

SABA BAZARGAN-FORWARD 

VESTING AGENT-RELATIVE PERMISSIONS IN A PROXY*

(Accepted 19 September 2018)

ABSTRACT. We all have agent relative permissions to give extra weight to our own well being. If you and two strangers are drowning, and you can save either yourself or two strangers, you have an agent relative permission to save yourself. But is it possible for you to ‘vest’ your agent relative permissions in a third party a ‘proxy’ who can enact your agent centered permissions on your behalf, thereby permitting her to do what would otherwise be impermissible? Some might think that the answer is ‘no’; it is definitive of agent centered permissions that they apply only to the individuals ineliminably referenced in the content of that reason, which means that they lack reason giving force for any ostensible proxy. The purpose of this paper is to (i) show that vesting agent relative permissions is indeed possible, (ii) provide an account of how agent relative permissions are vested by considering the structure of rights more generally, and (iii) show that we have a right to vest such permission in this way.

I. INTRODUCTION

You and two strangers are drowning. You can save yourself or the two strangers but not all three. Suppose that you have an agent-relative permission to save yourself. All things being equal, though, a third party in a position to save either you or the two strangers would not be permitted to save you. After all, your agent-relative permissions apply solely to you. But is it possible for you to ‘vest’ your agent-relative permissions in a third party – a ‘proxy’ – who can thereby enact your agent-relative permissions on your behalf?

Some might think that the answer is ‘yes’; the sheer importance of agent-relative permissions suggests that we should be able to delegate them (as we do many other rights) to third parties in cases

* I’d like to thank Victor Tadros and Helen Frowe for invaluable criticism of an earlier version of this paper presented at a 2017 workshop on defensive violence in the University of Warwick.

where we are unable to exercise them in our own person.¹ Others might think that the answer is ‘no’; it is definitive of agent-relative permissions that they apply only to the individuals ineliminably referenced in the content of that reason,² which means that they lack reason-giving force for any ostensible proxy.

I argue that it is indeed possible to vest your agent-relative permissions in a proxy, thereby permitting the proxy to do on your behalf what would otherwise be impermissible for the proxy to do. The sense in which agent-relative permissions are ineliminably first-personal is compatible with the possibility of vesting such permissions in a proxy. In arguing for this view, I present (in Section IV) a formal account of how agent-relative permissions are vested. I then argue (in Section V) that vesting such permission is permissible, in that doing so does not in and of itself violate the rights of, or treat unfairly, the individuals whose wellbeing the proxy is thereby permitted to forego. I will conclude by very briefly considering the practical relevance of this account. I begin, though, (in Section II) with a discussion of agent-relative permissions as such, followed by an overview (in Section III) of what it means to vest such permissions in a proxy.

II. A PRIMER ON AGENT RELATIVE PERMISSIONS

If agent-neutral reasons, such as the reason to promote overall wellbeing, are left unchecked, they will overwhelm our lives by forcing us to abandon our sub-optimific personal projects. Some of these personal projects are ground-level, long-term commitments constitutive of our practical identity. An account of morality forcing us to abandon sub-optimific personal projects is neither psychologically plausible nor normatively desirable.

To address this worry, some have argued that agent-neutral reasons are defeasible. Each of us has agent-relative permissions to give extra weight to our own legitimate interests in our moral calculus. This gives us the moral space to pursue a range of impersonally sub-optimific projects. There are, at the broadest level, two

¹ Cecile Fabre makes this point in Fabre (2012, pp. 201–202). Others who argue that it is possible to vest an agent-relative permission in a third party include Quong (2016, pp. 817–818) and Lazar (2016, pp. 219–223). See also Bazargan-Forward (2018).

² Nancy Ann Davis famously makes this point in Davis (1984, pp. 192–193).

types of argument in favor of agent-relative permissions: derivative and non-derivative arguments.

According to derivative arguments for agent-relative permissions, consequentialist reasons to make things go impersonally best ground agent-relative permissions to give special weight to personal projects. On this view, consequentialism is partially self-effacing; we are permitted to act *as if* there is no absolute requirement to maximize the impersonal good because deliberating from that standpoint is more likely to maximize the impersonal good. Peter Railton is best known for this kind of argument for agent-relative permissions.³

According to the most well-known *non*-derivative arguments for agent-relative permissions, our defeasible moral permission to pursue impersonally sub-optimific personal projects is grounded either (i) deontically in the fact the project is one's own, or (ii) axiologically in the intrinsic personal value manifest in the permitted project. Samuel Scheffler is the most well-known proponent of the first view;⁴ Thomas Nagel is the most well-known proponent of the second.⁵

On all of these characterizations, agent-relative permissions allow us to give extra weight to our own legitimate interests in our moral calculus. There is room for disagreement, of course, regarding which interests count as legitimate, how much extra weight the agent-relative permission confers, and whether this extra weight depends on the kind of legitimate interest in question. But since my focus is on the possibility of vesting agent-relative permissions *per se*, I will prescind from those complications by arbitrarily stipulating that each of us has an agent-relative permission to grant our own legitimate interests *five times* the weight which would be permitted absent that permission. I do this to keep the discussion tractable.

There is also disagreement on whether the agent-relative permission is sensitive to the *mode of agency* by which that permission is exercised. In particular, the amount of extra weight you are permitted to give your own wellbeing might be sensitive to whether the conduct under consideration is an instance of causing a harm or merely allowing a harm. To avoid unnecessary controversy, I will

³ Railton (1984).

⁴ Nagel (1989).

⁵ Scheffler (1982).

limit myself to the claim that we have agent-relative permissions to forego doing the greater good, rather than to do more harm.⁶

I will assume, then, the following: when deciding between saving your own life or the lives of two strangers, you have an agent-relative permission to save your own. Is it possible to vest this permission in a third-party who is thereby permitted to vicariously exercise your agent-relative permissions on your behalf? To answer this, we need a better grasp of what it means to ‘vest’ an agent-relative permission.

III. THE CONCEPT OF VESTING AGENT RELATIVE PERMISSIONS

A. A First Gloss

Typically, agent-relative permissions protect conduct performed by the person who harbors the agent-relative permission in question. So if *I* have an agent-relative permission to do ϕ , then that permission allows *me* to do ϕ . In that way, my agent-relative permission protects my conduct. But when I vest an agent-relative permission in you, my agent-relative permission protects your conduct.

Crucially, vesting an agent-relative permission is not the same as *transferring* that permission. Suppose I vest in a proxy my agent-relative permission to give my own life extra weight. I do not thereby ‘lose’ the right to directly exercise that agent-relative permission. I am still entitled to act in furtherance of it, in my own person. Rather than transferring agent-relative permissions, vesting *extends* them.

In Section IV.A I provide a more thoroughgoing and rigorous analysis of how agent-relative permissions are vested; in doing so, I argue that such vesting does not violate any formal constraints on agent-relative reasons as such. For now, my purpose is only to present a rough grasp of what it means to vest an agent-relative permission. To that end, it’s helpful to distinguish between a) vesting an agent-relative permission in a proxy, and b) an agreement in which a third party promises to enact an agent-relative permission. The latter possibility is uncontroversial. Suppose you have an agent-relative permission to do ϕ , and I promise to do ϕ for you. I have

⁶ Those who argue that we can have agent-relative permissions to impose harm include Frowe (2008) (who has since changed her view), Quong (2009), and Lazar (2013). Skeptics include Lefkowitz (2009) and Tadros (2011, pp. 202–208).

thereby promised to do what you have an agent-relative permission to do. But when I promise to do Φ for you, your agent-relative permission to do Φ does not apply to me. The promise I make gives me a special reason to do Φ ; but your agent-relative permission does not undergird that special reason. Put differently, it's the promise rather than your agent-relative permission that 'does the work' in requiring me to do Φ . This is important because our strongest agent-relative permissions permit conduct which no promise can warrant. To see this, consider the following case.

Suppose I promise to save your life should it ever be endangered. You are otherwise a stranger to me. Since you have an agent-relative permission to give extra weight to your own life, I have thereby promised to do what you have an agent-relative permission to do. Now suppose that I chance upon you drowning; two others are drowning as well. I can easily save either you or the two others but not all three. I presumably have no permission, let alone a duty, to save you given the opportunity costs of doing so. Those costs just describe a *defeater* of the promise I made to you. You, however, would be permitted to save your own life (if you could) even at the opportunity cost of allowing two others to drown. So if you had vested those agent-relative permissions in me, then I would indeed be permitted to save your life given the opportunity costs. This suggests that an agent-relative permission to do Φ , when vested in a proxy, permits what promising to do Φ does not. We cannot, then, analyze vesting by simply recasting it as promising. These are related but importantly different practices.

In Section IV.B I further distinguish vesting an agent-relative permission from promising to enact an agent-relative permission. But in what immediately follows I consider and respond to arguments purporting to show that it is conceptually impossible for agent-relative permissions to have reason-giving force for the proxy in whom such reasons are ostensibly vested.

B. Conceptual Challenges to Vesting

At first, the possibility of vesting agent-relative permissions might seem unproblematic. Suppose the only way for me to prevent myself from drowning is by pulling a lever that will save you, though at the

opportunity cost of failing to save five others. And the only way for you to prevent yourself from drowning is by pulling a lever which will save me, though at the opportunity cost of failing to save five others. Suppose each of us is permitted to save him or herself instead of five others (but no more). Surely, then, we can ‘trade tasks’ – I save you and you save me.⁷ This seems to show that we can indeed delegate agent-relative permissions to others.

To diffuse the intuition that ‘trading tasks’ in this way is permissible, a skeptic of vesting can simply turn this case on its head. If my act of saving you is counterfactually dependent on your act of saving me, and *vice versa*, then *ten* lives go unsaved when we ‘trade tasks’.⁸ So the only way for you to ensure that the lives of ten innocents are saved is by refusing to save the stranger, who consequently refuses to save you. Since you’re permitted to save yourself at the opportunity cost of allowing no more than five others to die, it seems, then, that you’re required to forego saving yourself and the stranger. This effectively re-describes the ‘trading tasks’ case, but in a way eliciting a contrary intuition. The moral here is that we cannot rely on brute intuitions to determine whether it is possible to vest agent-relative permissions.

Any account which allows for the possibility of vesting agent-relative permissions in a proxy must confront a fundamental challenge: the proxy must be able to justified acting sub-optimifically to those whose welfare she foregoes. As discussed in Section II, I can invoke the agent-relative permission I have to give my own interests extra weight as a defeater for the presumptive agent-neutral reason I have to promote the impersonal good. But it seems that my proxy can make no such claim on my behalf when she must choose between my life and the lives of the strangers. From her perspective, the choice is not *her* life or the lives of the strangers. Rather, the choice is the life of one stranger or the lives of several strangers. What consideration, then, defeats the presumptive reason the proxy has to save the greater number?

⁷ Lazar uses variations of this example in Lazar (2016).

⁸ I thank an anonymous referee for pressing me on this. Of course, we might deny that all ten deaths go in each of our respective moral ledgers, given the moral relevance of intervening agency. But this is difficult to accept given that each of us foregoes saving five lives specifically at, and because of, the other’s request.

There might seem to be none. Agent-relative permissions are, after all, agent-relative. For the sake of simplicity, suppose that agent-relative reasons are reasons ineliminably referencing the agent to whom the reason applies (whether that reason is to act with respect to or in furtherance of a particular aim).⁹ So it might seem that the proxy cannot invoke my agent-relative permissions in justifying her decision to refrain from saving the greater number.

In an attempt to prove that point, one might mistakenly argue as follows. People cannot be reasonably expected to be neutral about their own welfare (or the welfare of their nearest and dearest). Such a requirement asks too much of us; any action-guiding moral theory cannot ignore the fact that for the vast majority of us, it is prohibitively difficult to act solely from the impersonal good. In that sense, most of us are to some degree unavoidably egoistic (or self-referentially altruistic).¹⁰ Agent-relative permissions accommodate that fact – or so it might be thought. Since it is overly onerous to require of me that I remain neutral about my own welfare, I am accordingly permitted to save my own life even when doing so comes at the opportunity cost of failing to save several others. I can justify my decision to the strangers whose welfare I forego by pointing out that I am not required to act optimifically, since such a requirement demands a kind of neutrality that I cannot be expected to muster. This sort of justification, however, is not in the offing in a case where my proxy must choose between my life on the one hand, and several strangers on the other. Since she can be neutral in such a case, she is required to be neutral. The presumptive reason to act optimifically is undefeated. So it's impermissible in such a case for my proxy to enact my agent-relative permissions. Or so it might be argued.

But that argument rests on a mistaken view of what grounds agent-relative permissions. On that view, it's prohibitively difficult to be impartial about one's own welfare, so morality allows some personal partiality in the form of agent-relative permissions. But if

⁹ There are various competing accounts of agent-relative reasons. Some claim that agent-relative reasons are derived from a general rule or (universally quantified) principle: Nagel (1970, 1989), Parfit (1984), McNaughton and Rawling (1995). Some forego any reference to a general rule or principle: Pettit (1987). Some presuppose that agent-relative reasons are necessarily teleological (Nagel, 1970, 1989; Parfit, 1984). Others make no such presumption (Scanlon, 1998, pp. 79–107; Korsgaard, 1996, p. 300). These differences are largely immaterial to this discussion.

¹⁰ See Broad (1930, pp. 54–55).

that were true, then a saint capable of perfect impartiality would be *required* to act accordingly. All of her actions would be normatively determined solely by whatever so happens to maximize overall welfare. Inasmuch, she would be morally prohibited from living a life recognizably her own – i.e., flowing from commitments *she* adopts. Agent-relative permissions, though, are supposed to provide *everyone* with moral protection against the tyranny of impersonal reasons. The purpose of such a permission is not (or is not simply) to accommodate our egoistic natures, but to safeguard a particular kind of value – personal value. Our projects have a value derived from and applicable to only from the positional stance of the person undertaking the project in question. Any account of morality failing to capture this kind of value – i.e., any account of morality capturing only impersonal value – is impoverished.¹¹

So even where we are perfectly capable of acting neutrally in furtherance of an impersonally optimific outcome, we are still permitted to act in accordance with personal values which agent-relative permissions protect. The fact that my proxy is eminently capable of acting impartially does not itself undercut her permission to enact my sub-optimific agent-relative permissions.

We might accordingly claim that my proxy acts in furtherance of the personal value attaching to my life and projects. But that move fails. Recall that personal value, by definition, is reason-providing only for the person whose ground-level projects generate that value. The personal value inhering in my ground-level, long-term commitments provides me with a reason to prioritize my projects; it does not concomitantly provide my proxy – or anyone else – with such a reason. The upshot is that my proxy, in justifying the decision to save me at the opportunity cost of failing to save several others, cannot invoke the personal value inhering in my life and ground-projects. So it seems my proxy has no justifying reason for giving my life the extra weight that I am permitted to give it. This in turn suggests that she is not permitted to vicariously discharge my agent-relative permissions after all.

¹¹ Some cast doubt on the cogency of agent-relative value as such; see in particular Schroeder (2007). Those so inclined can recast the point I make here by stating it in terms of reasons rather than value: our projects yield reasons derived from and applicable to only the positional stance of the person undertaking the project. But for ease of exposition, I will continue to write in terms of agent-relative value.

But in what follows I argue that this argument fails; my proxy is indeed permitted to enact my agent-relative permissions even given that the agent-relative value inhering in my ground projects are not reason-providing for her. In explaining why, I describe in greater detail how agent-relative permissions are vested in a proxy.

IV. THE POSSIBILITY OF VESTING AGENT RELATIVE PERMISSIONS

A. *How Vesting Works*

How is it that my proxy is permitted to enact my agent-relative permissions even though the personal value inhering in my life has no reason-giving force for her? To answer this question, it is helpful to begin with a discussion of rights, since agent-relative permissions are a species of right.

There are, of course, competing accounts of what the function¹² and grounds¹³ of rights are. These disputes are largely orthogonal to the claims I make here. For ease of exposition, I will write in terms of will-based accounts of the function of rights, and status-based accounts of the justification of rights. But nothing substantive turns on this choice. What I say can be translated into the language of competing accounts without undercutting my principal thesis: that we vest agent-relative permissions in a proxy by exercising a right that we have against third parties that they treat the proxy in a certain way.

Rights are held against third parties. This is obvious when it comes to claim-rights. I have a claim-right to ϕ if and only if the person against whom the claim is held has a duty toward me that she ϕ . But it's also true of liberty-rights. I have a liberty-right to ϕ if and only if I have no duty not to ϕ , in which case I do not wrong anyone

¹² According to will-based accounts, the function of a right is to give the right-holder the authority to control the duties others have within a certain domain of activity. See for example Hart (1982), Steiner (1994), and Wellman (1995). According to interest-based accounts, the function of a right is to further the right-holder's interests. See Lyons (1994), MacCormick (1982), Raz (1986), and Kramer (2001). There are a host of other accounts which do not fit into these categories; they are too numerous to list here.

¹³ According to status-based accounts of rights, we possess self-regulative and reasons-responsive capacities yielding constraints on how others can treat us. Nozick (1974) is best known for revitalizing this account of rights. According to instrumentalist accounts, rights are those constraints which, when observed, best achieves an optimal distribution of advantages. Various versions of indirect consequentialism exemplify this view. On contractualist accounts, rights are the principles resulting that properly situated and motivated agents would choose or which they could not reasonably reject. See Rawls (1971, p. 298) and Scanlon (2003). Other accounts abound as well.

by doing ϕ . Accordingly, third parties cannot legitimately demand that I cease doing ϕ ; they lack the authority to enforce such a demand. Rights, then, protect the conduct to which the right applies by requiring certain kinds of treatment from third parties.

Rights have at least three elements: the demand (*D*) the right makes to certain treatment from third parties, the conduct (*C*) that the right protects, and the status (*S*) that grounds the right. Each of these relations – *D*, *C*, and *S* – will take potentially several relata. *D* will specify the third parties against whom the demand is made as well as the treatment required of them. *C* will specify the conduct the right protects. And *S* will specify the particular status grounding the right. But I will suppress those relata in my quasi-formal treatment in order to emphasize a particular relatum which all three relations specify: the person to whom *D*, *C*, and *S* applies. That is, each relation specifies in turn the person who has the demand, the person whose conduct the right protects, and the person whose status grounds the right. Normally, each of the three will take the *right-bearer* as the person to whom the relation applies. So when I have a right, it can typically be characterized like this:

$$R_{me} = D_{me}, C_{me}, S_{me}$$

This says that *my* right is a demand that *I* can make to certain kinds of treatment from third parties, protecting *my* conduct, grounded in *my* status. If the right in question were, for example, the right I have to destroy my property, then the treatment I could demand would be non-interference from those who have no competing claim to the property. The conduct the right protects would be my act of destroying the property. And the status grounding the right would be my status as a person possessing the sort of inviolability which the property-right protects.

This analysis of rights is hardly comprehensive. Its simplicity, though, helps throw light on what it means to vest a right in a third party. When I vest a right in you as my proxy, the conduct *my* right protects is *your* conduct. The right, though, is still mine in the following sense: *I* am entitled to demand certain kinds of treatment from third parties, and it is *my* status – not yours – which grounds the right. My vested right, then, looks like this:

$$R_{me} = D_{me}, C_{you}, S_{me}$$

This has several important implications. Suppose a third party wrongly prevents you, the proxy, from engaging in the sort of conduct which the right protects. The right, however, is grounded in my status. So by preventing you from engaging in the protected conduct, the third parties *eo ipso* wrongs *me*. Of course, it might be that what the third party does wrongs you as well, or any number of other people for that matter. The point, though, is that I am *necessarily* wronged, since it is my status grounding the right in question. Since the third party's treatment of you wrongs me, I can demand that the third party treat *you* differently.

We are now in a position to see why the irrevocably first-personal nature of agent-relative value poses no threat to the possibility of vesting agent-relative permissions in a proxy. The worry was that my proxy has a strong presumptive reason to act optimifically, thereby requiring that she rescue the greater number of strangers instead of me. My agent-relative permission to give my own life priority derives from the personal value that my life has as to me. The agent-relative value that my life has is not reason-providing for anyone other than me; *a fortiori*, it is not reason-providing for my proxy. Thus the proxy would be unable to justify saving me given the grossly sub-optimific opportunity costs of doing so.

But on the given analysis of what it means for me to vest a right, it's not my proxy who can demand of third parties that they refrain from interfering with her conduct – rather, it's me. Suppose I had no proxy; if I could rescue myself or the several strangers, no one could permissibly force me to do the latter at the cost of my life. I have a right to save myself where that right, grounded in my status as a person whose life has personal value, consists in a demand that others refrain from forcing me to save the several strangers. Now return to the case in which my proxy must choose between saving me and the two strangers. My right to save myself, grounded in my status as a person whose life has personal value, consists in a demand that others refrain from forcing *my proxy* to save the two strangers.

Notice that this is consistent with the claim that the personal value attaching to my life is not reason-providing for my proxy. What permits my proxy to save me is *not* the added value that my

life has in the proxy's moral calculus. No such added value is available to my proxy. Rather, what permits her to save me is this: I have a right against third parties that they refrain from interfering with the conduct which my right protects. That right defeats any presumptive claim that others have against the proxy that she save the two strangers. Put simply, forcing my proxy to save the lives of the two strangers would violate my right to save myself.

Still, one might argue that though others cannot legitimately demand of my proxy that she save the two strangers, this still leaves her free to decide what to do – and what she ought to do is act optimifically. As I pointed out, the personal value attaching to *my* life is not reason-providing for my proxy; she cannot bring them to bear in her moral calculus. At the same time, though, I have a right against third parties that they refrain from interfering with my proxy. So on the one hand, in deciding what to do, my proxy cannot bring to bear the personal value attaching to my life. On the other hand, I have a right that no one prevent my proxy from saving the lives of the three strangers. Where does this leave the proxy?

At first it might seem that the proxy has at best a *non-enforceable duty* to save the two strangers. To understand this position, consider the treatment Samaritan duties receive under standard version of Right-Libertarianism. On that view, no one can be legitimately forced to engage in Samaritan assistance (absent a contractual or compensatory obligation – in which case the duty would not count as Samaritan). Such Libertarians can consistently claim that in cases where assistance incurs little cost, we have a moral duty to provide such assistance; it's just that the duty is non-enforceable in the sense that no one can permissibly force you to do what morality requires you to do. Put differently, though you ought to provide assistance, the individuals in need of assistance have no enforceable claim-right against you that you assist them. The sense in which you 'ought' to provide assistance is, of course, largely toothless on this view. But the point is that when we ask what the would-be Samaritan ought to do when deciding whether to provide life-saving assistance at little cost to herself, the Right-Libertarian can coherently and consistently say that she ought to render that assistance.

My proxy might seem to be in a similar position: I have a right that no one prevent her from saving me at the opportunity cost of

allowing the two strangers to die. But that's just a duty that others have to refrain from interfering. That still leaves open what my proxy should do. And what she should do is save the two strangers, since the personal value attaching to my life has no reason-giving force for her.

If this account is correct, then I have not shown that my proxy ought to save my life. Rather, all I have shown is that third parties have a duty of non-interference with respect to her decision. This is still an interesting conclusion, in that absent her role as a proxy, third parties could and should compel her to act according to her duty to save the two strangers. That is, her duty to save the two (when doing so costs her little) is an enforceable duty (Libertarianism notwithstanding). But where she *does* function as my proxy, my rights prohibit them from compelling her to choose optimifically.

But I want to show more: my proxy does not have merely a permission to save me, but an obligation (albeit a defeasible one, of course). Suppose my proxy does indeed have a non-enforceable duty to save the two – a *duty* on the grounds that it does substantial good at little cost to her, and a *non-enforceable* one on the grounds that I have a claim against third parties that they refrain from interfering with my proxy. There is, even in such a case, someone who does indeed have an enforceable claim on my proxy: me. I have, after all, contracted with her. Since her duty to rescue the two strangers is non-enforceable, and her contractual duty to save me is indeed enforceable, this suggests that her duty to me wins out. Hence, my proxy has an enforceable duty to save me rather than the two others.

To recap: we have enforceable duties to act optimifically when doing so comes at little cost to us. But when an individual contracts with me to act as my proxy in that she agrees to vicariously enact the agent-relative permission I have to give my life extra weight, I have a right against third parties to permit *my proxy* to do what my rights permit, which means she can permissibly save my life at the opportunity cost of allowing several others to die, if necessary. The contractual obligation my proxy has toward me requires that she avail herself of that permission.

So far I've presented an account of how vesting agent-relative permissions works. In what follows I outline what it is two indi-

viduals have to do in order for one of them to vest her agent-relative permissions in the other.

B. *Establishing a Proxy*

For me to vest in you my agent-relative permission to do ϕ requires that we form an agreement in which you promise to enact my agent-relative permission to do ϕ . The normative force of the subsequent reason you have to enact my agent-relative permission is derived from the normative force of the reason you have to fulfill the promise you made. At first, this seems problematic. After all, as I indicated in Section III.A, promises do not warrant acting in the sub-optimific way which agent-relative permissions countenance. So what explains the fact that my proxy can forego saving two in order to save me, when no promise is stringent enough to justify acting in that way?

The answer lies in the machinery I outlined in the Section IV.A. You, as my proxy, are permitted to save me rather than the two others not because the promise you made is so strong that it outweighs your obligation to save the two others, but rather because I exercise a right against third parties that they treat your acts as mine. The result is that no one (including the two others) can legitimately demand that you act optimifically should doing so transgress my permission to give my welfare priority. Hence, as I indicated in Section IV.A, any duty you have to act optimifically is at best non-enforceable – which conflicts with the enforceable duty you have to fulfill the promise you made to me. So though your promise to save my life is not strong enough to warrant doing so at the cost of failing to save two others, your promise *combined with* my decision to exercise my right against third parties that they treat your acts as mine permits you to save me at the cost of failing to save two others.

However, merely promising to do ϕ for me is not enough to make you my proxy. When I vest in you my right to do ϕ , you not only agree to do ϕ but also agree to do so on the understanding that I can legitimately demand of third parties that they treat you in a particular way. Recall that on the account I presented, when I vest a right in you as my proxy, the conduct the right protects is *your* conduct. But the right is still mine in that I am entitled to demand

from others certain kinds of treatment toward you. If the relevant third parties violate that demand, then *I* am wronged.

In what way can I demand that they treat you? Recall the demand that I legitimately make of third parties when you vicariously enact the agent-relative permission I have to save my own life: the demand I make is that they refrain from interfering with your attempt to save my life. Such a demand certainly does not impose a burden on you. Indeed, agent-relative permissions – as one might expect from their status as permissions – will in general require little more than forbearance from third parties.

Given that what I demand of third parties is, in effect, that they refrain from interfering with your conduct, why is it important that you recognize those demands? It is precisely recognition of that demand's legitimacy that enables you to permissibly enact my agent-relative permissions. Without that recognition, you wouldn't be able to see yourself as acting permissibly *qua* proxy. Suppose you do not recognize that I can legitimately demand of third parties that they refrain from interfering with you when you save me instead of two others. In such a case, you would act knowing that you should be stopped. Clearly, then, you are not functioning as a proxy (or you are only nominally a proxy) insofar as you take yourself to be doing something that I – but not you – have permission to do. Put differently, if you fail to recognize that I can demand of third parties that they refrain from interfering with your conduct, then you cannot take yourself to have the sorts of rights which acting *qua* proxy confers.

So for me to vest in you my agent-relative permission to do Φ , you must not only agree to do Φ , but also recognize that *I* have a right to do Φ , and that I can legitimately demand of others that they permit you to do Φ for me, where absent that demand they would have a right or even an obligation to prevent you from so acting.

In addition, vesting a right in a proxy must satisfy a publicity-condition. When I vest a right in you as my proxy, I can legitimately demand of third parties that they treat you in the ways that the right specifies. But those third parties have to be in a position to recognize that their acts are normatively constrained in that way. That is, the third parties need to be able to recognize you as my proxy. The onus will typically (but not necessarily always) be upon me to inform the

relevant third parties that you are acting as my proxy. If I gratuitously fail to do so, I cannot reasonably expect those third parties to treat you as my proxy. In such a case, should their actions interfere with your conduct in a way forestalling my efforts to vicariously enact a right I have, they would not be to blame for the setback to my interests; rather, I would, insofar as the onus was upon me to inform the relevant third parties that I am choosing to exercise a claim against them via you. Vesting, then, will typically involve communication – not just with the proxy, but with the third parties whose conduct is normatively constrained by vesting the right.

This is not to say, of course, that I need to verbally communicate with the relevant third parties either one-on-one or *en masse* when I vest my right in a proxy. It is enough if I make it known to others, by some sign or signal, your status as my proxy. The importance of making your status known will presumably vary with the importance of the agent-relative permission in question and with what is at stake in vicariously discharging it. If communicating with the relevant third parties is impossible or prohibitively difficult, presumably no one is to blame for the resulting setback to my interests should third parties interfere with my proxy's conduct. (These are all *ceteris paribus* claims – for particular kinds of vicariously enacted rights, we might adopt norms specifying whether the onus is upon me to inform third parties that you are my proxy, or whether the onus is upon them to presume as much).

So vesting in you my agent-relative permission to do ϕ , requires that a) you agree to do ϕ , b) you recognize that I can legitimately demand of others that they permit you to do ϕ for me where otherwise they might have a right to stop you, and c) you are in a position to evince your status as a proxy to relevant third parties. To be clear, though, vesting an agent-relative permission in a proxy is not an all-or-nothing affair. It is possible, for instance, for me to vest an agent-relative permission in you, but with only half the reason-giving force that it provides for me. In this respect, vesting is scalar. Indeed, in certain circumstances it might be impossible to *fully* vest an agent-relative permission in a proxy, specifically when the proxy has competing commitments.

Given the account I've presented so far of how agent-relative permissions are vested, one might raise the following puzzle. The

agent-relative permissions *I* have are mere permissions. Yet the agent-relative reasons my proxy has to promote *my* ends are duties (albeit defeasible ones). This might suggest that whatever the content of my proxy's duty toward me is, the fact that it is a *duty* means we cannot characterize the end to be promoted as an agent-relative reasons *I* have, because those reasons are for me discretionary in a way that my proxy's duty is not.

This puzzle is easily solved. It is perfectly cogent to say that one can have a duty to promote a permission. Most ordinary promises take this form. For example, suppose my friend decides to fly to a conference. This is something she is morally permitted but not required to do. If I promise to give her a ride to the airport, I thereby have a defeasible duty to help her do what she is merely permitted to do. Similarly, if I have an agent-relative permission which I vest in a proxy, then my proxy has a duty to act accordingly. This means, though, that the duty she has as my proxy retains its status as a duty only insofar as I authorize her to exercise the concomitant agent-relative permission. Suppose I decide to willingly sacrifice myself to save the lives of the two strangers, and I communicate my desire to do so to my proxy. She consequently is no longer in a position to vicariously enact my agent-relative permission to choose my life over theirs, precisely because I have chosen to forego that privilege. There is in this case no operative agent-relative permission for her to vicariously enact.

V. THE PERMISSIBILITY OF VESTING AGENT RELATIVE PERMISSIONS

So far I have argued that it is conceptually *possible* to vest agent-relative permissions in a proxy. In what follows, I argue that it is *permissible* to vest. There might seem to be reasons for thinking otherwise, precisely in the sorts of cases I've described. It might seem fundamentally unfair to vest agent-relative permissions when doing so vitiates the rights others have to be saved. It seems unfair for you to vest your agent-relative permissions on occasions in which doing so means that two other victims will thereby lose their right to be saved. Of course, if you have a *right* to vest your agent-relative permissions in a proxy, then the two others are not wronged or treated unfairly when your proxy acts accordingly. But we cannot just stipulate that you have a right to vest in such a case; rather, that

claim needs to be shown. What we need is an argument demonstrating that a right to vest is morally attractive in the first place.¹⁴ That is my task in this section. To show as much, I argue that a regime in which we are permitted to vest agent-relative permissions better protects our status as inviolable beings. There are, accordingly, contractualist-based reasons for thinking that we have a right to vest in that way.

A. *Vesting and Status-Protection*

As Thomas Nagel noted, we place importance not just on axiological values (such as happiness, knowledge, preference-satisfaction, beauty, etc.) but on our *status* – specifically, our protected status as beings which cannot be permissibly sacrificed at the bar of just any greater good.¹⁵ As suggested in Section II, this status is what grounds agent-relative permissions; they enshrine our inviolable status by allocating to each person some space in which to give impersonally sub-optimal priority to his or her own well-being. Inasmuch, I do no wrong should I save my own life instead of the lives of two strangers. The two strangers have no basis for claiming that my choice was unfair to them.

Each individual's inviolable status is better protected given that each individual has a right to vest that permission in a third party. After all, if we have such a right, we can enlist the assistance of third parties in furtherance of enacting our own agent-relative permissions, thereby extending our ability to prioritize our own interests. In contrast, if we lack a right to vest our agent-relative permissions, then each individual is fundamentally and inescapably 'on her own' in enacting her agent-relative permissions. On this view, if their circumstances prevent them from doing so without assistance then they are out of luck.

So which view is correct? One way to determine whether our inviolable status does in fact yield a right to vest agent-relative permissions in a proxy is by adopting a contractualist framework. We can ask: given suitably situated and aptly motivated contractors, would they prefer a *permissive* regime in which we possess the right

¹⁴ I thank a referee for pressing me on this point, on and for articulating the challenge in these terms.

¹⁵ Nagel (2007).

to vest agent-relative permissions? Or would they prefer a *restrictive* regime denying such a right? If they would prefer the former, then two strangers whose lives my proxy foregoes saving in favor of my own would have no basis for claiming that my proxy's choice was unfair to them, given that they themselves would have agreed to a regime permitting that sort of choice. (Of course, the two strangers happen to end up worse off under the permissive regime. But each contractor's knowledge of whether she will *personally* benefit or suffer under the candidate regime would be suppressed.)

Against the claim that the permissive regime better protects our status as inviolable beings, one might point out that the regime is likely to make things go impersonally *worse*. Allowing a greater range of protections means there will be more instances in which others can, for my benefit, forego acting in furtherance of the greater good. That benefits *me*. But it also means that there will be more instances in which others can, for the sake of *someone else's benefit*, forego acting in furtherance of the greater good. In that case, I end up worse off since I am not a beneficiary of the vicariously enacted agent-relative permission, but rather among those who would have benefited by an increase in the impersonal good.

But this argument proves far too much in that, if correct, it would undercut agent-relative permissions *simpliciter*. If our goal is to choose a regime that maximizes the expected impersonal value accruing to any arbitrarily chosen person, then the best regime is one void of *any* agent-relative reasons, since they by definition permit acting sub-optimifically. Yet we are willing to accept lower expected impersonal value if doing so is necessary to ensure our protected status as inviolable beings who cannot be sacrificed for just any greater good. This is because, as Nagel puts it, “[w]hat actually happens to us is not the only thing we care about: What *may* be done to us is also important, quite apart from whether or not it *is* done to us – and the same is true of what we *may do* as opposed to what we actually do.”¹⁶

His point is that if there were no rights, the very fact that we would be *in principle* sacrificeable in furtherance of just any greater good would diminish us, even if we happened to suffer none of the transgressions that the rights protect against. By preventing us from

¹⁶ Nagel (2007, p. 108).

being diminished in that way, agent-relative permissions protect our status as inviolable beings.

A permissive regime, in which we have the right to outsource the means by which our agent-relative permissions are enacted, provides that much greater personal protection against being sacrificed for the benefit of the greater impersonal good. A restrictive regime, on the other hand, insofar as it prohibits us from vicariously enacting our agent-relative permissions, undermines the importance of the protections which the agent-relative permissions vouchsafe. These are reasons to think that we would adopt the permissive regime if we could.

Of course, we do not place infinite importance on our status as inviolable beings. We are willing to accept a regime that permits sacrificing an individual when doing so is necessary to avert a *catastrophic* harm because we do not think that any single person's status is so important as to warrant protection under those circumstances. Though we do not want a regime that maximizes the expected value accruing to any arbitrarily chosen person, neither is it the case that we want a regime that ignores expected value (or more precisely: that relegates them to the role of tie-breaker among candidate regimes in which our status is equally protected). The result is that though we want some protection against sacrificeability, we do not want absolute inviolability.

The trick, of course, is to strike a balance between protecting our status as inviolable beings on the one hand and promoting agent-neutral values on the other. It might seem that the permissive regime fails to achieve that balance. In what follows I consider and respond to reasons why one might think that the permissive regime is too permissive.

B. Vesting and Fairness

There are a pair of related reasons for thinking that we might reject the permissive regime on the grounds that it is too permissive. I will consider both before responding to them univocally.

First, the permissive regime seems to unfairly advantage those who are already the best off. If anyone can, given the requisite resources, vest their agent-relative permissions in anyone else, then

this would further empower the most powerful in an already dismally inegalitarian world. A wealthy individual could potentially hire an army of proxies, each permitted to vicariously discharge her agent-relative permissions. This is clearly an undesirable outcome. Contractors might reject the permissive regime on these grounds.

Second, the permissive regime seems to unfairly advantage those who happen to have access to third parties in whom to vest our agent-relative permissions. Suppose, again, that you and two others are drowning; a third stranger passing by can at little cost to herself rescue either you or the two strangers, but not both. The stranger is about to save the two others, when you ask for her to save you instead. She points out that she has a presumptive duty to save the greater number. But then you make clear to her that you would like to vest the agent-relative permission you have to save your own life in her, thereby permitting her to save you. The other two victims are unable to vest their rights in the stranger, either because a) they are unable to communicate with the stranger, b) you outbid the two victims by offering more cash than they possess, or c) the stranger arbitrarily chooses to act as your proxy. In each of these cases it is counterintuitive, given the described circumstances, to suggest either that the stranger is permitted to save your life or that you can vest your agent-relative permission to save your own life in the stranger. To the extent that the permissive regime countenances this outcome, contractors might reject it.

The answers to these two worries that the permissive regime is too permissive are one and the same: a permission to vest agent-relative permissions is not meant to operate in a moral vacuum, but instead against a backdrop of competing rights and values, of which fairness is one. Take the first worry. The fact that an unfettered right to vest agent-relative permissions in others would lead to an unacceptably inegalitarian outcome is reason for thinking that there are egalitarian constraints on exercising the right to vest. This is characteristic of rights in general.

Consider the following analogy. The fact that an unfettered right to private property would yield unacceptably inegalitarian consequences is not itself an argument against the right to own private property. Instead, it is an argument against an *unfettered* right to own private property. The fact that an unfettered right of that sort would

lead to a grossly inegalitarian outcome can be a reason for thinking that there are egalitarian constraints on exercising that right. Likewise, the fact that an unrestricted right to vest agent-relative permissions would yield unacceptably inegalitarian consequences is not itself an argument against the right to vest. Instead, it is an argument against an *unrestricted* right to vest. In particular, there might be an egalitarian constraint on vesting agent-relative permissions in cases where doing so further empowers the best off, especially if doing so comes at costs to the significantly worse off.

With respect to the second objection, fairness requires that the stranger give due consideration to the competing interests of the potential victims in deciding whether to accept a role as a proxy. To be clear, once she *is* a proxy she is perfectly permitted – indeed, she has a duty – to give priority to the interests of the individual who has vested rights in her. After all, the very purpose of an agent-relative permission, which the proxy is tasked with discharging, is to enable each of us to give defeasible priority to her own interests. So once an individual accepts a role as a proxy, fairness plays a substantially diminished role in deciding whom to favor (for reasons articulated in the previous section). But considerations of fairness apply in full force when deciding *whether* to accept a role as a proxy in the first place. And it would be presumably unfair for the stranger to do so in cases ‘a’, ‘b’, and ‘c’.

The two worries, then, are mirror images of each other. In response to the first worry, I pointed out that there are fairness constraints on our right *to vest* our agent-relative permissions in a proxy. In response to the second worry, I pointed out that there are fairness constraints on our right *to accept* a role as proxy. Once these fairness constraints are articulated, the permissive regime would, presumably, be far more attractive to contractors who might otherwise worry that the regime is too permissive.

Ultimately, then, a comprehensive account of vesting would have to provide fairness constraints which delineate the *specific conditions* under which we can permissibly vest agent-relative permissions. That is beyond the purview of this paper, which is to show that (i) vesting such permission is conceptually possible, and (ii) we have a right to vest such permissions.

I started Section V with a challenge to my account: even if it is conceptually possible to vest agent-relative permissions, why believe that we have a right to do so? I argued that such a right more effectively protects our status as inviolable beings, and that as such, there are contractualist reasons for thinking that we have such a right – albeit one that operates against a backdrop of competing rights and values, of which fairness is one. I conclude by considering the applications of a right to vest.

VI. APPLICATIONS

So far I have argued that the sense in which our agent-relative permissions are first-personal does not preclude the possibility of vesting them in a proxy. I have also provided reasons for thinking why a right to vest is morally attractive. For want of space, I have not laid out a comprehensive account of the *specific* conditions under which it is permissible to vest. But it is worth exploring, in closing, some possible applications of such a permission.

Some individuals, such as persons with physical disabilities, have more difficulty achieving their ends (especially when society is rigged in ways that disadvantage them by catering solely to the capabilities of the fully abled).¹⁷ Consider, then, a paraplegic unable to move about independently. She needs assistance accomplishing basic tasks. If the account I have provided here is correct, then she can vest in a proxy her agent-relative permission to give priority to her own wellbeing, thereby effectively enabling her to vicariously exercise the basic rights which we all have in virtue of our status as persons. Absent such vesting, third parties would not be permitted to give her basic interests the full weight that *she* is permitted to give them. An implication of the account I have presented here, then, is that the physically impaired are not irrevocably and unfairly ‘on their own’ when it comes to enacting their agent-relative permissions.

Or consider the context of warfare. What is a soldier fighting in a just war permitted to do in furtherance of protecting the lives of the civilians for whom she is fighting? Innocent civilians have an agent-relative permission to give priority to their own welfare over the welfare of others. Arguably, such a permission not only allows us to

¹⁷ See the modern classic, Anderson (1999).

forego doing the impersonally greater good in cases where that option would be detrimental to our interests, but also permits us to do more harm than would be impersonally permitted.¹⁸ If we can licitly construe the civilians as having vested *that* agent-relative permission in the soldiers fighting on their behalf, then the soldiers might thereby be permitted to do more harm in furtherance of winning that war than they otherwise would. That is to say, if soldiers are acting as proxies for their civilians, then that weakens the constraint of proportionality, which specifies the amount of harm that can be done in furtherance of achieving a particular good. The result is that soldiers waging a just defensive war are permitted to do more than might initially be thought (which is not to say that the amount of harm that they typically do is permitted).¹⁹

Finally, consider the duty of those who have temporary guardianship over children. Suppose a parent has an agent-relative permission to give defeasible priority to the well-being of her own infant child. This parent, due to circumstances beyond her control, temporarily places the infant in the care of another. The temporary guardian has a duty to give defeasible priority to the infant's well-being over that of strangers. One way to make sense of this duty is by characterizing the temporary guardian as a proxy; she has a duty to enact *the parent's* agent-relative permission to give priority to the well-being of the infant. The guardian must, for example, save the infant under her care even if doing so means allowing two other infants (to whom she has no more than Samaritan duties) to die. It is difficult to make sense of the peculiar strength of the duty the guardian has, absent the possibility of vesting. (Recall from Section III.A that a mere promise cannot do the requisite normative work.)

This, of course, only scratches the surface when it comes to the potential applications of a permission to vest agent-relative permissions. Moreover, I have said nothing about whether the individual whose permissions are vested is morally responsible for what her proxy does. A fuller treatment of the applications and ethical implications of vesting agent-relative permissions, I leave for another time.

¹⁸ See fn. 6.

¹⁹ In addition to the works by Lazar I've cited in this paper, see Bazargan-Forward (2018).

REFERENCES

- Anderson, E. (1999). What is the Point of Equality? *Ethics*, 109(2), 287–337.
- Bazargan-Forward, S. (2018). Weighing Civilian Lives in War: Domestic Versus Foreign. In L. May, *Cambridge Handbook on the Just War* (pp. 186–198). New York: Cambridge University Press.
- Broad, C. D. (1930). *The Philosophy of Broad*. LaSalle: Open Court.
- Davis, N. (1984). Abortion and Self-Defense. *Philosophy and Public Affairs*, 13(3), 175–207.
- Fabre, C. (2012). *Cosmopolitan War*. Oxford: Oxford University Press.
- Frowe, H. (2008). Threats, Bystanders and Obstructors. *Proceedings of the Aristotelian Society*, 108(3), 365–372.
- Hart, H. (1982). *Essays on Bentham: Jurisprudence and Political Philosophy*. Oxford: Oxford University Press.
- Korsgaard, C. (1996). *Creating the Kingdom of Ends*. Cambridge: Cambridge University Press.
- Kramer, M. H. (2001). Getting Rights Right. In M. H. Kramer, *Rights, Wrongs, and Responsibilities* (pp. 28–95). London: Macmillan.
- Lazar, S. (2013). Associative Duties and the Ethics of Killing in War. *Journal of Practical Ethics*, 1(1), 3–48.
- Lazar, S. (2016). Authorization and the Morality of War. *Australasian Journal of Philosophy*, 94(2), 211–226.
- Lefkowitz, D. (2009). Partiality and Weighing Harm to Non-Combatants. *Journal of Moral Philosophy*, 6(3), 298–316.
- Lyons, D. (1994). *Rights, Welfare and Mill's Moral Theory*. Oxford: Oxford University Press.
- MacCormick, N. (1982). *Legal Right and Social Democracy*. Oxford: Oxford University Press.
- McNaughton, D., & Rawling, P. (1995). Value and Agent-Relative Reasons. *Utilitas*, 7(1), 31–41.
- Nagel, T. (1970). *The Possibility of Altruism*. Princeton, NJ: Princeton University Press.
- Nagel, T. (1989). *The View from Nowhere*. Oxford: Oxford University Press.
- Nagel, T. (2007). The Value of Inviolability. In P. Bloomfield (Ed.), *Morality and Self-Interest* (pp. 102–113). Oxford: Oxford University Press.
- Nozick, R. (1974). *Anarchy, State, and Utopia*. New York: Basic Books.
- Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press.
- Pettit, P. (1987). Universality Without Utilitarianism. *Mind*, 72, 74–82.
- Quong, J. (2009). Killing in Self-Defense. *Ethics*, 119(3), 507–537.
- Quong, J. (2016). Agent-Relative Prerogatives to Do Harm. *Criminal Law and Philosophy*, 10(4), 815–829.
- Railton, P. (1984). Alienation, Consequentialism, and the Demands of Morality. *Philosophy & Public Affairs*, 13(2), 134–171.
- Rawls, J. (1971). *A Theory of Justice*. Cambridge: Harvard University Press.
- Raz, J. (1986). *The Morality of Freedom*. Oxford: Oxford University Press.
- Scanlon, T. (1998). *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Scanlon, T. (2003). *The Difficulty of Tolerance*. Cambridge: Cambridge University Press.
- Scheffler, S. (1982). *The Rejection of Consequentialism*. Oxford: Oxford University Press.
- Schroeder, M. (2007). Agent-Relative Value, and 'Good'. *Ethics*, 117(2), 265–295.
- Steiner, H. (1994). *An Essay on Rights*. Oxford: Blackwell.
- Tadros, V. (2011). *Ends of Harm*. Oxford: Oxford University Press.
- Wellman, C. (1995). *Real Rights*. Oxford: Oxford University Press.

Department of Philosophy
 UC San Diego, La Jolla, CA, USA
 E-mail: sbazargan@ucsd.edu