiable victim bias constitutes immoral discrimination (note, however, that easy rescue is still permissible and even required). Equal chances of being saved should be given to identified versus statistical victims so as to avoid unfair, unjust discrimination. This argument would yield a knockdown of the hRR in any case where the identifiable victim bias leads to a preference for the same or greater number of identifiable lives compared to statistical lives. In the later deontological section we will return to this argument, and I will aim to refute it.

The second form of the discrimination bias argues that many instances of following a rescue bias leads to discrimination against certain sub-populations, like women, racial or religious minorities, or geographic regions. I will argue in the deontological section that these considerations are valid, but that they do not constitute an in-principle issue with the hRR or the identifiable victim bias. Rather, they are contextual considerations.

However, before I embark on theoretical arguments on either consequentialism or deontology, it is essential to provide what empirical context I can. It is necessary to develop some empirical understanding of what the “human impulse” underlying the RR really is. To do this, I will next explore the psychological literature on the instinct to rescue identifiable individuals. I will describe what is known of this impulse, and I will subject the findings to two general questions: do our psychological tendencies point to some intrinsically moral (or amoral) underlying value? Separately, are our psychological desires in rescue so strong that they merit an important place in a consequentialist framework as a means for promoting well-being? This discussion will serve as the empirical framing for the following chapters, which deal with

theoretical arguments for and against the hRR when making institutional healthcare allocation decisions.

3. The Psychology of Rescue

Psychologists have been investigating the human predisposition to assist imperiled individuals for decades. While the psychology of rescue impulses is descriptive rather than normative, it is helpful to first know what the impulse *is* before one can know if that impulse should be indulged or rehabilitated. Fortunately for research, but unfortunately for this paper, the literature that is relevant to rescue psychology is voluminous and includes work in economics and marketing.^{12, 13} Accordingly, I will focus on describing what I take to be the key findings in psychology around the RR. I will also offer brief thoughts on to what extent these psychological findings do or do not have bearing on whether the RR is ethically sound.

Overview of the Psychology of Rescue

Jenni and Loewenstein's seminal study in 1997 was the first major breakthrough in studying the RR. Jenni and Loewenstein summarized two previous studies in which research participants read descriptions of various scenarios and rated the importance of reducing harm to the victims. The scenarios were modified to track each of four potential causal factors. These four factors had been commonly cited explanations for the identifiable victim bias up to that time. First, they tested for situations in which the person in need of rescue is vividly identified, as

¹² For example, there is considerable marketing research on the economic impact of portraying identifiable victims in charity campaigns. Consider another piece of marketing literature: "spokesperson-product congruence", or research on identifying the features of a spokesperson that are most important for eliciting trust, sympathy, and empathy in the audience (Kamins and Gupta, 1994). Spokesperson marketing generally does not depict a rescue situation – but it does in some cases, and it certainly has bearing on rescue psychology.

¹³ I am indebted to Deborah Small's attempts to summarize psychological research around rescue impulses, although her focus is explicitly on the identifiable victim bias. For her discussion of the psychology of identifiable versus statistical lives, see her excellent book chapter: Small, D.A. (2015). On the Psychology of the Identifiable Victim Effect. *Identified Versus Statistical Lives: An Interdisciplinary Perspective*. p. 13-24. Oxford University Press.

defined by the emotional salience and the time-sensitivity of the situation. For example, a child drowning would evoke a powerful emotional response and the child's condition would appear highly time-sensitive to would-be rescuers. Second, the certainty of the threat to the identified victim versus the probability of the threat to the statistical victims. Our inclinations bend towards preventing certain losses rather than probabilistic losses, a behavior called "risk seeking with regard to losses" (Kahneman & Tversky, 1979; *Ibid.*, 2000). Third, the distinction between evaluating harm before it occurs versus after. Humans have a rather narrow time horizon, and it is well-known that humans often struggle to delay gratification—the RR so understood would be a sort of impatience in wanting to help sooner rather than later. Fourth, the proportion of the reference group that can be saved. If validated, this would mean something like that people perceive the deaths of 10 individuals out a village of 100 as a greater loss than 11 individuals in a village of 1,000.¹⁴ An identifiable victim is considered the perennial case of a need for concentrated help: they are a reference group of one.

The results of Jenni and Loewenstein's paper were highly counterintuitive: they found no evidence that the first 3 commonly cited features of rescue impulses play a role in our compulsion to rescue. The only significant effect they found was with respect to the proportion of the reference group that could be saved. The greater the proportion of the reference group that could be saved, the greater the bias was towards helping that group. Their research indicated that sympathy—as tracked by vividly and emotionally recounted stories—was not responsible (at least in a strong sense) for the RR; rather, a quirk of human quantitative thinking was.

¹⁴ I offer the 10/100 and 11/1000 case as an example, but it need not be factually accurate for the bias to hold. There could be a bias for a larger proportion of the reference group saved that discriminates only based on more drastically different cases (for example, it could be a preference to save 10/100 rather than 30/30,000,000) than the example I provide.

Further research on the RR over the last 2 decades has shown that the reference group effect is not sufficient in explaining our preference for identified lives—and sympathy made a reappearance. The 1997 experiment focused on vividness (i.e. beefing up the emotional appeal of a story) to see if sympathy played a causal role in the RR. Later studies found that it is not vividness that causes a sympathetic response that elicits the identified victim effect, but rather the simple act of being identified. Small and Loewenstein (2003) found that any form of identification leads to what has subsequently been called “the identification effect”—even just assigning a victim a number increases their emotional appeal, and can double giving. In a follow-up study, people who were asked to donate to Habitat for Humanity to help a family get a new home gave much more if they were told that a family *had* been selected rather than *would* be selected. Giving to an identified family (a concrete, present family) was much greater than to an unidentified family (a family that was abstract, in the future).

Additional studies have also identified what has been termed the “singularity effect” (Kogut and Rotuv, 2005a). Kogut and Rotuv showed that a single identifiable victim gets more support than a small group of identifiable victims. Another major finding of Kogut and Rotuv was when they showed that, when presented with the choice between assisting a single victim versus a group in need, people make the utilitarian decision to help the group—but they donate less than people who were only shown the option of helping the single victim (Kogut and Rotuv, 2005b)! These seemingly contradictory results indicate that sympathy—and therefore generosity—unconsciously favors single individuals, while consciously made allocation decisions favor groups because people attempt to optimize social utility.

There are two final principles that have been offered to explain the identifiable victim bias. Psychological distance refers to perceived closeness (i.e. shared traits like nationality,

ethnicity, age, etc.), and is likely related to evolutionary psychology—the general idea is that evolution would select for traits that lead us to help our family members (thereby promoting the spread of our genes), and our desire to help generalizes to those we perceive as similar to us. Considerable evidence substantiates the notion that humans allocate resources and generosity based on the perceived closeness of others—that is, we act on in-group biases (Rachlin and Jones 2006; Ibid., 2008).¹⁵

Finally, a finding on rescue psychology diverges a bit from the above considerations: research into why there is such an outsized show of support in the case of sudden disasters (like a hurricane or earthquake) as compared to ongoing, but just as easily improved disasters (like the burden from neglected tropical diseases). Findings here support the concept of “reference dependent sympathy”, which holds that people are most inclined to help in situations where a negative change from the baseline occurs (Small, 2010). For example, we are more likely to feel an impulse to assist a newly impoverished person than a long-time poor person.

A possible explanation for all of these effects relies on “dual-process theory,” developed by Kahneman and Frederick (2002). They hold that much of human decision-making happens along two different systems: in system one, the implicit/unconscious system, decisions are driven by emotions and instincts, leading to rapid action and strong motivations. In system two, the conscious/deliberative system, decisions are driven by rational considerations and lead to slow, weak motivations. System two can overrule system one, but only if the person is inclined to take the time to carefully consider an issue. In rushed decision-making (which constitutes the vast

¹⁵ For example, consider the following evidence: Individual allocation decisions favor members of the same group—even if the group is entirely made up as part of the experiment! (Dovidio et al. 1997). Similarly, empathic priming leads to more generous allocation decisions (Stotland and Dunn 1963; Batson, Early, and Salvarani 1997). An intuitive aspect of psychological distance is that we are predisposed to help people who have the same medical condition as someone personally close to us, like a family member (Small and Simonsohn 2008).

majority of our decisions) system two never comes online, and system one dominates.¹⁶

Numerous empirical studies have provided evidence in support of dual-process theory as an explanation of rescue impulses. Small, Loewenstein, and Slovic (2007) found that presenting victim statistics alongside charity materials for identifiable victims decreases donations. As Small says, “a picture and a statistic are worse than a picture” (Small, 2015, p.17). Further, priming participants to think in a statistical manner decreased donations to identifiable victims, while priming participants with emotional materials failed to increase giving to statistical victims. Essentially, thinking analytically overwhelms thinking emotively with respect to giving, exactly as dual process theory predicts. The paradox is that priming healthcare stakeholders to think analytically could increase the reasonableness of their decisions per System two’s superior rationality, while decreasing the overall generosity of their behavior per System two’s weaker motivational powers (Ibid., 2015).

To briefly summarize: research has highlighted the importance of reference group effects, identification, the singularity effect, psychological distance, and reference dependent sympathy as pieces of the explanatory puzzle around the bias to help identifiable victims. These notions are generally well-accommodated by a dual-process theory that says that unconscious and immediate emotional responses provide strong motivation to help and that more deliberate, rational processes can overrule the emotional response but lead to less motivation to help.

¹⁶ Interestingly, this delineation of unconscious and forceful thinking versus rational but less motivational thinking has a long history in philosophy—extending from Plato’s tripartite soul to Hare’s theory of intuitive versus critical moral thinking to Kagan’s theory of vivid/cold ideas.

Do the Psychological Impulses have Normative Force?

We can now ask the question: do any of these psychological mechanisms appear to have normative force in support of the hRR? The answer appears to be no, and worse a few of the psychological mechanisms appear explicitly *immoral*. First, the reference group effect does not seem a particularly moral consideration — it does not track concerns frequently considered moral, like efficacy, fairness, sympathy, or empathy. Rather, it seems predicated on quirks in human statistical thinking. Second, the identification effect begs the question as to whether it is moral, since it is what is under investigation. Moreover, the fact that to identify a victim by a number versus leaving them unidentified has such a profound effect seems indicative of a weakness of moral imagination, rather than tracking typically moral considerations like justice or need.¹⁷

Third, making allocation decisions based on psychological distance is discriminatory in the context of healthcare. There have been strong arguments that certain “agent-relative prerogatives” like physical proximity and special relationships may be morally noteworthy (Kamm, 2000), but these concerns are not admissible in medical decision-making situations. Indeed, considering psychological distance or agent-relative prerogatives in healthcare allocation would be deeply immoral and unjust. For example, it would be profoundly unethical for a leader within the United Network for Organ Sharing (UNOS) to arbitrarily move their mother up the kidney transplant list due to agent-relative prerogatives, and it would be depraved to decide that

¹⁷ Peter Railton presents an interesting counter-argument—he argues that, as social animals, we have learned that once something is identified it is much more likely to occur or be realized, and that those things that are non-identified are not likely to occur and so should be discounted. For example, home owners undergoing a home renovation frequently disbelieve the dates and times by which contractors “guarantee” projects will be completed, and for good reason! Thus, Railton says our bias for *identified* entities (be it people or projects) is perfectly sensible, but not well suited to healthcare rationing where statistics are far more precise than our commonly encountered social promises. Railton therefore believes that much of the bias in favor of the RR, as it were, can be trained out of individuals by exposure to rationing situations (Railton, 2015).

Southerners lives should be discounted because they lie at a further psychological distance from the Northern-majority leaders of UNOS.¹⁸ Fourth, the singularity effect also seems discriminatory and arbitrary with regard to identified lives, particularly given evidence that presenting statistical lives as unified plays on the same singularity effect and increases donations (Smith, Faro, and Burson, 2011). Here, we see a case in which diminishing the amount of need—all other features of being identifiable held constant—increases sympathy. Reference dependent sympathy is an interesting case—perhaps a sudden loss is worse than a continuation of a long-time problem—but is primarily relevant to large scale social context changes, like a disaster scenario, rather than to healthcare allocation decisions within a set context.

Finally, dual-process theory points to an apparent catch 22: we have a tendency to misinterpret need in various situations due to our having a strong emotional/irrational response, but using our rational faculties makes us less charitable toward those in desperate need. It seems *prima facie* that the best situation would be to have an emotional/motivational response in proportion to need, but this is empirically not what people do.

All of this is not to say that these psychological mechanisms couldn't be instrumentally useful for increasing giving and therefore warranted. Dual-process theory above is a good example of how the RR may be instrumentally useful by increasing the motivation to provide assistance to those in need. Rather, none of the psychological factors appear to be intrinsically moral. Indeed, many seem inherently unfair or discriminatory.

¹⁸ The tension between agent-relative prerogatives and institutional duties are nicely depicted by the ancient Confucian scholar, Mencius. Confucians are renowned for emphasizing filial piety (a form of agent-relative prerogative) as a primary virtue. In one account, Mencius remarks that if the father of Shun, the legendary Sage, committed murder then Shun as his son should aid his father in escape—but only after abdicating his throne, so that he can be a son without any conflict with his duty as just ruler. Here, even the filially pious Mencius acknowledges that individuals leading institutions accept duties that may go against filial piety, and that in those cases they cannot wield the authority of their position with a filial bias (*Mencius*, 1983).

Summary of the Psychology of Rescue and the Possibility of Modifying Rescue Impulses

After a trek through such a varied empirical environment, summary remarks are certainly in order. We have seen that the impulse to rescue identifiable lives is actually a cluster of interrelated psychological tendencies: reference group effects, identifiability, singularity, psychological distancing, and reference dependent sympathy. Explanations of these causal factors generally utilize dual-process theory, in which system one thinking is emotional and instinctual but highly motivational, while system two thinking is rational and can override system one, but is only weakly motivational. This leads to a paradox in which rational decision-making to maximize impact seems to lead to necessarily weaker motivations than emotional decision-making evoked by a singular, identified person.

Interestingly, some researchers have discussed potential methods for bypassing this paradox. First, the singularity effect can be leveraged to level the instinctual playing field between identifiable and statistical victims. Smith, Faro and Burson (2011) have shown that portraying statistical victims as a coherent unit causes the singularity effect to occur for the statistical victims, effectively cancelling out the singularity bias towards identified individuals. A prominent moral philosopher, Peter Railton, believes that system one thinking is largely influenced by experiential social training that teaches us to prioritize the here and now. He therefore believes we can retrain non-ideal system one behaviors by providing lay-people (and, more realistically, healthcare decision makers) opportunities to experience simulated allocation scenarios (Railton, 2015). This could affect a greater number of the psychological effects, including the reference group, identifiability, and psychological distancing factors.

In sum, we have gained considerably from reviewing what is currently known about the psychology of rescue impulses. Psychological analysis has simplified and reduced the vague RR

impulse, and the identifiable victim bias, into a series of separable impulses that can be more easily considered. This reduction allowed us to see that none of the impulses seem likely to be tracking considerations of moral substance in the context of applying the RR to the context of healthcare. Further, the possibility of manipulating the known psychological factors in order to minimize rescue concerns will be of interest in consequentialist theoretical considerations.

However, note the substantial limitations of current research. Psychologists have not performed studies of large-scale medical trust effects, or of effects on community solidarity and well-being, due to system-level healthcare allocation decisions. Rather, research has focused on how the individual views a specific rescue situation. This perspective is useful when considering whether specific cases of a bias to rescue have intrinsic moral worth. However, they do not have much to say by way of the societal well-being and positive public opinion generated by institutional rescue allocation decisions. In the following section I will argue that this case, of healthcare institutions making allocation decisions, is the most relevant and important to consider with respect to the hRR. So, until empirical research can find some footing on the social value of institutional allocation decision-making, we will have to consider its value as an open question while relying on the anecdotes and evidence that exists. These societal/communal considerations will be at the crux of theoretical arguments about the Rule.

Chapter Two: Considerations from Ethical Theory

The following discussion will detail theoretical issues with applying the hRR as a moral dictum. First, I will offer preliminary remarks on the “shape” of the obligation the hRR could impose by outlining the various actors (the general public, healthcare providers, and institutions) who could be thought to have a duty to follow the hRR. Here, I will make the case that the relevant actor to discuss as potentially bearing a duty to rescue is the healthcare institution, rather than healthcare providers or lay-people. In the later two sections I will test the hRR (in particular, the identifiable victim bias) by both consequentialist and deontological critiques. I will adopt what I take to be the default perspective of an institutional actor: relying on consequentialism primarily, while incorporating deontological side constraints as a check on consequentialist thinking. I will conclude that the hRR and in particular the identifiable victim bias can be thought to reasonably satisfy theoretical complaints, but that the form of the rule that satisfies these complaints is complex and is best followed by discussing it by committee on a case-by-case basis.

1. To Whom Might the Rule of Rescue Impose a Duty?

As previously noted, the rule of rescue describes the impulse most people have to rescue others, when possible. The descriptive feature of the hRR, that most people have this impulse, is not up for debate. Rather, our concern is whether this psychological fact carries moral weight. If it does, then it would impose certain duties to rescue on the healthcare system. Apparent duties to rescue have been extensively discussed in philosophical and policy literature. The discussion has largely focused on differential obligations due to varied societal positions – for example, the differences between the duties to rescue that we all owe each other, that governments owe to citizens, or that specific groups of professionals (like doctors or firefighters) owe to their

populations served. Here, I wish to provide a (very brief) discussion of this literature in order to identify the most relevant and tractable audience for hRR considerations.

The two valences of rescue obligations are 1) scope, answering the question: to whom do I have a duty to rescue? and 2) force, answering the question: how far must I go in acting on my duty to rescue? Using the clarity of scope and force obligations as a test, I will quickly outline best attempts at a general duty to rescue (also called the duty of easy rescue), the medical professional duty to rescue, and the healthcare institution duty to rescue. I will conclude that the institutional duty to rescue is where the bulk of the controversy around the hRR lies in healthcare, allowing me to focus on the hRR as it applies to institutions for the bulk of the later theoretical sections.

The General Duty to Rescue

The most generally applicable and broad account of the scope and force of rescue is given by the “duty to help”, which says that one has a moral duty to prevent serious harm to another person so long as the personal cost is minimal. The duty to help is normally held to apply to all moral agents—that is, to every adult person (Rulli and Millum, 2016). However, the duty to help faces a number of serious issues, such that it is nearly impossible to provide a consensus (or even pluralist) descriptive account of the obligations it imposes. First, the force of the duty is ambiguous: how can one define minimal cost? Peter Singer disagrees with the notion of minimal cost altogether, and famously argues that the obligation to rescue should continue until the rescuer is facing an equal tradeoff with the rescued (Singer, 1972). The duty to help is also likely perceived quite differently by different cultures, further muddying any clear description of its force.

The scope of the duty to help is also very vague: it is unclear if the duty to prevent serious harm applies to everyone, regardless of the distance between rescuer and the person in need, and regardless of the number of similarly situated rescuers and victims. For example, consider a child in a developing country facing imminent starvation, who could be saved with a donation of \$40 (Save the Children Fund, Accessed December 2018).¹⁹ This is a cost that would be minimal for most middle-class Americans, but should any middle-class American who does not opt to pay it be considered morally reprehensible? There is little agreement on this question. Indeed, some thinkers, like Frances Kamm, have developed theories to demonstrate the moral significance of distance by arguing for the existence of agent-relative prerogatives (the common-morality notion that individuals have special responsibilities for what happens near them rather than far from them) (Kamm, 2000).

There are additional situations for which it is even more difficult to provide a satisfactory account for the scope of the duty of easy rescue. Rulli and Millum give the example of a middle-class person in Calcutta passing numerous starving beggar children on a daily basis (Rulli & Millum, 2016). Typically, we would not say that the middle-class person has an obligation to provide food for the individuals they pass. But distance isn't the issue here; the issue is something more like the saliency of the victims' need. Intuitively, we seem to act as though extreme need that is "normal" for an area no longer tokens the rescue duty. If we accept this, then we accept that the duty of easy rescue is enormously dependent on the cultural and historical considerations about the location in question. It appears that the duty to help—at the

¹⁹ This is a rather ironic example, because these charities utilize a version of the RR to solicit donations. They put a face and a name on their fundraising appeals, personalizing the desperation of the people for whom they are acting. Their fundraising would not be nearly as effective if they just sent out a generic "ask" to help "the starving children in [insert country]".

cost of brevity and some intuitive appeal—would be better called “the duty to help those in exceptionally bad circumstances relative to their context.”

There are even more complicated considerations in the Calcutta thought experiment. If we do accept that we should provide money to a starving beggar, then we also need to determine how much to give. Aside from the question of how much to give, what should we do for her 15 co-beggars on the same street, and the 30 on the next street? Should we give until we have no more (or just enough to feed yourself a meal) as Singer might suggest? Further, if we had \$50 that we were planning to donate, would it be best to give it all to one person and raise their standard of living considerably, or give 50 people \$1 each? Or 25 people \$2 each? These issues are not easily resolved and are at the crux of ongoing scholarly debate.

Given the difficulties in providing a satisfactory account of the scope and force of the duty of easy rescue, we should prefer other options for providing an account of the RR as it relates to the medical establishment. Fortunately, there are more specific versions of a duty to rescue that are of greater relevance to medical institutions. These alternatives are the medical professional duty to rescue, and the healthcare institutional duty to rescue.

The Medical Professional Duty to Rescue

Both the medical professional and healthcare institutional roles are person-dependent features of a rescue situation. Whereas the duty to easy rescue applies to any person who finds themselves in a position to assist in a rescue (i.e. person-independent) both the professional and institutional duty applies to only specific groups of people (individuals in certain professional roles and in certain institutional decision-making roles, respectively).

First, consider the medical professional duty to rescue. This applies to physicians, nurses, paramedics, EMTs, etc. One can reasonably ask if a medical provider assisting an imperiled

patient can be counted as a “rescue” in the morally admirable sense, given that the professional is being paid to provide the assistance, and given that they willingly and knowingly enter into a profession with well specified duties and ethical codes. Further, the weight of their duty is increased by considerations of public trust. For a professional to act against the specified duties of their profession risks causing fear and mistrust among the public towards the profession in general, a highly undesirable outcome—especially in healthcare, where trust is so essential. These considerations are important in differentiating solely moral duties from financial, legal, or professional duties. However, for my purposes I am interested in what *rescue* duties at large the physician has in terms of assisting imperiled and identifiable individuals (as outlined in my definition of “rescue” in Section 2), regardless of whether those duties are acquired for solely moral or also financial reasons. So long as their overall rescue duties are clear, the physician is no longer a source of controversy in determining rescue decision making. My goal is to argue for exactly that conclusion.

While “rescue” as a morally laden term may therefore not be appropriate for labeling the medical professional duty, still there is a clear answer to the force and scope of medical professional’s obligations in acting to help those in sudden peril. The general structure of their professional obligation to rescue hinges on the sanctity of the patient-medical provider relationship. As Phillip Rosoff notes, “the essence of the duty that physicians have to their patients is devoted to bettering the latter’s interest” (Rosoff, 2017, p. 137). Many other thinkers agree (Beach et. al, 2005; Honeybul et. al, 2011; Kohn et al., 2011). The unique patient-physician relationship includes both a professional obligation and a financial obligation that physicians do everything within their power to save the life of their patient or prevent future disability, so long as patient autonomy is being provided for. This includes favoring the patient

over other similarly situated individuals, given that they are not the physician's patients as well. As such, it is relatively uncontroversial to articulate rescue-like obligations as they apply to on-the-job physicians: they should act to rescue their patients, even at the cost of a greater number of others' lives. While working, the scope of their duty to rescue is limited to those with whom they have established a physician-patient relationship, and the force of their duty is to act to rescue their patients from all life (or limb) threatening medical emergencies.²⁰ They have no special obligations to rescue non-patients. While off the job, physicians are generally regarded as having an ethical (if not legal or financial) duty to provide what emergency medical assistance they can to an injured person (AMA Code of Ethics, 2016). So the scope and force of rescue duties for medical professionals seem (as far as such things can be) relatively clear.²¹

The Healthcare Institutional Duty to Rescue

A concern one could voice against physician's acting on their professional obligation to rescue is that to save *every* patient who requests treatment cannot be cost-effective given the limited resources available in healthcare. Consider an example: two physicians are arguing over who should get an operating room, one to save a 20-year-old suffering from a sudden hemorrhagic stroke, the other to save an 85-year-old in imminent need of a catheterization to

²⁰ There are a few notable exceptions, in which the physician has separate public health duties. For example, most physicians are mandatory reporters of elder abuse or of certain infectious diseases, like Ebola. In these cases, the physician's duty to public health interests are greater than their duties to patients (in this case, duties of confidentiality). Still, even here the duties gain their force because they are willingly agreed upon aspects of entering into the medical profession.

²¹ There are, of course, still many aspects of physician responsibilities that are controversial. For example, there have been ongoing debates about the "duty to care", or the responsibility of medical providers to provide care to anyone who presents themselves as a patient. This issue came to a head in the SARS epidemic in 2003, when a large number of nurses (and some doctors) refused to care for SARS patients out of fear for their own health. The SARS scare also led to conversations around whether the duty to care, if it exists, is equally strong for all providers—are the obligations of doctors greater than nurses whose obligations are greater than techs (Reid, 2005; Ruderman et. al, 2006)? All of this to say: there are many complicated issues around physician obligations that remain controversial. However, the duty of physicians to act in the best interest of their patients (a duty to rescue established patients) is widely acknowledged.

treat a heart attack (i.e. both facing imminent life threats with the same likelihood of recovery). Given the professional duty to rescue, both physicians should argue on behalf of their patients for the operating room. However, it is up to institutional policy to decide to whom the room should be assigned; in the case above, most would agree that the institution should elect for the 20-year-old, based on cost-benefit (in this case, life expectancy) and given that all else is held equal.

Further, there are differences in the ability of physicians to advocate for their patients that could be seen as unjust. For example, prior to 2002 the order of liver organ transplant lists was determined by physician arguments rather than by objective scores (Moylan et al., 2008). This led to racial and social class based inequities, since the more socially privileged generally had access to better resources and more persuasive physicians. Various health institutions, whether the United Network for Organ Sharing or a hospital, establish policies and arbitrate between providers to avoid these sorts of undesirable outcomes. This is where the institutional duty to rescue enters.

Institutional policy decisions may seem at first glance to be improper instances in which to consider the RR because policy decisions are made in board rooms, not in the face of imperiled, identifiable individuals. However, policy decisions must be made given the consideration of later rescues. Institutions must anticipate the choice situation of real providers, and the responses of their patients to institutional policies for or against rescue (Hadorn, 1991).

So the institution is responsible for handling allocation decisions that implicate the hRR, freeing the physicians to defend the public trust dependent on the special physician-patient relationship. Accordingly, we now need an account of the scope and force of the institutional duty to rescue. Fortunately, it is relatively easy to delineate the scope of the institutional responsibility. The scope is the specific population the institution is designed to serve. In the

case of the Ventura County Community Hospital, their duty to rescue extends to the population of Ventura County.²² For medical institutions with a public service mission (like a public health department or a not-for-profit hospital) their scope is their local population.²³

However, describing the force of the institutional duty is quite difficult. I will provide brief comments here; however, laying out in full considerations for the institutional setting will be the focus of the remainder of this treatment. It is uncontroversial that institutions should provide for rescue in situations where rescue is cost-effective and there are few (if any) tradeoffs. For example, if there are a surplus of surgeons/operating rooms and a patient presents with a life-threatening appendectomy, they should receive the surgery. The controversy arises when discussing in which (if any) situations a medical institution should favor policies that provide for the rescue of identifiable individuals over a greater number of other lives. Consider the famous case of Coby Howard. The Oregon Medicaid system had recently made the decision to move funds from an expensive bone marrow transplant program for children to expand prenatal maternal care, based on cost-benefit evidence of lives predicted to be saved. Coby was an eight-year-old with leukemia impacted by the decision, and he quickly became an identified individual via sympathetic national media coverage. Should the Medicaid system bow to the demands of the more easily identifiable patients (those with bone cancer needing clear treatment) at the expense of a more cost-effective, but less identifiable program (prenatal care to prevent downstream health issues for mothers and children)?

²² The hospital of course also has a duty to individuals passing through Ventura County—for example, the commuter driving through who has a car accident and is transported to the hospital. So a more nuanced description of the scope would be, “the population of patients that present to the Ventura County Community Hospital.”

²³ The scope of responsibility for a national pharmaceutical company, however, is more difficult to describe—in particular because obligations to stock-holders muddy a clear description of a service population. Therefore, my discussion here is most relevant to organizations with an explicit communal and altruistic focus.

The high-level answer I will argue in the following sections is that institutions should elect for cost-effective rescues in most cases while giving some deference to the public trust and value placed on rescuing identifiable individuals, and also make certain that decisions adhere reasonably well to standards of equity and fairness. This answer is in line with the general public's opinion on the hRR: a recent patient committee called to describe their perspective on the hRR as it applies to medical institutions said that the hRR should rarely guide health policy, except for in instances that threaten to denigrate the value of human life (Citizen's Council Report: the Rule of Rescue, NICE, 2006).

2. Putting the Rule of Rescue to the Test: Consequentialism and Deontology

Consider this situation. It is 2025, and the next flu pandemic is raging. You are a New York hospital administrator, and the intensive care unit staff have just informed you that all ventilators in the hospital are fully occupied. The line of patients needing ventilators is continuing to increase, with many patients at risk of dying prior to a spot opening for them. Worse, every hospital in the state is in the same situation, so none of your requests for additional ventilators have been met. This situation is likely to last at least a few more months, with hundreds or thousands of more flu victims who will die without ventilator support—far more than you can provide. It is up to you to determine the criteria for who receives ventilators. You have many options: you could decide on a first-come first-serve scheme, rationalizing the choice on the basis of autonomy and respect for persons. You could allocate ventilators based on the human interest appeal of various patients (for example, prioritizing all working adults over the homeless). Or, you could allocate based on the expected years of life saved — you could consider how likely a person is to survive their disease, and how long they are likely to live after that point.

Most likely, you will fall back on existing recommendations, like those issued by the New York State Guidelines Task Force Influenza Allocation in 2007 (New York State Guidelines, 2007). This plan recommended allocation of ventilators based solely on an objective scoring of medical survivability, with regular monitoring of patients for improvement or decompensation. “Medical survivability” means how likely a patient is to survive the illness, without regard for life expectancy after that point (unlike a pure consequentialist scheme). Most notably, if a patient fails to improve or further decompensates upon reevaluation they can be extubated – essentially, abandoned to die – in order to free a ventilator for a patient with a better likelihood of survival. The New York allocation plan therefore brings up questions not only of rescue, but also abandonment. The plan has no incorporation of a rule of rescue—which would dictate that you rescue each identifiable victim as they come, i.e. in a first-come-first serve scheme without considering the flu victims as statistical and so devising an allocation scheme. Instead, it utilizes survivability as the sole standard. The question that we shall ask: is it right to ignore the hRR? Does the rule of rescue have any bearing on how you, the hospital administrator, should set up your allocation scheme?

Here, I will explore this question by applying two traditional ethical frameworks: consequentialism and deontology. I will attempt to mimic default institutional decision-making by using consequentialism as the foundational theory that provides both positive commands and negative constraints, and I will consider deontology only as offering moral side constraints. In both ethical theories, I will argue that there both is and should be room for rescue considerations, but only minimally and within set bounds. Throughout the discussions, I shall come back to the example of a ventilator shortage in a flu pandemic to illustrate key points.

Consequentialist Considerations

Consequentialism is highly influential in healthcare institutions and policy making broadly. This is largely because its result is so simple, intuitive, and relatively implementable: consequentialism states that we ought to optimize wellbeing. In the context of healthcare, consequentialism generally reduces to allocating money or resources to whatever will lead to the most lives saved (considering both lives saved and overall life expectancy) and suffering averted. The argument to get to this conclusion is brief:

- 1) We should act so as to maximize overall wellbeing (the base tenet of consequentialism).
- 2) Life and health are among the greatest sources of well-being, and the sources of well-being with which healthcare is concerned, such that other sources of well-being are relatively negligible.
- C) In the medical context, we should act so as to maximize life and health

If we accept C), then in any instance in which the hRR comes in conflict with saving a greater number of statistical lives the hRR should be abandoned. In short, C) means that the identifiable victim bias is unethical. Fried and Schelling, two of the earliest writers on rescue biases, are notable for adopting this simple consequentialist approach to the rule (Fried, 1969; Schelling, 1968). They anticipate that there are potential externalities on issues of rescue that affect overall well-being, like public trust or sense of societal caring, but both argued that the primacy of human lives must come before these considerations. To use the example of the flu pandemic, they would argue for a pure cost-effectiveness scheme in determining the allocation of ventilators.

Other thinkers have argued that the societal effects of the rule of rescue merit some weighting in a consequentialist well-being analysis (Rosoff, 2017; Brock, 2015; McKie, 2003; Maclean, 1986; Daniels, 2012). These theorists believe that a full consequentialist approach to the RR would incorporate societal preference in a meaningful way, rather than dismissing it as negligible compared to lives saved. In the flu pandemic case, they might incorporate some degree of partiality for those first identified in the flu pandemic, as opposed to considering future statistical victims of the flu as of equal consideration with current victims.

Fried himself presents this idea well. He notes that “often too we encounter people as actual persons” and “we demonstrate the value we place on human life” when we give some preference to rescuing identified victims, whereas ignoring their suffering undermines the “symbolic value” of human life (Fried, 1969, p. 1430). Note that this “symbolic value” argument is still a consequentialist argument: it is discussing the consequences of the RR in terms of maximizing overall wellbeing, arguing that symbolic value is an externality that carries greater weight than the simple consequentialist camp will acknowledge. It seems intuitively undeniable that, in instances where a person is identifiable, to deny them rescue is a particularly difficult decision. It is to look coldly on as a person suffers and dies. This is a personal, and therefore sympathetic, omission of care, and one that seems likely to greatly concern lay-people if such acts were regular occurrences in the healthcare setting.

To see the strength of symbolic value, consider the following scenario. Note at the outset that the scenarios are not plausible—various policies are in place to make certain that neither case can occur. Rather, I have selected the scenarios in order to control for the symbolic value of rescue as the sole determiner of our intuitive reaction to each case.

First, a child becomes severely ill with the flu, but the child is poor and turned away from a hospital due to inability to pay, and ultimately dies. Second, many poor children are unable to get the flu shot because the hospital didn't allocate sufficient funds to guarantee access for uninsured people in their area. One child subsequently becomes ill with the flu and receives a heroic level of care from the hospital, but still dies. Both cases feature hospital administrators devaluing the lives of the poor, at the cost of one life. Yet the first case (the case of an identified victim being denied care) seems more alarming and disconcerting. Identifiability, singularity, and reference group effects, among other factors, all conspire to make us feel that way. The second case (of a statistical victim) is intuitively less alarming—we may blame the hospital for negligence or discrimination, but we would not blame medical personnel for extreme callousness or heartlessness. If the second case were widely publicized, the public might demand greater accountability for vaccination programs broadly. But if the first case were widely publicized, angry citizens would be furious with the hospital. They would doubt the trustworthiness and beneficence of the hospital—perhaps even of healthcare facilities generally. To deny statistical victims access to interventions does not amount to denying a particular person care, and therefore lacks the psychological draw of singularity or identifiability, or the intuitive demand for compassion that a personal interaction has.

The above example illustrates one form of symbolic value realized by following the rule of rescue, what we can call “norms against dehumanization” or “norms that promote human fellowship”. People see acting with sympathy towards someone in an acute emergency as indicative of humanity. This is true from a perspective of the lay-person, but it is also true from a healthcare provider standpoint. As Norman Daniels (2012) notes, rescue squads or emergency physicians could easily become unmotivated or angry if required to act according to a “strictly

rational” budget, in which they are constrained from attempting feasible rescues that are costlier than a strictly rational budget allows.

A general critique has occasionally been levied at the validity of considering social norms in moral decision-making: there are many apparently natural norms that lead to clearly amoral decision-making, like our nativist (i.e. racist, nationalistic) tendencies (Daniels, 2012). From the consequentialist perspective, following these impulses may be enjoyable for the racist or the misogynist, but are clearly harmful (in a greater sense) for those discriminated against. The conclusion of this argument is supposed to be that social norms are not admissible when considering consequences. However, rescue impulses are qualitatively quite different from those undesirable social norms. In the first place, they are not inherently negative. For example, racist thoughts are hateful in a literal sense, whereas an impulse to rescue is literally altruistic (the remaining debate being whether it is immoral due to an omission of altruism for statistical lives). Further, rescue impulses have intrinsic worth — indeed, Allan Gibbard writes that ignoring rescue impulses “interferes with the direct psychological rewards of human fellowship”, where fellowship is understood as essentially person-to-person (Gibbard, 1986, p. 101). Given that rescue impulses do not have the inherent negativity of motives like racism, the only remaining complaint against the symbolic value consequentialist position is that it could lead to a greater number of lives lost. But the position already takes into account lives that might be lost due to incorporating symbolic value (it is a consequentialist position, after all!). The argument for the inadmissibility of social norms therefore seems to be advocating for either a double-counting of harmful effects or the dismissal of the RR impulse based on a mistaken equivalency with racism/misogyny, when in fact the symbolic value consequentialist has already accounted for all relevant considerations in their initial calculus.

In addition to the value derived from following humanistic norms, there are two other consequentialist externalities that give support to the rule of rescue. First, the professional role of healthcare providers is of symbolic value. As I noted earlier, this role is meaningful to the healthcare providers themselves. But the sanctity of the patient-physician relationship is also immensely important to patients. Indeed, to act against a presumed professional role has negative implications for the public perception of the profession at large, so it is essential that physicians are seen as allies of their patients, and not themselves forced to act against patients' interests.²⁴ The New York State flu plan accommodates this concern—a patient's physician does not participate in the decision to extubate a decompensating patient, nor are they required to perform the extubation themselves. Similarly, the initial triage decisions are determined by triage physicians and nurses who do not meet any of the patients, and therefore never enter into the symbolic physician-patient relationship.

The second externality argues that utilizing the rule of rescue frequently encourages good behavior and motivations. For example, recall the psychological study that found that playing up the identifiability of a person in need increases donations, whereas utilizing statistics about needs drives down donations (Small and Loewenstein, 2003). Similarly, it could be the case that a hospital can secure greater donations and community support for a pediatric bone marrow transplant program, which helps identifiable children, rather than to support their flu vaccine drives, which aim to prevent harm to statistical victims.²⁵ Per dollar, the flu vaccine drive may be

²⁴ As mentioned in a previous footnote, there are a few, well-accepted exceptions. The duty that physicians have to their patients is balanced with other duties they have to society. Physicians, for example, are required to report instances of certain infectious diseases to public health authorities, regardless of concerns around patient confidentiality.

²⁵ Beyond our general bias in favor of identifiable victims, people also have a bias in favor of drama and excitement broadly. Blood marrow transplants are a more dramatic and gripping form of treatment than the relative boredom of mass vaccination efforts. But this exciting-boring dichotomy is true of most non-emergency public health measures, and so can be considered as its own feature of the identifiable victim bias.

more cost-effective, but it could be that dollars are more easily secured for the bone marrow program such that the hospital can actually save more lives by focusing on marrow transplants—meaning that consequentialism rules in favor of bone marrow donation drives. Beyond these specific cases, it could also be true that a general societal proviso for encouraging an identifiable victim bias encourages overall life-saving behavior, such as greater overall altruism, empathy, and trust throughout society.

While it seems clear that there are externalities—norms of human fellowship, professional duties, and encouraging altruism—that may lead the consequentialist to incorporate the RR into their framework, the effects of these considerations are not as easily quantifiable as lives saved.²⁶ This is exactly the critique of the non-symbolic value consequentialist. They will admit that symbolic value effects exist, but say they pale against concerns about life-saving and further that it is unproductive to consider them since they are unknown. Symbolic value consequentialists disagree and think the externalities are significant and warrant consideration. Ultimately, it is a currently unknown empirical question how much the rule of rescue impacts social wellbeing, trust in professional duties, or overall altruistic behavior for people. This makes it very tricky to determine how much weight healthcare leaders should give the hRR during decision-making, and will be a major component in my later conclusion that only a deliberative process on a case-by-case basis can consistently provide adequate and sensitive ethical decisions on RR dilemmas.

²⁶ The exception could be the increase in altruism due to rescue impulses. Quantifying this increase is the focus of many of the previously discussed psychology articles (see Jenni and Loewenstein, 1997; Small, Loewenstein, and Slovic, 2007; Kogut and Rotov 2005a; Ibid., 2005b). Still, it is far more difficult to quantify the increase in altruism realized specifically in healthcare due to prioritizing the hRR. Potential research could look at donations for treatments versus preventive programs (for example, bone marrow transplant services versus prenatal checkups).

Decreasing the Pull of the Rule of Rescue

A frequently ignored implication of the symbolic value argument is that the amount of symbolic value a society derives from rescue is not a constant. Rather, it is relative to the amount of symbolic value a particular population places on the hRR, and to the broader context of the rescue decision itself. If that population had less regard for symbolic value without a decrease in sense of wellbeing, faith in the health professions, or net altruism, then their health institution should change its ethical calculus to lean more in favor of the non-symbolic value consequentialist position. If we could encourage this, then it would be a good – desirable effects (lives saved) would increase, and negative effects (symbolic value lost) would not occur.

Indeed, it seems likely that there are ways to decrease the psychological pull of rescue for all three of the above considerations. For example, it has been noted that the hRR is likely strengthened by the public's image of health insurance companies as "heartless and faceless" organizations that "[make] decisions in an opaque manner" (Rosoff, 2017, 146). Accordingly, creating transparent, reasonable, and responsive systems for adjudicating issues around the hRR, open to public input, could help decrease some of the RR's pull. As previously noted, the philosopher Peter Railton believes that greater exposure to rationing situations could cause changes in how lay-people and healthcare administrators think of rescue versus effective allocation (Railton, 2015). Professional role duties can also be accommodated by building safeguards around patient-physician relationships, such as how the New York ventilator allocation scheme delineates between treating physicians (those who care for patients) and triaging physicians (those who never assume patient care but make rationing decisions). Finally, as discussed in the psychology section, certain psychological factors have been shown to be

malleable in determining rescue impulses. In particular, portraying statistical victims as a unified group decreases the pull of the singularity affect for identifiable victims.

Finally, consider again the argument that the rule of rescue may play a role in driving overall altruistic behavior, and therefore should be safeguarded. This may well be true at a personal level, but I am skeptical that healthcare moving more towards cost-effective rescue versus indiscriminate rescue would greatly influence the general altruism of lay-people. In the first place, the vast majority of rescues will always be attempted by healthcare institutions—the US is a wealthy nation who can bear those costs, and the vast majority of acute emergencies are cost-effective to respond to. Further, identifiability can always continue to be leveraged in fundraising or storytelling around healthcare, maintaining funding levels and the sense of individual caring that is indubitably essential to medicine. Individuals and private charities could fill gaps left by healthcare institutions. Further, identifiability and rescue would also continue to be visible in myriad non-medical areas of society, like police response and fire rescues. As discussed in the psychology section, rescue impulses are inherently human, and for that reason there will always be symbolic value ascribed to the rule of rescue—but the above arguments demonstrate that we also can and should act to increase a general rather than focused approach to distributing health benefits.

There is a final consequentialist objection to incorporating symbolic value that merits brief consideration. Cookson et. al (2008) writes that, given the difficulty of predicting the externalities associated with implementing a cost-effectiveness analysis, the policymaking default should be to act from a pure life-optimizing standpoint and only consider symbolic value if and when the public objects with reasonable force to a policy decision. This view is problematic because it would rationalize allowing ethical discussions in the healthcare setting to

be indeterminate and reactive. Healthcare institutions need to be able to articulate their ethical reasoning and demonstrate consistency, so as to maintain respectability and exhibit a principled approach to their mission to help their communities. A possible litmus test is the following question: would I be comfortable explaining this decision, and the decision-making process, to the public? In order to answer in the affirmative, full consideration of both lives saved and symbolic value should be considered at the outset and defined as clearly as possible for the question at hand. If later public response generates valid considerations that were not previously discussed, it can warrant a change in policy. However, the default for the healthcare institution should be to make defensible, measured decisions.

Summary of Consequentialist Considerations

In sum, the consequentialist approach boils down to an argument between those who believe in a pure life-saving approach and those who argue that there is “symbolic value” to showing some partiality towards identifiable victims. This symbolic value is evident in the sense of well-being and social trust engendered by rescue considerations, the importance placed on the sanctity of patient-physician relationships, and the boost in altruism evoked by rescues. There is reasonable disagreement about how much these symbolic value considerations should be weighed in real decision-making situations, and so a transparent, deliberative process should occur for each major decision (later I will discuss in greater depth how this process might look). Moreover, any action that would allow institutional actors to decide closer to the act-consequentialist position without negatively affecting symbolic values should be taken. These actions could include establishing a public and fair deliberative process in decision-making and placing strong safeguards around patient-physician relationships.

The New York State ventilator allocation plan provides an excellent example of how this consequentialist approach might look in practice. Remember that this is a case in which there is a scarcity of ventilators due to a flu pandemic. Allocation is determined ultimately by a consequentialist scheme in which patients are rated by survivability²⁷ and then given a “trial period” on ventilator support. If ventilator support upon reevaluation does not appear to be leading to improvement, patients are removed from the ventilator and the device is transferred to another patient who is likely to benefit more. While the arrangement of removing patients from life-saving care (effectively abandoning them to die) seems like it would raise major concerns about symbolic value, the workgroup made strategic decisions to allay these concerns: they clearly outlined the reasoning behind their plan, invited public comment on their plan and reasoning prior to implementation, ensured that treating physicians would maintain the traditional physician-patient role, and developed an appeals process for any patients or physicians who might feel unfairly treated. Finally, special care was made to ensure that ventilator allocation would be decided as equitably as possible. This means that triage physicians would only see objective markers of health when determining allocation, not subjective stories, photos, or arguments.

This was an essential provision to the plan, since the medical establishment, unfortunately, has a long and tragic history of discrimination in care. Indeed, the primary deontological concern aimed at the hRR hinges on discrimination and unfairness—both that identifiability leads to inherent biases against certain groups (i.e. out-groups), and also that a bias

²⁷ As previously mentioned, “survivability” as their core criteria is the only deviation from a traditional consequentialist approach because it does not consider total years of life saved. While the allocation plan considers co-morbidities like end stage cancer or advanced heart failure when considering “survivability”, it is not thinking explicitly in terms of years of life saved. For example, there is no preference shown based on age. Interestingly, the North Carolina ventilator plan is largely based off of the New York State plan but includes age as a consideration when determining allocation (Draft: North Carolina Triage Protocol, 2010).

in favor of the identifiable is inherently discriminatory against statistical lives. These arguments are the focus of the following section.

Deontological Considerations

Beyond considerations of maximizing utility (or good consequences) there are also non-consequentialist arguments against the hRR. Here, I will detail two such arguments. I will counter the first – on deontological grounds – and agree with the second. Both arguments rely on a very general and widely accepted principle: that all lives deserve equal respect, where equal respect is defined as avoiding morally irrelevant discriminatory considerations. Due to this shared starting point, I will call both arguments a kind of “discrimination argument”.

There are two primary varieties of the discrimination argument against rescue impulses: the first, and more general, of the discrimination arguments states that acting on the identifiable victim bias constitutes discrimination against statistical lives (or alternately, statistical persons). This argument, if true, would always lead to the conclusion that the identifiable victim bias is discriminatory and therefore immoral by the equal respect principle. The primary discussion in the literature has focused on this version of the discrimination argument, which I will call the “general case of the discrimination argument”. I will argue that appeals to distributive justice, which are the commonly utilized counter to this argument, may be valid, but that even if these appeals are accepted the final conclusion only allows an insignificantly small bias in favor of identified lives. However, I will go further and propose, using a Rawlsian veil of ignorance argument, that discrimination against statistical lives is incoherent in and of itself—that statistical lives are not the sort of lives that can be discriminated against.

The second, and more narrow, of the discrimination arguments states that oftentimes rescuers will elect to help one group over another similarly situated group based on in-group

biases or system-level biases, and so discriminate based on race, gender, geographic reason, etc. I will agree with the narrow discrimination argument, but I will emphasize that it is not a general takedown of the hRR or the identifiable victim bias; rather, it requires case-by-case contextual consideration.

Using Distributive Justice to Bypass Deontological Critiques: An Impractical Argument

The classic move in the literature to get by the general version of the discrimination argument is to appeal to distributive justice. Distributive justice, as formulated to support the identifiable victim bias, is the notion that the most equitable distribution of survival odds is preferable over an unequitable distribution of survival odds (Daniels, 2012; Hare, 2012; Lübke, 2019; Frick, 2015). If distributive justice is a moral consideration, then it is possible to say that a certain degree of preference for identified lives over statistical is not a case of discrimination against statistical lives, but rather giving appropriate preference to the imperative to equally disperse harm. But many deontological philosophers (and some consequentialists) have engaged in vigorous debate about whether distributive justice (“risk pooling”) is a moral consideration.

The disagreement relates back to the aggregation problem, a larger ethical issue. The aggregation problem asks the question: is it better to help one person a great deal, or a very many people a tiny amount? Is it better that one person die so that a million people don't have to endure the nuisance of a mosquito bite? Most deontologists agree with the common-morality intuition that it is better to help one person a great deal, and its corollary that it is better for many people to suffer a little than for one person to suffer a lot (this position is called “anti-aggregationist”).²⁸ Defenders of the morality of the identifiable victim bias against the general

²⁸ This, of course, is given that the harm and benefit really can be distributed among many in a meaningful way. It is often the case in healthcare that to distribute aid equally may not be of substantive help to anyone. For example, a

form of the discrimination argument generally support this train of thought. If you believe that distributing harm across many people is better than concentrating the same amount of harm towards one person, it is congruent to believe that sufficiently diffuse harm, even if greater in absolute terms, may be preferable over more concentrated harm.

Take the example below, adapted from Daniels (2012). Note that in both cases the expected outcomes in terms of lives lost are equal, at 1 expected life lost.

Great harm to one: Alice has a disease and needs five tablets of a medication (the only tablets available) to recover. She will certainly die without the tablets.

Commensurate harm to many: Five others have been exposed to the disease but have not shown symptoms. They can each be saved by taking one tablet. Without taking a tablet, each has a 20% chance of acquiring the full disease and dying.

Daniels contends that we have a greater obligation to help Alice in the above scenario. He argues that this is because Alice faces certain death without treatment, whereas the others face much better odds. Alice, he says, is facing greater harm. Per the general solution to the aggregation problem, we should disperse harm across the many rather than allow one to face great harm. Accordingly, distributive fairness demands that, in cases of equal expected outcomes of harm, we distribute risks rather than concentrate them.

Many theorists disagree with Daniels and company, arguing that this is a misapplication of distributive justice. Alice is not facing greater harm than her five friends these critics say; the relevant harm is *lives* lost, not the risk of dying (Adler, 2015; Eyal, 2015). If one person dies out of the five, then that person is harmed just as badly as Alice would have if she was denied the medication. Both died due to lack of the medication. This is a 1:1 comparison, not an instance of

novel vaccine regimen may not be effective unless a full dose is taken. It would then be useless to everyone to “distribute” a full dose of the scarce vaccine evenly among 100 people.

aggregating harms across many. As Nir Eyal says, “equality is fairer than inequality, but that ought to mean equality of good and bad outcomes, not equality of risks and good prospects.” (2015, p. 97).

However, the debate does not end here. Hare (2012) counters Eyal by arguing that reasoning in healthcare allocation can only operate by using counterfactual conditional probabilities, and so this is the perspective from which aggregation must be considered. Thus, risk distribution is the correct way to determine fairness. Since healthcare allocation always relies on conditional probabilistic depictions of harm, we must consider those probabilities as depicting real harms in cases like Alice’s.

Another thinker, Lübbe, argues that the correct perspective to adopt in thinking about these issues is that of the rescuer—in his terminology, “the choice-situation of the decider.” (Lübbe, 2019, p. 55). Say that Mary is deciding who to save between Alice and the five others. If she chooses to save the five, and Alice dies, it can correctly be said that Mary left Alice to die. However, if Mary decides to save Alice and one of the five dies, it cannot be said that Mary left that person to die. Rather, Mary failed to lower that person’s risk of dying. Since this harm was not located in a specific person’s life, it seems less morally reprehensible from the perspective of the decider (that is, Mary).

I personally find Hare’s argument moderately compelling, and Lübbe’s argument less so, as it seems to me to rest largely on semantic considerations. Regardless, even if proponents of this form of distributive justice are correct, it would still be either a remarkably weak argument for the identifiable victim bias or an overtly strong one. If distributive justice was taken to the extreme, then all help would need to be directed towards those at highest risk of harm regardless of the inefficiencies incurred. This is untenable. Placing limits on how much preference should

be given to distributive justice reduces most compellingly to tie-breaking (or very nearly tie-breaking) cases. Indeed, it would be difficult to argue that if the Alice case was modified to place the 5 friends at a 40% risk of death (expected deaths = 2) that we should then opt to help Alice. With distributive justice as only a tie-breaking (or nearly tie-breaking) consideration, it only provides a sliver of justification for bias towards individuals, and certainly nothing like the robust preference that individuals and institutions display in aiding the identified over the statistical in many instances. Distributive justice considerations would therefore only very rarely apply in the real choice situations of healthcare institutions.

A New Counter to the General Discrimination Argument: A Rejection of the Possibility of Discriminating Against Statistical Lives

So I have argued that distributive justice considerations can do little to rescue the identifiable victim bias, in a practical sense, from the general discrimination argument that the identifiable victim bias constitutes discrimination against statistical lives. However, another argument can move past the general discrimination argument: one that outright states that statistical lives are not the sort of lives that can be discriminated against. Here, I contend that the identifiable victim bias does not constitute differential treatment due to systemic bias against a distinct group of persons. I argue that there *is* no discrimination.

Before I present the principled reasoning for believing this argument, I wish to motivate the case with some general appeals. As noted at the outset, deontological theories of fairness generally result from a consideration that all human persons warrant respect (Kantian theories) or that all human persons enjoy a fair array of rights and protections, without consideration of social differences that arise by chance (Rawlsian theories). Notice that both veins of thought focus on

respect for distinct persons, who may have varying identities and goals as autonomous agents, and so move to argue that it is unethical to discriminate based on the identities of those persons.

From here, the cases of discrimination that theorists (and lay-people utilizing common-morality) are primarily concerned with are those in which there is a distinct, fixed identity group that is the focus of differential treatment within a community. Racism, sexism, nativism, classism, even ageism all discriminate against (relatively) clearly defined groups. A person *is* black, or female, or a coal miner, or a migrant, or an old/young person. It is a feature of their identity. A person is a black person. However, there is no such thing as a statistical person. Someone may be considered from a statistical perspective on one health issue (say, they are at risk of developing heart disease) and be considered from an identified perspective on another health issue (their emphysema treatment). To discriminate against black persons is to fail to follow the deontological provision to treat all persons as worthy of respect. But to discriminate against statistical lives as opposed to identified lives is not to discriminate against any person or even group of persons in particular.

Another way to see the strangeness of considering “statistical” as a category for discrimination is to consider how deontologists typically reason about rescue versus allocation cases, such as the New York ventilator allocation case. The general discrimination argument would lead to the same conclusion as the pure consequentialist position: that it would be discriminatory to provide care preferentially to the early arrivers as opposed to the statistical lives of those who could be predicted to soon arrive. However, while deontologists are concerned that we do not discriminate against persons, their theories tend to be skeptical of thinking statistically (as this is associated with consequentialism, and with aggregation). Yet in this case the claim is that thinking only from the individualistic perspective is discriminatory and evidence

of a lack of respect for statistical lives. The tension here is apparent – either deontologists are mistaken in holding together both a skepticism towards statistical morality and a high value for equal respect for persons, or else the general discrimination argument misses the mark. Consider two examples from famous deontologists, both of which depict a preference for the identifiable at the expense of the statistical. Kant is famous for holding that each human life is of near-infinite value such that a person can never be treated as a means to an end. The result would be that no patient could be removed from a ventilator in order to make it available for another, because to do so would be to treat the first patient as a means to an end. Kant’s position entails that one should never engage in prognosticating around optimizing for lives saved when to do so requires removing rescue assistance from currently treated persons (Kant, 1785). Similarly, Thomas Scanlon argues that the individual is the correct focus of moral concern— in the case of the identifiable victim bias, this would mean that a preference should be shown for the most ill victims, rather than aggregating statistical outcomes to ensure non-discrimination (Scanlon, 1998).²⁹

Finally, consider an argument derived from Rawlsian theory itself. The basis of Rawlsian moral theory is the “veil of ignorance” argument, in which we are to imagine a rational person placed in the “original position” who has to develop a social contract (Rawls, 1971). In the original position, the rational person knows the consequences of various decisions, the general types of people that will exist in society, and facts about human psychology, but they do not

²⁹ Kant and Scanlon’s positions do have different results in the case of identified versus statistical lives. Both theorists would allow for triage in the case that care has not yet been initiated (indeed, Scanlon’s moral account would require it). Kant’s position is that a person cannot be used as a means for an end, which results in the conclusion that a person could not be removed from care in order to make space for another. Scanlon would allow care to be transferred to a patient in greater need; however, his account focuses on aiding the individual in greater need, regardless of aggregation concerns (aggregation is only permitted if there are multiple individuals with the same degree of need).

know which person within the society they are. From this original position, Rawls holds that all cases of group-based discrimination, like misogyny, racism, etc. would be considered immoral and against the arrived at social contract. This result occurs because the decision maker would wish to avoid the chance of being the one discriminated against.

However, their position on the identified victim bias is more complex. It certainly is a different case than for racism or misogyny. The decision maker would consider that they will likely be both identified and statistical throughout their lives, rather than occupying a specific identity with respect to identifiability in the way that persons occupy a specific racial identity. Given their innate knowledge of human psychology, they would consider the humanity inherent in the impulse to rescue, and its effects on social trust, overall community altruism, and individual well-being. They may decide that a certain degree of preference for the identifiable victim bias is reasonable.

Another way to see this conclusion is by analogy. Imagine that the decision maker is considering whether to permit a bias to help one's family over other individuals. This bias could be called "discriminatory" in favor of the family over other members of the society. And yet it seems quite likely that the decision maker would allow for some degree of family bias out of an understanding that human psychology demands it. To not allow for the family bias would be to misunderstand basic human needs and special human duties to those close to us. Similarly, to not allow for some degree of the identifiable victim bias would be de-humanizing.

Thus taking the approach of Rawls' veil of ignorance leads to a similar conclusion as the symbolic value consequentialist: one should indulge the identifiable victim bias insofar as the symbolic value of rescue as a human psychological reality outweighs other concerns.

These “other concerns” that the Rawlsian arrives at are essentially those of the narrower discrimination argument. To see this, let us continue with Rawlsian theory. Rawls argues that the decision maker will arrive at two general principles through which to structure society: 1) the Principle of Equal Liberty, which states that each person has an equal right to the most extensive liberties compatible with similar liberties with all; 2) the Difference Principle, which states that social and economic inequalities should be arranged so that they are a) to the greatest benefit of the least advantaged persons b) attached to offices and positions open to all under conditions of equality of opportunity (Rawls, 1971). Of these two principles, the Difference Principle is the relevant source of concern for the Rawlsian considering the identified victim bias. Using the Difference Principle, the Rawlsian will focus on whether certain instances of the identified victim bias lead to undue harm to the least advantaged person, where “least advantaged” is measured by the metric of overall wellbeing (including intellectual ability, social position, etc.). Thus the Rawlsian will be concerned with limiting unnecessary inequality among the socially bad-off — the poor, marginalized minorities, the disabled — rather than terming the identified victim bias discriminatory writ large.

Discrimination Based on Social Identity

So we arrive at the narrower discrimination argument, that discrimination against the socially marginalized is morally undesirable and should, all else equal, be avoided when deciding to act on rescue impulses. At this point, I am in full agreement that this type of discrimination is valid. And it is a ubiquitous and heterogeneous form of discrimination, with powerful implications for rescue.

For example, perhaps an intervention to provide cancer treatment to identifiable individuals could be complicated because a minority group is less likely to report early to

hospitals for that form of cancer, due to mistrust of the medical system. In that case an identifiable approach would reach relatively less minority individuals than a statistical prevention approach, so opting for the identifiable intervention could be unfair to the minority group. Similarly, many individuals are predisposed to certain health outcomes based on their social identity. An initiative to improve cosmetic therapies rather than to target obesity prevention could therefore discriminate in favor of the wealthy at the expense of minority groups and the poor.^{30,31} Further, there may be cases where, due to socially unattractive features of an individual, the identifiable victim bias is turned on its head. This is the so-called “Rule of Abandonment”, in which people have a prejudice not to rescue certain undesirable individuals, like ex-prisoners (Rosoff, 2017, p. 147). The narrow discrimination argument would correctly find this rule unjust.

There are also individuals who could benefit by being uniquely desirable. For these individuals, their ability to generate publicity and garner public sympathy may be unjustly distributed, and therefore discriminatory. Consider the case of Sarah Murnaghan, a 10-year-old girl with cystic fibrosis in need of a lung transplant. Sarah was placed on a waiting list for pediatric donor lungs but was deemed ineligible for adult lungs due to a scarcity of evidence from which to determine their transplant efficacy in a child. As it became clear that Sarah would not receive pediatric donor lungs in time to save her life, her parents engaged the national media, petitioned Congress members, and enlisted the help of a prominent law firm (DeSante et al., 2014). Ultimately, Federal Judge Michael Baylson ordered that Sarah be placed on the adult lung

³⁰ This only holds true if the money being allocated is coming from a common pot. This is the case in centralized health systems, but is not the case in the United State’s fractured system.

³¹ Interestingly, if the choice was between a treatment and prevention initiative targeted at the same disease, then there could be no discrimination (unless there is a difference in how frequently a particular group accesses care, as in the preceding example). This is because the relative percentage of the groups predisposed to the disease would be constant relative to either decision.

transplant list. Many commentators (and I share their opinion) consider this sort of one-off exception to objective criteria an instance of discrimination: another person, even a child, in Sarah's position may not have had parents with sufficient media and legal know-how to overturn standing allocation guidelines (DeSante et al, 2014; Evans, 2013).

The narrow discrimination argument has a final instance in which it is realized: that of system-level differences. Different hospital systems have different rationing schemes. For example, the New York ventilator allocation scheme has subtle differences from the North Carolina scheme (the primary difference being that North Carolina weighs age in their allocation algorithm, whereas New York's scheme does not consider age at all) (New York State Guidelines, 2007; Draft: North Carolina Triage Protocol, 2010). The disjoint nature of US healthcare means that different decisions of life-or-death importance are being realized in the same situations due simply to different leadership and legal environments between hospitals.

Summary of Considerations Around Discrimination Arguments Against Rescue Impulses

In sum, the discrimination argument takes two forms. The more forceful version states that it is discriminatory to favor identifiable lives over statistical lives. The traditional rebuttal, that distributive justice can provide a rationale for a bias in favor of identifiable lives, is potentially valid. However, it only applies in cases where there is very nearly the same expected number of identified versus statistical lives at risk. A more general response to the forceful discrimination argument is available, which is to claim that statistical lives are not the sort of lives that can be discriminated against. This claim, in turn, can be motivated through considerations of typical notions of discrimination cases, an analysis of disagreements within well-known deontological theories on how to handle the value of predicted lives versus contemporary lives, and through Rawlsian moral theory.

This leaves the weaker discrimination argument remaining as valid. The argument, that institutional decision makers need to weigh the potential for unfair discrimination against certain groups of people in rescue situations, is indubitably morally salient insofar as concerns about racism, sexism, classism, and the like are moral concerns. Notice, however, that this weaker discrimination argument is not an in-principle argument against the identifiable victim bias in all cases. Rather, it rests on empirical questions about what sorts of effects electing for (or against) a particular rescue allocation decision might yield for particular groups. Further, there is no clear-cut method for weighing the relative severity of a particular discriminatory decision; instead, the most morally sensitive manner of identifying and weighing these considerations is through a fair, transparent, deliberative process. In the following, and final, section I will elaborate on what exactly this deliberative process might look like, and I will summarize the lasting points of debate around rescue impulses as a checklist, pulling from both the preceding consequentialist and deontological sections. I will conclude by demonstrating how the checklist could be used in a variety of historical case studies to arrive at morally sensitive decision-making.

Chapter Three: From Theory to Practice

1. Procedural Justice and a Checklist for Institutional Use

In the previous consequentialist and deontological sections, I have aimed to demonstrate the complexity of moral reasoning around the hRR and the identifiable victim bias in particular. From a consequentialist perspective, I have argued that there exists some symbolic value to the identifiable victim bias which must be taken into account. However, the exact amount of this symbolic value is empirical and unknown, and both highly context-dependent and largely open to debate. From a deontological perspective, the identifiable victim bias could lead to (or mitigate) immoral discrimination against certain groups of individuals. Again, the exact type and moral weightiness of such discrimination need to be adduced for each case at hand, and there exists room for disagreement on any given case.

Given this room for contextual variety and reasonable disagreement, procedural justice (sometimes called “accountability for reasonableness”) is required to arrive at a mutually acceptable, morally sensitive, and publically defensible resolution to each rescue allocation case (Daniels and Sabin, 1997; Daniels, 2000; Daniels, 2012). Daniels and Sabin, in addition to John Rawls, have written extensively about procedural justice in a democratic society. A general definition is as follows: procedural justice involves a “fair, deliberative process in which conflicting views are considered and rationales are developed” and is “a fully transparent process that involves appropriate stakeholders who broaden the deliberation about what reasons should be considered, and that allows decisions to be revised in light of new evidence and arguments” (Daniels, 2012, 42). In essence, procedural justice looks to bring relevant decision makers (including community members) who may have varying perspectives on a nuanced issue to the table and allow sufficient time for varying views to be considered and mutually acceptable

resolutions adopted. Further, the resolution reached is both transparent to the community impacted and “defeasibly fair”, meaning that the wider community has the ability to comment on the initial resolution,³² and that the initial resolution is open to reconsideration if additional evidence or viewpoints are raised (Daniels, 2012, p.42). Thus, democratic procedural justice requires 1) a group of decision makers that broadly resemble the community impacted by the decision 2) fair consideration of all views, and an agreement that is mutually acceptable to different parties 3) transparency to the wider community 4) defeasible resolutions, which can be reconsidered in light of new evidence.

Such a group is particularly useful for providing morally sensitive decisions in contentious issues around the hRR. In terms of identifying the symbolic value of a rescue, a group that resembles the community affected is more likely than a homogenous group (say, only of physicians or health administrators) to correctly determine the symbolic value of that rescue decision to the community. Similarly, a diverse group of decision-makers is more likely to identify decisions that have the potential to lead to discrimination. Finally, the presence of varied community members on the committee adds to the moral legitimacy of the final decision and decreases the potential for complaints that an institution is out of touch, or only focused on the bottom line.

³² There is disagreement about how much publicity of rationing decision-making is good, and how much is bad. For example, the NY ventilator allocation plan was aired for public comment, but on a government website and without substantial fanfare (New York State Guidelines, 2007). The North Carolina ventilator allocation plan was not aired for public comment at all, for fear of needlessly inciting concern in the public. These considerations are valid in today's media climate, in which sensationalized headlines often are more valuable to media outlets than objective reporting. Still, providing transparency about a decision-making process and employing a democratic process in which the public is permitted comment are essential to the procedural justice of a decision in which there is reasonable disagreement (Draft: North Carolina Triage Protocol, 2010). “Getting ahead” on media coverage can help mitigate potential bad coverage as well. Thus, decisions on rescue versus statistical lives need not be widely publicized, but public comment should be at least minimally permitted and solicited.

From a practical standpoint, a procedurally just taskforce to make rescue allocation decisions could borrow their structure from existing patient advisory committees, or groups like the Duke Critical Drug Shortage Taskforce and NICE's Citizen's Council (Citizen's Council Report: the Rule of Rescue, NICE, 2006; Rosoff, 2012; Rosoff et al., 2012). The ideal group would include medical experts and technical experts on the question at hand (for example, pharmacists in the case of a drug shortage), ethicists, and community members, and are structured such that fair deliberation can occur. These taskforces should be called for any major contentious allocation decision on rescue versus statistical lives—for example, a ventilator allocation task force, a critical drug shortage task force, or even a budgeting task force on allocating funds to opioid prevention versus bone marrow transplant.

For such discussions it would be useful to have a clear framework for the controversial issues at play. The committee could then move through the framework, filling in empirical or context-dependent features of the situation as they see them. Below, I have included just such a framework for use, in the hope that it both increases the moral sensitivity of final decisions, as well as promotes speedy and mutually understandable discussion among committee members. There is certainly room for additional features of the case that may be beyond the scope of my checklist; nevertheless, I am cautiously optimistic that the checklist captures the points that are both most difficult to resolve and most integral to a morally correct decision in the majority of cases.

Conflict Between Rescue Impulses and Lives Saved: A Framework for Allocation Decisions

1. Identify the Contentious Rescue Impulse

Is the allocation decision an instance of rescue impulses versus statistical lives?

Allocation decisions that feature a conflict between:

- Allocating resources to the rescue of identifiable victims, i.e. the victim is presented as an individual, requires immediate care, and there is a clinically meaningful intervention available
- Allocating resources to an equal or greater number of expected statistical victims, i.e. victims presented as one of many who are predicated to face life shortening or life-altering consequences

2. Consider Health Cost-Benefit Analysis

What is the relative efficacy, in terms of years of lives saved and quality of life, for each decision?

3. Consider Symbolic Value

What is the societal value to assisting the identifiable victims in the case?

- Will either decision affect societal well-being? Individual level well-being? To what extent?
- Will either decision detract from societal level altruism or charitable behavior?
- Will either decision impact healthcare providers sense of fulfilling their own professional duties?
- Will either decision impact the medical trust of patients? To what extent?
- Can any of these concerns be mitigated, for example through the articulation of clearly reasoned guidelines or by proactive transparent processes?

4. Consider the Potential for Discrimination

Is there potential for systematic discrimination against or for a particular class or group?

- Will either decision cause any specific class or group of individuals be overly harmed or benefitted relative to another? To what extent?
- Is media attention revealing an ethical issue that should be resolved writ large (a generalizable change in policy)?
- Is the media pandering solely to sympathetic psychological impulses, but impulses that are not morally salient and indeed discriminatory in favor of the socially attractive?

5. Consider Additional Practical Concerns

Is procedural justice being followed? Are there externalities to either decision, such as the potential for additional revenue that could be used for life-saving purposes?

- Is the deliberative process fair, diverse, transparent, and defeasible?
- What quantifiable externalities might result from either decision, such as revenue or technical acclaim?
- Consider qualitative externalities, such as whether a decision contributes to research or has epidemiologic implications.
- Is irrational exuberance building around a particular issue? Is this distorting the ethical conversation?
- Consider precedent/sustainability: is the rationale behind the final resolution, if generalized to all cases of the type in question, sustainable? Does it set an acceptable precedent?
- What other considerations might be relevant to the particular case?

6. Weigh the Considerations

Compare the forcefulness of the cost-benefit analysis, symbolic value, discriminatory considerations, and additional practical concerns to arrive at a morally justifiable decision.

Other frameworks have been proposed for resolving issues of identified versus statistical lives; however, these frameworks have generally been proposed by either pharmaceutical groups for the specific case of experimental drug use (PBAC Guidelines, Last Revised 2016) or were developed by patient committee rather than through systematic theoretical considerations (Citizen’s Council Report: the Rule of Rescue, NICE, 2006). Thus the above framework is novel both in form and in the process through which it was derived. All frameworks for structuring discussion around the hRR have been rightly criticized for the broadness of the checklist items and the unfortunate fact that moral weights are left unspecified (Cookson et al., 2008). However, asking the correct questions is invaluable in focusing deliberation and arriving at morally acceptable, and justifiable solutions.

Further, I will highlight the practical utility of the framework by applying it to various historical (or likely to occur) case studies. I am not a taskforce myself, of course, and so cannot mimic the types of discussion that might result from those sessions. Rather, I will step through each of the six questions in each case, propose what I think to be the reasonable answer (and weight) for each question, and thereby arrive at a morally justifiable resolution. Each case study has different salient characteristics, and I have included both cases in which the rescue bias is supported and rejected.

2. The Rule of Rescue and Exemplar Case Studies

Case One—Ventilator Allocation: Saving Lives Versus Symbolic Value

First, we will consider the New York ventilator allocation case that was discussed in both the consequentialist and deontological sections. To recap: a flu pandemic has broken out, and the demand for ventilators to rescue patients in respiratory failure has outstripped supply. How should ventilators be allocated among patients? The two clearest possibilities are either to

allocate on a first-come-first-served basis (as a pure identifiable victim bias would contend) versus allocation based on an objective scheme, in which predicted incoming patients would take precedent if they seem more likely to survive on ventilator support than a current patient. Below I apply the framework to this case:

1. Identify the contentious rescue impulse: The case involves conflict between current (identified) patients in need of immediate rescue from a life-endangering condition with a viable treatment (ventilator use), and other patients who are not yet in the process of being rescued but who could benefit even more from care.
2. Cost-benefit analysis: Significantly more years of life will be saved in the case of an objective scheme. For argument's sake, let's say that an analysis finds that 10-20% more years of life will be saved with the objective scheme.
3. Symbolic value: Individual patients and families may be angered if they are not prioritized for ventilator support; however, society at large is likely to understand the need to prioritize those who are more likely to survive and who have a longer expected life span (i.e. children) in the case of scarcity. A more concerning aspect of the objective scheme, however, would be the removal of ventilator support from a current patient to a new patient with better survival chances. This constitutes harming the current patient and risks larger societal backlash as well as the potential for conflicts with the mandate for doctors to act in the best interest of their patients. So there are moderate-to-significant symbolic value concerns at play. A key question to consider, however, is whether there are means to mitigate those concerns. In this case, there are: triage physicians can make determinations of removal of ventilator support rather than treating physicians, freeing the treating physicians from conflicts with their duty to care for their patients. Further,

transparent and proactive discussion and publicizing of the plan can be conducted proactively in the community, allowing concerns to be voiced, addressed, and potentially mollified. Utilizing procedural justice would also increase the justifiability of the final resolution if critiqued by the public or the media.

4. **Discrimination:** There is no overt discrimination in either policy. The objective scheme may disproportionately exclude minority groups from ventilator support, since minority individuals are more likely to have pre-existing medical conditions that could be rule-outs. This is unfortunately the case in organ transplantation as well (Higgins and Fishman, 2006). Still, this degree of discrimination is likely minor within the realm of health inequities, especially compared to the number of lives saved. Further, it is likely that more minority life years will be saved under the objective scheme than the first-come-first-served scheme, even if the objective scheme may disproportionately exclude minority individuals.
5. **Practical concerns:** Procedural justice, in the form of an institutional IRB discussion, should be followed in determining the outcome of the case. There are no clear quantifiable externalities, issues with the precedent that could be established, etc.
6. **Weighing the considerations:** In sum, following the objective criteria would save a significant number of additional lives, cause moderate symbolic value concerns that could be mollified through transparency and community inclusiveness, and faces minor discrimination issues. Overall, the value of the additional lives saved appears to outweigh the relatively minor symbolic value and discrimination concerns such that the objective criteria resolution should be adopted.

Case Two—Sarah Murnaghan and Media Attention

Sarah Murnaghan was a 10-year-old girl with cystic fibrosis in need of a lung transplant. She was placed on a waiting list for pediatric donor lungs but was deemed ineligible for adult lungs due to a scarcity of evidence from which to determine their transplant efficacy in a child. Her parents actively petitioned on her behalf and enlisted the help of the national media, members of Congress, and legal aid. An exception was ultimately made for Sarah, placing her on the adult lung transplant list. Was this decision correct?

1. Identify the contentious rescue impulse: Sarah Murnaghan is clearly facing a rescue situation and benefitting from her identifiability through media coverage. Further, the decision to not place Sarah on the adult transplant list would be a decision to prioritize other, statistical lives.³³
2. Cost-benefit analysis: Due to the scarcity of evidence on the effects of transplanting adult lungs into children, it is uncertain if Sarah’s inclusion on the transplant list would lead to greater or fewer years of life saved. It appears that transplant experts believed years of life saved would be greater for an adult transplant recipient than Sarah due to concerns about the efficacy of transplanting adult lungs into pediatric patients. I will proceed with this assumption.
3. Symbolic value: Sarah’s plight generated significant media and government attention.

However, this attention was generated by pandering to individual level psychological

³³ From the perspective of the public, Sarah was an identified life, and the nameless others on the organ transplant list were statistical lives. However, from the perspective of UNOS it was known who the others above Sarah are on the list—they were identified. This illustrates a complicated point – depending on the perspective adopted, it can be difficult to sort through the meaning of “statistical” and “identified.” I would argue that even to UNOS Sarah became identified relative to the others on the list (who were likely thought of in aggregate, i.e. as statistics). However, these are nuanced considerations. Greater clarity—and perhaps more gradations—is needed in delineating *statistical* from *identifiable* and determining the resulting implications for policy makers.

impulses such as reference group effects, identifiability, singularity, and psychological distance. As discussed in the psychology section, none of those effects are admissible in an institutional decision-making process, because they all play to either discriminatory impulses (psychological distancing) or errors in numerical thinking (reference group effect, identifiability, and singularity). Still, negative societal level trust and well-being effects could occur if Sarah were excluded from the list. Efforts to mollify these effects could be attempted; however, Sarah's case is a prime example of the fraught nature of such attempts. The United Network for Organ Sharing, the agency that manages organ transplantation lists, provides public information on their processes and welcomes public comment. Yet these measures had been insufficient to stop the media fervor around Sarah's case.

4. Discrimination: It is discriminatory to rule in favor of Sarah if the general rule (that pediatric patients should not be allowed on the adult transplant list) is considered medically valid. If that rule is maintained and yet Sarah is made an exception, it is an explicit instance of the well-positioned receiving preferential behavior based on their ability to garner public sympathy.
5. Practical concerns: UNOS, the DHS, and other relevant agencies should follow procedural justice in reviewing their rules around pediatric patient's eligibility for adult organ transplant. There are no other obvious practical concerns in this case.
6. Weighing the considerations: Overall, Sarah presents a complicated case. The symbolic value to saving her life was likely substantial. However, deciding to save her is not as effective from a life-saving perspective, nor is it life optimizing if adopted as a general rule. Moreover, to make an exception for Sarah would be explicitly discriminatory

because her ability to garner media attention was dependent on her social position and chance. Thus, I am inclined to recommend that Sarah's petition be denied.

Case Three—The Original Draft of the Oregon Health Plan

On May 2nd, 1990, the Oregon Health Services Commission (OHSC) unveiled their draft priority list of health services to be covered for Medicaid patients. The plan was purely based on cost-effectiveness: every health intervention was listed by cost per quality-adjusted life year (QALY) saved (a measure of both the number of years of life saved and the quality of that year of life) (Hadorn, 1991). The result was highly counterintuitive to many, with “surgical treatment for ectopic pregnancy and appendicitis rated just below, or as less important than, dental caps...and splints for temporomandibular joint disorder” (Hadorn, 1991, p. 2219). Thus clearly life saving, and nearly 100% effective, procedures were rated as lower priority than relatively minor conditions that are far cheaper to treat. Is this allocation scheme justifiable?

1. Identify the contentious rescue impulse: There is a strong intuitive rescue impulse to assist individuals in need of surgery for conditions like appendicitis or ectopic pregnancy, whereas the persons in need of dental caps are not rescue victims but rather persons predicted to face life-altering consequences in terms of quality of life.
2. Cost-benefit analysis: QALYs are maximized in the case of the cost-effectiveness allocation scheme. It is difficult to say how much more they are maximized in this scheme relative to a scheme that gives preference to life-saving procedures; for sake of argument, let us say that it is to a fairly significant degree (~10-20% improvement).
3. Symbolic value: The negative symbolic value of an entire health coverage plan prioritizing quality of life concerns over life saving, routine surgery is immense. It is hard to imagine the public understanding the denial of an appendectomy, a well-known and

life-saving surgery, so that they have improved dental coverage. This symbolic value is compounded by each life-saving procedure that is deprioritized.

4. Discrimination: There is no obvious discrimination in the allocation scheme.
5. Practical Concerns: As always, democratic procedural justice must be followed.

Here there is a case of significant QALY gains versus massive negative symbolic value. The amount of symbolic value and societal mistrust and confusion that would result from the cost-effectiveness plan outweighs the QALY improvement in this case. Thus, the Oregon case provides a general lesson: health coverage plans should always include some room for rescue impulses. Notice, however, that QALYs should not be discarded entirely in deference to life-saving interventions. The exact amount of room would be determined by weighing the point at which the relative value of decreasing QALYs saved equals the value of increasing symbolic value.

Case Four: EMTALA and National Policy

On April 7th, 1986 the Emergency Medical Treatment and Active Labor Act (EMTALA) was signed into law. EMTALA was written with the explicit goal of combatting patient “dumping”, a practice in which hospitals either refuse to care for nonpaying patients or transfer those patients to (usually community) hospitals without providing needed stabilization. It requires that all hospitals with an emergency room conduct a medical screening exam to determine if an emergency medical condition is present, provide appropriate care to any patient with such a condition, and, if patient transfer is determined to be in the patients’ best interest, to conduct an appropriate transfer of care. EMTALA was therefore intended to ensure that all patients, regardless of means, ethnicity, gender, or any other considerations, would have access to emergency medical care. Notably, it is also an unfunded mandate—meaning that hospital

systems were left to their own devices to develop systems and allocate funds towards compliance.

The law was passed for various reasons, including public pressure, financial incentives, and concern that the expiration of past and implementation of new policies could lead to gaps for vulnerable populations. Notably, there was little research into patient dumping prior to the public backlash about its occurrence and the subsequent passing of EMTALA. Rather, anecdotes, op-eds, and news coverage, with the support of a few studies, was sufficient to stir up public sentiment against dumping practices (Friedman, 1982; Smith, 1984; Dowell, 1984). The majority of the evidence, and the anecdotes, came from just one community hospital system: Parkland Memorial Health and Hospital System in Dallas, Texas. Researchers there kept records of 1,897 patient transfers between October 1982 and September 1984. Over 500 patients arrived at Parkland without any advanced notification of the transfer (Reed, 1986). Recordings of physicians requesting transfers to Parkland included quotes declaring “this is a private, capitalist, money-making hospital. They’re on my back to have her transferred” and “Honey, we're not talking about ethical practice. We're talking about a lady that needs something done that doesn't have the money to do it with...that’s what your damned hospital is there for” (Friedman, 2011). News coverage of the Parkland findings in the popular press led to a 60 Minutes episode on March 17, 1985 titled “the Billfold Biopsy” (Goldin, 1985). EMTALA-like laws were passed in numerous states shortly thereafter, and then the national law followed in April of 1986.

With this as background, we can now ask: should EMTALA exist?

1. Identify the contentious rescue impulse: EMTALA was explicitly motivated by and passed to ensure that necessary medical rescues would be provided to all individuals. Since EMTALA is an unfunded mandate, the alternative to EMTALA would simply

- be that medical systems are left to their own devices to determine how to allocate their resources between emergent patients and other aspects of healthcare.
2. Cost-benefit analysis: There is disagreement in the medical and public health community as to whether EMTALA has improved emergency medical care. Some argue that EMTALA provided a much needed financial reprieve for community hospitals, who previously bore the brunt of the burden in caring for poor emergency patients, and that it also improved the quality of processes in emergency rooms (Lee, 2004). Others argue that complying with EMTALA from a documentation and staffing perspective was a greater drain on resources than it was a benefit, and further that public knowledge of EMTALA requirements has contributed to unnecessary and costly overcrowding of emergency rooms (Wanerman, 2002; Derlet and Richards, 2000). Given these opposing arguments, I will consider this point about even on either side of the debate.
 3. Symbolic value: There is great symbolic value to guaranteeing evaluation for any person, and to appropriate care if an emergency is identified. The positive publicity of the law, and the public outcry that led to its adoption in states and then the national government, are sufficient evidence of this positive symbolic value.
 4. Discrimination: EMTALA is anti-discriminatory, in terms of providing greater assistance to the social disadvantaged. It is likely that before it was passed minority individuals in particular were more likely to either not be offered a medical screening or receive an inadequate screening. EMTALA would not fully change the latter case, but it was undeniably a positive step forward.

5. Practical considerations: Democratic procedural justice would need to be observed.

The number of additional externalities in EMTALA's case are enormously complicated, ranging from inadequate research on the problem of patient dumping to considerations about its impact on other policies like the Hill Burton Act and the (at the time) new Medicare Diagnosis Related Groups. I will consider these overall wide-ranging implications to be neutral.

The passage of EMTALA is therefore relatively neutral from a QALYs perspective, substantially positive in terms of symbolic value, and positive in terms of decreasing discrimination. Thus, its passage was morally justifiable.

3. Conclusion and the Need for Further Work

Among the most powerful moral impulses is the instinct to rescue someone in imminent peril for whom help is possible. However, limited resources and discrimination for or against certain groups provide reasons to limit the instances in which that impulse is followed. In this paper, I aimed to provide the first (reasonably) comprehensive theoretical treatment of these issues within the context of healthcare allocation.

My method was to adopt the perspective of the default institutional decision-maker, which I assume to be a largely consequentialist framework that features deontological side constraints. I outlined the psychological literature around rescue impulses and argued that psychological factors appear unconvincing in and of themselves in providing moral reasons to defend rescue impulses that trend towards non-life optimizing policies. Rather, the societal value of rescue and the deontological command to respect all persons as equal provide a rationale for preferring the rescue of identifiable victims over a greater number of statistical victims, but only in certain cases.

In the end, there is reasonable disagreement in determining the relative importance as well as the objective values associated with maximizing lives saved, the societal value placed on rescue, and the potential for unfair discrimination. I proposed a framework for considering these questions within a diverse and transparent deliberative space, and I provided four case studies as examples of how the framework could be used in practice. Hopefully, this framework is helpful in facilitating conversation and enabling deliberative bodies to arrive at morally sensitive solutions. I am also optimistic that the points of the framework, and the variety of specific features that contribute to each point, can be refined greatly through future use and critiques.

There are various logistical and theoretical limitations to the arguments covered in this paper. Logistically, hospitals and health systems already have many regulatory hurdles and committee structures in place when making allocation decisions. Adding rescue allocation taskforces could be seen as an additional burden. Further, the rationing decisions covered only apply to situations in which funds could be freely allocated between rescue versus statistical decisions. For example, in the case of Coby Howard, the child with leukemia, money that would be directed to his case as an exception to the previously established Oregon Medicaid process would be taken out of the general insurance pool, and thus taken from some other individual(s) who had been deemed at the outset to be more cost-effective to help. However, my paper has not discussed situations with restricted funds in which there is no freedom to re-allocate. Indeed, if funds could be provided for rescue without detracting from help to others (in the case of Coby Howard, perhaps by crowdsourcing his cause, or by his local community agreeing to a temporary tax for his benefit, etc.) then my arguments would not hold. This leads into a more general limitation of my arguments: that they primarily hold for public rescue, or for institutions that have a public service mission. I purposefully did not include the ethical obligations of private

rescue by citing the ambiguity of the scope and force of rescue obligations for private individuals and thus for private entities writ large. Additional work on the ethics of private rescue and the potential obligations imposed by the RR on private individuals and entities is therefore needed.

There is also a significant amount of additional discussion needed in order to better understand the applications of hRR in the institutional setting. Of the remaining areas requiring further study, three seem especially important following my general treatment of institutional hRR:

1. Psychological research into the symbolic value of rescue. Such research would be immensely difficult, but not impossible. Opportunistic studies could quantify and longitudinally follow well-being and group solidarity shortly after publicized rescues. Trade-off schemes could ask study participants how they feel towards various allocation programs, and could vary the degree of information provided about the programs to investigate the impact of information on symbolic value. Studies should also attempt to elucidate the relative impact of rescue decisions on each of the three components of symbolic value: societal well-being, medical trust, and increased altruism
2. Applications of the checklist from various institutional perspectives and for various recurrent scenarios. For example, one could elucidate applications of the checklist from the perspective of insurers versus hospitals versus healthcare NGOs. Furthermore, it would be beneficial to discuss applications for commonly encountered allocation scenarios, for example, specific checklist considerations for allocation of ventilators in flu pandemics. Moreover, published accounts of the checklist in real use would be invaluable for identifying practical issues and nuancing the various considerations.

3. Further ethical responses and critiques to the theoretical arguments I have presented.

Philosophical enquiries only progress through critique, re-evaluation, novel proposals, and synthesis. Perhaps another thinker can provide principled reasoning in favor of a more prominent role for distributive justice in an account of the Rule of Rescue, and thereby provide reasons for adjusting the checklist. If so, this would be significant.

I am hopeful that with additional work, the remaining issues around the Rule of Rescue and the identifiable victim bias can slowly be resolved. Only through continued research, theorizing, and the practical (if sometimes painful) gains derived from real-world decision making will we advance ever closer to morally justifiable solutions in the healthcare arena.

Acknowledgments

I would like to thank a number of individuals for their support, advice, and consistency in assisting me with this project. Dr. Philip Rosoff provided invaluable insight and nuance to the paper, and challenged me to strive for both clarity and subtlety in my analysis. Dr. Jennifer Hawkins was a consistent sounding board for my arguments, and supplied essential guidance on how to approach a long-term research project. Michelle Griffin and Emily Cohen were diligent copy editors. To Jenny Jiao fell the vital service of telling me when I had a bad idea or was failing to convey my meaning. Finally, my parents Brent and Jodi Flynn provided ongoing support and tolerated my musings throughout the project. To all of these individuals, I owe a great deal. Thank you.

References

- Adler, P. S., & Kwon, S.-W. (2002). Social Capital: Prospects for a New Concept. *Academy of Management Review*, 27(1), 17–40. <https://doi.org/10.5465/amr.2002.5922314>
- Adler, M.D. (2015). Welfarism, Equity, and the Choice between Statistical and Identified Victims. In G. Cohen, N. Daniels, & N. Eyal (Eds.), *Identified Versus Statistical Lives: An Interdisciplinary Perspective*. 53-77. Oxford University Press.
- American Medical Association Code of Medical Ethics. Principle VI. (2016). Accessed at <https://www.ama-assn.org/sites/ama-assn.org/files/corp/media-browser/principles-of-medical-ethics.pdf>
- Beach, M. C., Meredith, L. S., Halpern, J., Wells, K. B., & Ford, D. E. (2005). Physician Conceptions of Responsibility to Individual Patients and Distributive Justice in Health Care. *The Annals of Family Medicine*, 3(1), 53–59. <https://doi.org/10.1370/afm.257>
- Bochner, F., Martine, E. D., Burgess, N.G., Somogyi, A. A., & Garry, M. H. (1994). Controversies in treatment: How can hospitals ration drugs? *British Medical Journal*, 308 (6933). 901-905, 907-908.
- Brock, D. (2015). “Identified Versus Statistical Lives: Some Introductory Issues and Arguments.” In G. Cohen, N. Daniels, & N. Eyal (Eds.), *Identified Versus Statistical Lives: An Interdisciplinary Perspective*. 53-77. Oxford University Press.
- Child Starvation in Africa. (n.d.). Save the Children Foundation. Retrieved December 9, 2018, from <https://www.savethechildren.org/us/what-we-do/emergency-response/helping-starving-african-children>
- Cookson, R., McCabe, C. & Tsuchiya, A. (2008). Public Healthcare Resource Allocation and the Rule of Rescue. *Journal of Medical Ethics* 34: 540-544.
- Creadon, M. (1997). The Ocean Denied. *Time*, Vol. 149 (3), January 20.
- Daniels, N., Sabin, J. (1997). Limits to Health Care: Fair Procedures, Democratic Deliberation, and the Legitimacy Problem for Insurers. *Philosophy & Public Affairs* 26, 303–350. <https://doi.org/10.1111/j.1088-4963.1997.tb00082.x>
- Daniels, N. (2000). Accountability for reasonableness. *The BMJ* 321, 1300–1301.
- Daniels, Norman. (2012). Reasonable Disagreement about Identified vs. Statistical Victims. *Hastings Center Report* 42: 35-45.
- Derlet, R. W., & Richards, J. R. (2000). Overcrowding in the nation’s emergency departments: Complex causes and disturbing effects. *Annals of Emergency Medicine*, 35(1), 63–68. [https://doi.org/10.1016/S0196-0644\(00\)70105-3](https://doi.org/10.1016/S0196-0644(00)70105-3)

- DeSante, J., Caplan, A., Hippen, B., Testa, G., & Lantos, J.D., (2014). Was Sarah Murnaghan Treated Justly? *Pediatrics* 134, 155–162. <https://doi.org/10.1542/peds.2013-4189>
- Dowell, M.A. (1984). Indigent access to hospital emergency room services. *Clearinghouse Review* 18:483-499.
- Draft: North Carolina Triage Protocol for Allocation of Scarce and Critical Care Resources During an Influenza Pandemic. (2010). Received in correspondence with Dr. Phillip Rosoff; used by permission of the same.
- Eddy D.M. (1990) Clinical decision making: from theory to practice. Connecting value and costs. Whom do we ask, and what do we ask them? *JAMA : the journal of the American Medical Association* 264(13):1737-1739.
- Evans, R. W. (2013). The Sarah Murnaghan debacle: A health policy perspective on transplant candidate selection. *The Journal of Heart and Lung Transplantation*, 32(9), 868–870. <https://doi.org/10.1016/j.healun.2013.07.009>
- Eyal, N. (2015). Concentrated Risk, the Coventry Blitz, Chamberlain’s Cancer. In G. Cohen, N. Daniels, & N. Eyal (Eds.), *Identified Versus Statistical Lives: An Interdisciplinary Perspective*. 53-77. Oxford University Press.
- Frick, J. (2015). Treatment versus Prevention in the Fight against HIV/AIDS and the Problem of Identified versus Statistical Lives. In G. Cohen, N. Daniels, & N. Eyal (Eds.), *Identified Versus Statistical Lives: An Interdisciplinary Perspective*. 53-77. Oxford University Press.
- Fried, C. (1969). The value of life. *Harvard Law Review* 82(7), 1415–1437.
- Friedman, E. (1982). The 'dumping' dilemma: The poor are always with some of us. *Hospitals* 56(Sept 1):51-56. 3.
- Friedman, E. (2011). How EMTALA Transformed Health Care - Hospitals and Health Networks. *Hospitals and Health Networks*. Retrieved from <https://www.hhnmag.com/articles/5010-the-law-that-changed-everything-and-it-isn-t-the-one-you-think>
- Gibbard, A. (1986). “Risk and Value.” in D. Maclean (Ed.) *Values at Risk*. p. 94-112
- Goldin, M. (1985). Billfold Biopsy. 60 Minutes. CBS News.
- Grosse, Scott (2008) Assessing cost-effectiveness in healthcare: history of the \$50,000 per QALY threshold, Expert Review of Pharmacoeconomics & Outcomes Research, 8:2, 165-178, DOI: [10.1586/14737167.8.2.165](https://doi.org/10.1586/14737167.8.2.165)
- Hadorn, David. (1991). Setting Health Care Priorities in Oregon: Cost-Effectiveness Meets the Rule of Rescue. *Journal of the American Medical Association*, 265(17) 2218-2228.
- Hare, C. (2012). Obligations to Merely Statistical People. *The Journal of Philosophy*, 109(5/6), 378–390. Retrieved from JSTOR.

- Harris J. (1985). *The Value of Life. An Introduction to Medical Ethics*. London: Routledge & Kegan Paul; 1985. Chapter 5.
- Hawkins J. (2006). Justice and Placebo Controls. *Social Theory and Practice* 32(3): 467-96.
- Higgins, R. S. D., & Fishman, J. A. (2006). Disparities in Solid Organ Transplantation for Ethnic Minorities: Facts and Solutions. *American Journal of Transplantation*, 6(11), 2556–2562. <https://doi.org/10.1111/j.1600-6143.2006.01514.x>
- Honeybul, S., Gillett, G. R., Ho, K. M., & Lind, C. R. P. (2011). Neurotrauma and the rule of rescue. *Journal of Medical Ethics*, 37(12), 707–710. <https://doi.org/10.1136/medethics-2011-100081>
- Jenni, K., & Loewenstein, G. (1997). Explaining the Identifiable Victim Effect. *Journal of Risk and Uncertainty*, 14(3), 235–257. <https://doi.org/10.1023/A:1007740225484>
- Jones, B., Rachlin, H. (2006). "Social Discounting". *Psychological Science*. 17, 4.
- Jonsen, A. (1986). Bentham in a box: Technology Assessment and Health Care Allocation. *Law Med Health Care*; 14(3-4): 172-174
- Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica* 47(2) 263-91.
- Kahneman, D., & Tversky, A. (2000). *Choices, Values, and Frames*. Russel Sage Foundation, New York.
- Kahneman, D., Frederick, S. (2002). Representativeness Revisited: Attribute Substitution in Intuitive Judgment, in: Gilovich, T., Griffin, D., Kahneman, D. (Eds.), *Heuristics and Biases*. Cambridge University Press, pp. 49–81. <https://doi.org/10.1017/CBO9780511808098.004>
- Kamins, M.A., Gupta, K. (1994). Congruence between spokesperson and product type: A matchup hypothesis perspective. *Psychology & Marketing*, 11, 569–586. <https://doi.org/10.1002/mar.4220110605>
- Kamm, F.M. (2000). Does distance matter morally to the duty to rescue? *Law and Philosophy*, 19: 655-681.
- Kant, I., & Gregor, M. J. (1998; original 1785). *Groundwork of the metaphysics of morals*. Cambridge, U.K: Cambridge University Press.
- Kogut, T., and I., Rotov. (2005a). “The ‘Identified Victim’ Effect: An Identified Group or just an Individual? *Journal of Behavioral Decision Making* 18 (3): 157-67.
- Kogut, T., and I., Rotov. (2005b). “The Singularity Effect of Identified Victims in Separate and Joint Evaluation.” *Organizational Behavior and Human Decision Processes* 97 (2): 106-16.

- Kohn, R., Rubenfeld, G. D., Levy, M. M., Ubel, P. A., & Halpern, S. D. (2011). Rule of rescue or the good of the many? An analysis of physicians' and nurses' preferences for allocating ICU beds. *Intensive Care Medicine*, 37(7), 1210–1217. <https://doi.org/10.1007/s00134-011-2257-6>
- Latané, B., and Darley, J. (1970). *The Unresponsive Bystander: Why doesn't he help?* New York: Appleton-Century-Crofts
- Lee, T. M. (2004). An EMTALA Primer: The Impact of Changes in the Emergency Medicine Landscape on EMTALA Compliance and Enforcement. *Annals of Health Law*, 13, 145–178.
- Lübbe, W. (2019). Appeal to the Rule of Rescue in health care: discriminating and not benevolent? *Medicine, Health Care and Philosophy*, 22(1), 53–58. <https://doi.org/10.1007/s11019-018-9839-9>
- Maclean, D. (1986). Social values and the distribution of risk. In D. Maclean (Ed.) *Values at risk*. New Jersey: Rowman and Allanheld.
- McKie, J., Richardson, J. (2003). The Rule of Rescue. *Social Science & Medicine* 56 (2003). 2407-2419.
- Mencius, XIII. (1983). in *Si Shu Zhang Zhu Ji Zhu*, ed. Zhu-xi. 359-60.
- Mengzi: With Selections from Traditional Commentaries. (2008). translated by Bryan Van Norden. Hackett Publishing.
- Merritt M.W., Taylor H.A., Murray L.C. (2010). Ancillary care in community-based public health intervention research. *American Journal of Public Health*, 100(2) 211-16.
- Miller, F.G., Mello, M.M., Joffe S. (2008). Incidental findings in human subject research: what do investigators owe research participants? *Journal of Law, Medicine, and Ethics*, 36(2) 271-79.
- Moylan C.A., Brady C.W., Johnson J.L., Smith A.D., Tuttle-Newhall J.E., Muir A.J. (2008). Disparities in liver transplantation before and after introduction of the MELD score. *JAMA*, 300(20):2371-8.
- National Institute for Health and Clinical Excellence. (2006). *Citizen's Council Report: the Rule of Rescue*. NICE. London.
- New York State Workgroup on Ventilator Allocation in an Influenza Pandemic. Allocation of Ventilators in an Influenza Pandemic: Planning Document. (2007). http://www.health.state.ny.us/diseases/communicable/influenza/pandemic/ventilators/docs/ventilator_guidance.pdf. Accessed August 24, 2007.
- Nord. E., Richardson, J., Street, A., Kuhse, H., & Singer, P. (1995). Maximizing health benefits vs egalitarianism: An Australian survey of health issues. *Social Science & Medicine*, 41(10), 1429-1437

- Osborne, M., & Evans, T. W. (1994). Allocations of resources in intensive care: a transatlantic perspective. *The Lancet*, 343. 778-780.
- PBAC Guidelines | 5.4 Basis for any claim for the 'rule of rescue.' (2016). Australian Government Department of Health: The Pharmaceutical Benefits Advisory Committee Guidelines. <https://pbac.pbs.gov.au/section-5/5-4-basis-for-any-claim-for-the-rule-of-rescue.html>
- Peters, D.A. (1986.) Rationales for organ donation: Charity or duty? *Journal of Medical Humanities* 7, 106–121. <https://doi.org/10.1007/BF01117902>
- Piliavin, Jane A., et al. (1981). *Emergency Intervention*. New York: Academic Press.
- Rachlin, H., Jones, B.A. (2008). Altruism among relatives and non-relatives. *Behavioural Processes* 79, 120–123. <https://doi.org/10.1016/j.beproc.2008.06.002>
- Railton, P. (2015). “Dual-Process” Models of the Mind and the “Identifiable Victim Effect.” In G. Cohen, N. Daniels, & N. Eyal (Eds.), *Identified Versus Statistical Lives: An Interdisciplinary Perspective*. 53-77. Oxford University Press.
- Rawls, J. (1971). *A theory of justice*. Oxford: Oxford University Press.
- Reed., W., Cawley, C., & Anderson, R. (1986). Special Report: The Effect of a Public Hospital's Transfer Policy on Patient Care, 315 *New England Journal of Medicine*. 1428, 1431.
- Reid, L. (2005). Diminishing returns? Risk and the duty to care in the SARS epidemic. *Bioethics*, 9(4):348-361.
- Rosoff, P.M. (2012). Unpredictable drug shortages: an ethical framework for short-term rationing in hospitals. *American Journal of Bioethics*, 12(1):1-9;
- Rosoff P.M., Patel K.R., Scates, A., Rhea, G., Bush, P.W., & Govert, J.A. Coping with critical drug shortages: an ethical approach for allocating scarce resources in hospitals. *Archives of Internal Medicine*. 2012;172(19):1494-9
- Rosoff, P. (2017). *Drawing the Line: Healthcare Rationing and the Cutoff Problem*. Oxford University Press.
- Ruderman, C., Tracy, C.S., Bensimon, C.M., et al. (2006). On pandemics and the duty to care: whose duty? who cares? *BMC Medical Ethics*, 7(1):5.
- Rulli T, Millum J. (2016). Rescuing the duty to rescue. *Journal of Medical Ethics* 2016; 42:260-264.
- Scanlon, T. (1998). *What We Owe to Each Other*. Cambridge, MA: Belknap Press of Harvard University.
- Schelling, T.C. (1968). The life you save may be your own. In: Brookings I, Chase SB, eds. *Problems in public expenditure analysis: Papers presented at a conference of experts held Sept. 15-16, 1966*. Washington: Brookings Institution, 127-162.

- Singer, P. (1972). Famine, affluence, and morality. *Philosophy and Public Affairs*, 1(3), 229-243.
- 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. <https://doi.org/10.1023/A:1022299422219>
- 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. <https://doi.org/10.1016/j.obhdp.2006.01.005>
- Small, D.A. (2010). “Reference-Dependent Sympathy.” *Organizational Behavior and Human Decision Processes* 112 (2): 151-60.
- Small, D.A. (2015). On the Psychology of the Identifiable Victim Effect. In G. Cohen, N. Daniels, & N. Eyal (Eds.), *Identified Versus Statistical Lives: An Interdisciplinary Perspective*. 53-77. Oxford University Press.
- Small D.A., & Simonsohn, U. (2008). “Friends of Victims: The Impact of Personal Relationships with Victims on Generosity toward Others.” *Journal of Consumer Research*, 35: 532-42.
- Smith, P.C. (1984). Public hospitals fear costly 'dumping' of DRG 'losers.' *Hosp Tribune*, p. 11, Nov 21st.
- Smith, R., Faro, D., and Burson, M. (2011) , "Flocks, Herds, and Families: the Influence of Victim-Unitization on Charitable Giving", in NA - Advances in Consumer Research Volume 38, eds. Darren W. Dahl, Gita V. Johar, and Stijn M.J. van Osselaer, Duluth, MN : Association for Consumer Research.
- Squires, D. A. (2012). Explaining high health care spending in the United States: an international comparison of supply, utilization, prices, and quality. *Issue Brief (Commonwealth Fund)*, 10, 1–14.
- Wanerman, R. (2002). The EMTALA Paradox. *Annals of Emergency Medicine*, 40(5), 464–469. <https://doi.org/10.1067/mem.2002.126743>