**Patient Moral Luck**[1]

**Preston J. Werner**

*Forthcoming in Oxford Studies in Normative Ethics, Volume 14*

**{Please cite the published version}**

Word Count: 8946

Bernard Williams & Thomas Nagel (1976) pointed out something odd, perhaps even paradoxical: Our assessment of agents' moral character can come apart from voluntary aspects of their actions. The negligent driver who hits a child is more blameworthy than the equally negligent driver who does not, for example. Moral luck continues to be widely discussed today.[2] In this paper, I argue for a fundamentally different kind of moral luck, *Patient Moral Luck* (PML).[3] Unlike traditional moral luck, PML concerns the amount of moral consideration that different moral patients — that is, creatures (including human beings) with moral status — will be owed, independent of factors in their control. PML, I argue, entails that morality *itself* appears to sanction and even obligate actions which, along predictable patterns, involve repeatedly failing to equally consider certain moral patients - and repeatedly the same people - over sustained periods of time, through no fault of their own. And often these people will be members of groups who are *already* worse off through no fault of their own, thus exacerbating unjust inequalities.

In this paper, I introduce and elaborate on PML and why it generates a normative puzzle. This involves three claims which I call the *theoretical claim*, the *empirical claim*, and the *normative claim*. The *theoretical claim* is the claim that PML is a coherent and possible phenomenon, and that it can be separated from plain, morally relevant luck as applied to moral patients. The *empirical claim* is that PML is actually a pervasive phenomenon in our world. Finally, the *normative claim* is the claim that this phenomenon raises a serious

---

[1] For very helpful comments, I'd like to thank Daniel Berntson, Teresa Bruno-Niño, Brad Cokelet, Matti Eklund, Aaron Elliott, David Enoch, Jessica Fischer, Amy Floweree, Tom Hurka, Jacob Litz, Clara Lingle, Oded Na'aman, Ittay Nissan-Rozen, Anna Nyman, Adam Patterson, Olle Risberg, Holly Smith, Andreas Stokke, Sarah Stroud, Will Yu, and audiences at Uppsala University and the 2024 Workshop in Normative Ethics (WiNE) in Tucson, Arizona. And thanks very much to Uppsala University and Olle Risberg for graciously hosting me as I completed work on a first draft of this paper.
This paper is dedicated to Khalil Abu Yahia and Vivian Silver, who refused to see their national identity as making their lives more valuable than others, and were killed by those who did.

[2] Nelkin (2019).

[3] To be clear, I do not intend to suggest that PML is an exact parallel to traditional moral luck. Traditional moral luck concerns agents' being differentially blamed for differential outcomes of the very same action and intention. PML does not operate in this way. However, I call it "moral luck" because, like traditional moral luck, it involves a pervasive inequality which is sanctioned by morality itself, i.e. *is internal to morality itself*.
If the reader prefers, she can replace "Patient Moral Luck" with "Unjust Considerative Inequality" or some such. Nothing crucial hinges on my choice of terminology.

normative puzzle, since it conflicts with deep intuitions we have about what I'll call diachronic equal consideration of moral patients.

After establishing these three claims, I consider how normative theories may accommodate PML. Aside from introducing the idea of PML, the core aim of this paper is to establish that normative theories must either accommodate or explain away Patient Moral Luck. While the phrase "moral luck" that I use here is used in homage to Williams & Nagel, I do *not* intend to claim that PML is exactly analogous to traditional (agential) moral luck, that it has the same source, or that its manifestations will be mirror images of each other. I merely use the phrase "Patient Moral Luck" in line with the fact that it is an as yet undiscussed way in which common sense moral principles can sanction and promote unfairness for individuals on the basis of brute luck.

The plan of the paper is as follows. I begin in section 1 with a few cases to remind ourselves of some common sense intuitions concerning balancing equal consideration of persons with the associative reasons that meaningful relationships generate. In section 2, I introduce the conceptual machinery of "moral force", which helps us to focus on moral principles from the moral patients' point of view. This allows me to introduce the notion of Patient Moral Luck in section 3. In sections 4 and 5, I argue that PML is widespread and that PML is normatively problematic, respectively. In section 6, I sketch out the logical space of how to respond to the problem of PML, considering two radical approaches and two more moderate but still revisionist approaches. Finally, in section 7, I take stock.

## 1. Some Cases

Let's work up to our puzzle by considering some cases. First, consider a hackneyed one:

> *Tragic Choice 1.* On his way home, Chris sees two strangers drowning. He isn't sure if he will have time to save both or just one of them.

At this level of description, Chris is obligated to do his best to save both strangers, but his choice of which one to start with - as long as it isn't motivated by some moral vice like racism or sexism - will be permissible either way. Let's consider, then, another 'easy' case:

> *Tragic Choice 2.* On her way home, Vera sees two people drowning. One of them is a stranger, but the other is her brother. She isn't sure if she will have time to save both or just one of them.

In this case, it seems obvious that it is permissible for Vera to save her brother first. She need not select at random. She may even be obligated to save her brother first, depending on one's specific theory about the strength and nature of Associative Reasons. Associative Reasons, as I'll use the phrase here, concern any reasons which are grounded in special

relationships between agents, such as parent-child, friend, sibling, colleague, or fellow-national.

These two cases illustrate two truisms about common sense morality. First, all else equal, every human patient deserves equal moral consideration. This explains why, when Chris comes across two strangers, as in *Tragic Choice 1*, he is permitted to save either of them (assuming he cannot save both). Let's call this idea *Equal Consideration*. Second, one way in which things will *not* be equal is when a moral agent and a moral patient bear a special relationship to each other which can ground *Associative Reasons*. Associative Reasons can permit or obligate agents to favor some moral patient over another on the basis of their relationship to each other.

Now let's consider a third, more realistic, but still quite simplified case:

> *Financial Support 1*. Octavia and Solita are both in similar situations - they were laid off a few months ago and do not know where to get the money to pay next month's rent. Both are at risk of eviction if they do not find the money, a fate that would be equally bad for each of them. Octavia has several siblings that she is close with, while Solita is an only child.

Octavia and Solita are in similar situations, but Octavia is lucky in at least one way that Solita is not - she has several siblings that she can reach out to for financial assistance, whereas Solita does not. As with *Tragic Choice 2*, this again seems to follow from the truism that healthy sibling relationships generate Associative Reasons. In a moment, I'm going to argue that cases like these involve normatively worrisome patient moral luck. But first, let's pause to introduce some of the conceptual tools that are necessary to lay this argument out.

**2. Moral Force**

Lawyers and epistemologists sometimes talk of "evidentiary force". Roughly, we might think of evidentiary force as something like the amount of *epistemic pull* toward some proposition P that some piece(s) of evidence provide(s). The evidentiary force of some set of evidence will vary based on the broader context as well as what proposition its evidentiary force is being assessed for (the same set of evidence may provide strong force for *P,* weak force for *Q*, and no force for *R*). We could also think of evidentiary force in terms of how (and how strongly) a rational agent would weigh the evidence when considering whether *P*.

Now consider the moral analog, "moral force".[4] (The name is new, but I believe the idea is familiar.) We can think of moral force as the amount, strength, and direction (positive or negative) of moral reasons that one moral patient's well-being normatively exerts on

[4] In an earlier version of this paper, I was calling this "moral charge". As Holly Smith helpfully pointed out to me, "charge" doesn't have the right formal features for the metaphor to work - charge is an intrinsic property, not a relational one. Thanks to Aaron Elliott for help with coming up with "moral force" as a better (though not ideal) way of framing the phenomenon I am trying to capture.

others.[5] (This is neutral between different first order normative theories.) We can think of moral force as a dispositional normative property of moral patients whose strength and direction depends on a wide range of background features. Moral force is a bit like the flip side to moral reasons - reasons apply to agents qua agents, but moral force belongs to moral patients qua patients.[6]

Moral force can be thought of in terms of a relation between some moral patient and some moral agent. But we can also consider moral force as relativized to some moral patient's local or global surroundings.[7] Consider:

> *Island.* Imagine an island containing only three people: Reed and Jamie, who are siblings, and Natalie, who is a stranger. Suppose Jamie and Natalie find themselves in the same perilous position on a hiking trail and are in need of saving. Reed walks this path sometimes, so there is a chance he will come across their perilous position. However, even if Reed does walk by, he will only be able to save one of them.

In *Island*, Reed has not walked by and has no way of knowing the trouble that Jamie and Natalie are in. But nonetheless, Jamie and Natalie each exert a certain amount of as-yet-dispositional moral force on him, waiting to be manifested in the event he arrives. And we also know how these forces compare, and thus what obligations will arise if Reed does arrive. If Reed deliberates and acts as he ought to do, Natalie has no chance of being saved, while Jamie does have a chance of being saved. That may ring a bit unfair, but it follows from common sense intuitions about Associative Reasons.

The strength of the moral force that Natalie and Jamie have can be relativized to a spatiotemporal location. Jamie has stronger moral force in *Island* when we relativize the situation to only the residents of the island. If we stipulate and consider the additional fact that Natalie's sister Roni is in a sailboat off the coast of the island, Jamie and Natalie's total moral forces in their local environment will equalize (or near enough equalize[8]). If we generalize spatiotemporal location more broadly, then even if Roni is quite far away from the island, Jamie and Natalie's total moral forces will still equalize (or near enough equalize), even though Jamie still has a higher chance of being saved. Moral force is the property of

---

[5] A closely related concept is Broome's concept of *claims*, which are reasons *owed* to a particular agent herself, as opposed to reasons justified in terms of some more general good. On the terminology of "moral force", moral force is the disposition to generate claims on others. (See Broome 1991 for discussion.)

[6] To be clear, moral force is not dependent on a moral patient *asserting* or *declaring* her moral status. It is a merely normative power, not a power of concrete physical or legal authority.

[7] Arguably this could also be done with "evidentiary force", but since I only appealed to evidentiary force by way of illustrating moral force, I set questions about how to cash this out precisely aside.

[8] I am of course oversimplifying here. The strength of the Associative Reasons, even between siblings, will intuitively vary in a myriad of ways based on the nature and history of the given sibling relationship.

one patient's ability to generate moral deliberative weight on a particular agent, not a statistical property about *how likely* that deliberative weight is to kick in.

What is important here is to make sense of the idea that we can compare Jamie and Natalie's total situational moral force despite the fact that some of the moral force we are dealing with here is "patient relative" (in the sense that some of the reasons the force can generate are agent-relative). So it should make intuitive sense that, in the original *Island* case (i.e. when Roni isn't being considered as relevant), Jamie has stronger total situational moral force. Things get more complicated if we are trying to compare total situational moral force in cases where there are more agents, more patients, and more potential for Associative Reasons. If we add more agents that could come across Jamie and Natalie's predicament, all of which bear different relations and histories to Jamie and Natalie, their respective total situational moral forces will require adding up many different strengths of many different sources of moral force. As we will see below, my argument for the existence of moral patient luck doesn't require an answer about how to do this; it does, however, require that there be an in principle fact of the matter about moral patients' total situational moral forces. And one might worry that there is no way to make such questions make coherent sense.

I don't pretend to have an algorithm for generating total situational moral force. Such a project would require a paper in itself. Instead let me briefly gesture toward the idea that there will be a way to do this, in principle. Moral agents are often in a situation in which they must weigh a variety of reasons that they have, where such reasons will involve a combination of agent-neutral reasons as well as agent-relative reasons of varying strengths. While it is certainly possible that genuine moral dilemmas — in the sense of decisions where no possible action would be morally permissible — could arise from such situations, this kind of moral weighting is not often taken to be evidence that such dilemmas exist. It is just the kind of deliberation that arises when, for example, I am deciding who to offer my spare bedroom to when the inn is full. If such weighting is possible on the side of reasons, it isn't clear why similar considerations wouldn't apply to the side of moral force. Keeping with the spare bedroom case, we could compare different moral patients' moral force in terms of how much consideration, *if everyone is acting morally optimally*, will be given to them with respect to offering them a room. If we iterated the case, moral force would correlate with how likely, if everyone is acting morally optimally, each patient is to receive a room. But it is important to stress that moral force is not per se about outcomes, but about weight in consideration. Of course weight in consideration is correlated with likelihood of outcomes, but they can come apart, such as if some agents have more rooms to offer than others, or if some agents are weak-willed and thus can't act on their morally optimal reasoning, etc.

Again, there are a number of complications that I am ignoring here. These would have to be dealt with if we were trying to formulate a comprehensive theory of moral force. Moral forces may interact in complicated ways given changing background conditions, for example, or given what combination of agents are around. These are not easy questions, but they don't need to be answered here.

### 3. Patient Moral Luck and Morally Relevant Patient Luck

I said above that I aimed to show three things with respect to cases like *Financial Support 1:* (a) they are pervasive and systemic in the real world, (b) they entail luck which is *internal* to morality, and (c) they are uniquely normatively worrisome as a result. These are the *empirical claim,* the *theoretical claim,* and the *normative claim*, respectively. In this section, I defend the theoretical claim.

In the literature on moral luck as traditionally construed (what I'll call *agent* moral luck), David Enoch has drawn a distinction between moral luck and "plain" (morally relevant) luck. A clear case of the former sort is Bernard Williams' case of the two negligent drivers, only one of whom hits a child. A clear case of the latter sort might be the difference between two agents, only one of whom happens across a drowning child and thus must ruin their suede jacket. One of these two agents is unlucky in the sense that they had to ruin their suede suit, but this doesn't raise any of the paradoxes that makes agent moral luck a puzzling phenomenon - it is just plain luck that happens to be morally relevant. Enoch diagnoses the difference here:

> In paradigmatic cases of (purported) moral luck, the luck is *internal* to morality; morality itself draws a distinction that is a function of luck. In cases such as [the suede jacket case], however, morality is sensitive to non-moral circumstances, and because the latter are partly a matter of luck, morality is sensitive to luck as well. But here the luck is external to morality, and it is only morality's sensitivity to non-moral circumstances that incorporates the effects of luck — of things, that is, that are not under the control of the relevant agents — back into morality.[9]

This distinction between luck which is internal and luck which is external to morality is imprecise, as Enoch himself later admitted.[10] Nonetheless, the distinction seems important, and I know of no more precise proposal on offer. So I will take this rough metaphor as good enough for my purposes.

To see why I take *Financial Support 1* to raise unfairness which is *internal* to morality, compare it with a related case:

> *Financial Support 2.* Felix and Pat are both in similar situations - they were laid off a few months ago and do not know where to get the money to pay next month's rent. Both are at risk of eviction if they do not find the money, a fate that would be equally

[9] Enoch (2007), p.39, emphasis his.
[10] Enoch (2019), p.259.

> bad for each of them. Felix has a sister, Grace, that is willing to help him. Pat has a sister, Dina, who would be willing to help him, but unfortunately just spent her extra savings on a non refundable necessary but non-urgent purchase.

As with Octavia and Solita, Felix and Pat's outcomes are predictable: Octavia and Felix will be able to meet their rent payment, and Solita and Pat will not. However, Solita's misfortune is explained by her lack of total moral force, while Pat's is not about moral principles or morality *per se* at all. He was simply unlucky in that Dina happened to make a purchase that prevents her from assisting. Felix and Pat's moral charges are in relevant respects equivalent, but their ability to discharge these charges differs because of plain, worldly luck.

More generally, then, when some moral patient's harm or lack of a good is a result of their (uncontrolled) total moral force, we have what I am calling *Patient Moral Luck* (PML). On the other hand, when some moral patient's harm or lack of a good is a result of uncontrolled worldly facts that don't affect her underlying total moral force, this is morally relevant plain luck.[11] My concern in what follows is only with PML. However, we must be careful here to distinguish PML from the *consequences* of PML on a given moral patient's downstream welfare. Of course, unequal consequences of PML are concerning in themselves, but in principle they can be addressed as a special case of the application of egalitarian and/or prioritarian principles.[12] PML is about inequality in *total moral force* per se: In other words, the *internal-to-morality* aspect of PML is about the distribution in moral force that a given moral patient has on their broader environment over time. One can have or lack the relevant relational properties which ground their proper moral consideration in the deliberation of others, and these properties can be systematically unevenly distributed. It is in this sense that PML is moral luck and not simply morally relevant luck.

One more note about PML before proceeding: Whether some situation will involve PML will depend on one's first order normative commitments. The limit case of this, of course, is the pure agent-neutralist about moral reasons. The pure agent-neutralist rejects the existence of *any* non-derivative Associative Reasons. As such, for her, the puzzle of PML cannot get off the ground - there simply will be no PML, because all patients will at all times have the same moral force on all agents. However, even setting the agent-neutralist aside,

---

[11] Compare someone who is drowning in a lake far in the wilderness where there are no other people around with someone drowning in a lake that strangers often walk by. The former does not have less moral charge than the latter merely because the charge cannot manifest, so this is merely morally relevant patient luck.
Graph theoretically, we could model this in terms of edges between moral agents/patients as signifying moral consideration owed. On this modeling, unequal distribution of edges (in number and strength) would constitute PML, while unequal resources held by nodes in the graph represent morally relevant patient luck. (Formalizing such a model would take us outside of the scope of the present paper, but I hope to explore this in future work.)

[12] In fact I do think that this is an independently neglected issue.

there will be disagreement about which kinds of relationships (and under what circumstances) generate Associative Reasons. The scope of PML will, in the first instance,[13] depend on which Associative Reasons different agents have and why.

Things will get tricky when we think about the implications for moral force from complications that arise when reasons interact in non-additive ways. For example, consider Williams' case where one must choose to save a stranger or one's spouse. Williams argues that it is not only obligatory to save one's spouse, but to even have to weigh up one's reasons in such a case would constitute a moral failure.[14] One way to read this case is as one where the reasons to save the stranger are not just outweighed, but completely silenced. It's an open question how to think about such a situation from the standpoint of moral force. Consider the drowning stranger: Does she now lack *any* moral force vis-a-vis the rescuer? Or does she still have moral force, but her dispositional force is somehow *masked* from being manifested, like a fragile glass wrapped in bubble wrap?[15] If we wanted to develop a full account of moral force, we would need to consider all of the ways that forces can interact with each other just as we consider all of the ways that reasons can interact. Unfortunately, I do not have the space to address these questions here.

**4. The Pervasiveness of Patient Moral Luck**

I've defended the theoretical possibility of PML. This, in itself, should be of theoretical interest. However, at this point, it may be thought that, while this is an interesting philosophical phenomenon, it won't make any difference to our on-the-ground normative theorizing. The purpose of this section is to argue for what I've called the *empirical claim:* PML is quite pervasive, and pervasive in ways that systematically re-entrench inequalities that we independently find normatively problematic.

Let's begin with the relationship that I've used to motivate PML in the first place - the sibling relationship. The sibling relationship is the cleanest case to illustrate PML because (a) the number of siblings one has is not in one's control, and (b) it is one of the least controversial cases of a relationship that generates Associative Reasons. In itself, this is sufficient to show that PML is going to be ubiquitous. Similar considerations will apply to other aspects of one's familial relationships (how many parents one has in one's life, the size of one's extended family, etc.).

---

[13] In this paper, I am focused only on PML as it arises from the conflict between Associative Reasons and the entitlement to equal moral consideration that moral patients have. I think it is an open and complicated question about whether and how PML may arise from other ways that patients can gain greater moral charge (being the recipient of a promise, for example).Because of the complications here, I set these questions aside in what follows.

[14] Williams (1981), pp.1-18.

[15] See, e.g. Ashwell (2010) and Choi & Fara (2018) for discussions of dispositional masks.

Next consider friendships and romantic partnerships. Such cases are messier for a few reasons. First of all, who and how many friends we have is in many ways in our control. This is true both for practical reasons (more or less effort can be made to make friends) as well as reasons which reflect on our moral character (less people want to be friends with a jerk). However, there will be differences in amount and closeness of friendships that extend beyond the factors that are in an agent's control. Attempting an exhaustive list would be foolish, but here are some factors that will partially determine one's success in making meaningful friendships: (a) where one grows up, (b) one's gender and sexual orientation, (c) one's religious background, (d) one's hobbies and interests, (e) one's physical attractiveness[16], (f) one's intelligence, (g) one's skills, (h) one's charisma, etc. Of course, many of these things are, to some extent, within our control. If I live in a culture where juggling is very highly valued (no party is complete without a juggler, of course), and I have some minimum amount of time and motor skills, I can learn how to juggle, thus improving my opportunities. On the other hand, if I live in a culture where most socially valued skills are ones that are not easily within my physical or financial reach, I won't be able to gain these skills at all, or at least not without a lot of difficulty and strain. Similar things will be true of many of the factors listed above: While there will be some flexibility and control over one's skills, hobbies, intelligence, and so forth, (i) the amount of effort that one must put into altering these things will vary across agents, and (ii) there will still be practical limitations on how much alteration individuals can accomplish.

In short, not just the familial relationships we have, but the majority of our social connections and meaningful relationships depend on factors that are a matter of luck. In itself, this is "plain" luck - morality does not give us (moral) reasons to ensure that some people have less meaningful relationships than others, that is just an unfortunate fact of life. However, the upshot of this "plain" luck, given how Associative Reasons are generated, results in a similar pervasiveness in (genuine) PML. Again, morality appears to take no stand on how many Associative-Reason-generating relationships given individuals should have; this is external to morality. However, once the relationships are set, morality permits and requires partiality in moral consideration between individuals who bear certain relations to each other, *whether or not those relations are morally arbitrary*. So the variance in total moral force that different individuals will have will predictably and systematically vary along non-morally relevant lines. Thus, PML is pervasive.

*4.1 It All Evens Out in the End? Version I*

[16] Krantz (1987)

Before turning to argue that the pervasiveness of PML is normatively problematic, it's worth briefly addressing an objection one may have to the argument just given. It's true, one might argue, that all kinds of arbitrary and uncontrollable factors can determine the cluster of relationships we have, and that this has implications for the distribution of moral force that individuals have. However, this wouldn't be concerning unless we have reason to believe that, ultimately, individuals' total moral forces will widely differ (on arbitrary grounds) between each other. Consider a hypothetical world in which somehow all people had one and only one sibling. In such a hypothetical, while there would be many sibling-related Associative Reasons, human moral patients' sibling-related moral force would roughly[17] even out. So the problem of PML does not arise from the mere presence of Associative Reasons, but only with their (substantially) uneven distribution.

The hypothetical case in which all people have one and only one sibling is very far from the actual world, so let's now return to how the 'it all evens out' objection might go in the real world. The thought is that the factors that affect one's opportunities for meaningful relationships will - setting aside moral failings - more or less predictably even out. Perhaps, for example, someone being an only child increases the amount of time they have to dedicate to cultivating meaningful friendships. Maybe having an odd or unpopular skill can be spun into a uniquely interesting quirk which draws people to you. It is at least prima facie plausible, as pop culture claims, that introverts have deeper friendships even if they have fewer friends than extroverts, and maybe this results in introverts having more highly concentrated but equivalent overall moral force. And so on.

In response to this worry, I want to say two things. First of all, I certainly concede that the factors that influence what kinds of relationships we form are extremely complicated. And in some cases, the very same factor can harm our opportunities in some respects while it helps in others. There is no doubt about this, as any glance into the literature on the sociology of friendship will show.[18] In my claims above, I certainly didn't mean to suggest that the many factors listed as affecting relationship opportunities are simple causal stories, much less destiny. However, and second, this is not the point. The complications raised above are genuine complications, but they don't give us reason to think that the complications will uniformly balance out, and to suggest as such without a general argument to this effect strikes me as unjustified and radical optimism. Perhaps there could be some "Master Argument" to suggest that things will even out, but I am not sure what

[17] See fn.4.
[18] Kelley & Thibaut (1978), Van Lange, P.A.M. & D. Balliet. (2015), Campbell et al. (2015).

such an argument would look like without amounting to wishful thinking. So the argument for the pervasiveness of PML still holds.

*4.2 It All Evens Out in the End? Version II*

There is another version of this objection which, if successful, would provide a Master Argument that PML is not something to be worried about. This line of thought turns on the point that relationships are "a two way street".[19] Return again to Octavia and Solita. It's true that, *qua* moral patient, Octavia has a greater moral force than Solita. However, while in the present moment it is easy to forget, this actually comes with its own benefit for Solita: Solita's lower moral force comes along with less moral demands on her *qua* agent, since she has no sibling-related Associative Reasons, while Octavia has many. Insofar as relationships only "morally burden" as much as they morally benefit, perhaps things are, morally speaking, evened out. (Note that the objector is clearly *not* claiming that having a meaningful relationship is no better than not having one, since relationships have their own intrinsic benefits. So the "burden" in question here is purely in terms of one's moral obligations.)

This worry relies on the assumption that the Associative Reasons generated by relationships are consistently equal. This assumption is false. Consider the parent-child relationship, the mentor-mentee relationship, or the employee-employer relationship. There are many relationships which, by their nature, involve disparities in the Associative Reasons that they generate. This isn't to say that such relationships are not bidirectional: Of course a child has special duties with respect to her parent. (And at some points in the relationship, say with an aging parent and an adult child, it is the child that has much stronger duties than the parent.) But the weight and type of the Associative Reasons that being a child generates are much different than those of a parent. Similarly with many other types of relationships. And there is no reason to think that these will intuitively even out over time.

## 5. The Normative Puzzle of Patient Moral Luck

So far, I've defended the existence and pervasiveness of PML. My hope is that, already, you are beginning to feel some discomfort about this phenomenon. Associative Reasons are one thing, but if their distribution means that morality *encourages* deliberation which systematically treats some people as less worthy of consideration than others, that's *prima facie* problematic.

There are actually two puzzles that arise when considering PML. The first has to do with the systemic and predictable welfare consequences of the unequal distribution of moral consideration. The ways in which Octavia has people to turn to and Solita does not will

[19] I thank an anonymous referee for pressing me to say a bit more about this.

likely have recurring and potentially severe consequences for their respective welfare opportunities. Thought about in this way, the problem isn't exactly about PML *per se*, but about its predictable consequences. This resembles an objection first raised by Samuel Scheffler, who is concerned with the justification of Associative Reasons:

> Suppose, for example, that there are three individuals, A, B, and C, none of whom has any special tie or relationship to any of the others. Each has only general duties toward the others, which is to say that each's duties toward the others are distributed equally…Now, however, suppose that A and B…become members of some group of a kind that is ordinarily thought to give rise to associative duties. And suppose that C is not a member of this group, which we may call *The In Group*…A and B are now required to give each other's interests priority over the interests of C in a wide range of contexts…Clearly, then, the overall distribution of duty that now prevails is both inegalitarian and decidedly unfavorable to C.[20]

Scheffler appears to be as concerned with inequalities in "plain" luck that arise from special relationships. These are a result of what I am calling PML, but are not yet wholly "intrinsic" to morality. These inequalities in welfare opportunities are themselves a source of normative concern (as liberal egalitarians have long argued under the heading of "brute luck").

To see the unique puzzle of PML, compare two moral principles:

> *Equal Consideration* (EC) - All else equal, every human patient deserves equal moral consideration.
>
> *Equal Treatment* (ET) - All else equal, every human patient should have equal opportunity for welfare.

I take it that both of these principles are plausible, compatible, and even mutually supportive. However, they are separable - ET is given in terms of outcomes, while EC is given in terms of the level of moral consideration owed to each moral patient as such. It is difficult to motivate one without motivating the other, since meeting ET will almost always be a result of meeting EC. However, I think these are nonetheless distinct and each important in their own right. The distinction here is similar here to a point Thomas Nagel makes about the importance of the inviolability of rights:

> We can distinguish the desirability of not being tortured from the desirability of its being impermissible to torture us; we can distinguish the desirability of not being murdered from the desirability of our murder's being impermissible; we can distinguish the desirability of not being coerced from the desirability of its being impermissible to coerce us. These are distinct subjects, and they have distinct values. To be tortured would be terrible; but to be tortured and also to be someone whom it was not wrong to torture would be even worse.[21]

---

[20] Scheffler (1994), pp.9-10.
[21] Nagel (2007), p.111.

Return to the simplistic case of *Island*, where Natalie and Jamie are both in peril, and Reed, Jamie's brother, can at most save one of them. Now given a general situation in which someone can only save one of two *strangers*, we might say that the stranger who is not saved is unlucky, but has no grounds of complaint against the saver. Similarly, it would be odd for Natalie to complain about Reed choosing his sister over her. However, what also seems true is that Natalie can legitimately complain (to morality?) that her status as a human being was not even properly considered - her death was ensured, not by the fact that she was literally unsaveable, but by the fact that morality made it forbidden for her saving to even be considered. This will be especially true if Natalie and Jamie find themselves in similar (though perhaps less dire) circumstances quite often, and Natalie's interests are systematically and morally permissibly neglected in deliberation. The systemic moral justification of the discounting of her interests is its own harm. And, I claim, this would be true even setting aside the (also worrisome) non-moral downstream effects on Natalie's well-being.

In fact, even further along the lines of Nagel's distinction, on one standard and intuitive kind of claim about the morality of special relations, things are even worse for Natalie. Bernard Williams famously argued that a husband's choosing to save his wife and not a stranger from drowning reflects poorly on his character if his action was partly justified on the basis of his thought that *it is morally permissible to favor one's wife over a stranger*. This, Williams says, would involve "one thought too many". The decent person's motive for saving their wife is merely *because it is my wife*, says Williams. If one is convinced by Williams' argument here (at least once made precise[22]), it seems like it is blameworthy for the husband to even consider saving the stranger or considering her interests or well-being.

I want to be clear that I am not arguing that the Williams-influenced thought is mistaken - perhaps it *is* morally wrong to, even for a moment, consider rescuing a stranger over one's nearest-and-dearest. And thus, perhaps Reed is a wholly morally good person for unthinkingly rescuing his sister Jamie over Natalie. My claim, instead, is that, especially once ramified across all Associative Reasons, and given the imbalance in how those are distributed, this will result in distinctive moral harms to some moral patients. This doesn't reflect poorly on Reed, or on those with Associative Reasons generally. Instead, it suggests that, if morality is not going to endorse unevenly distributed moral force (which certainly looks like unfairness baked into morality), there must be more to the normative story.

As a one-off case here or there such as in *Island*, it may seem easy to shrug and not lose sleep over PML. However, once we move from these cases to the real world, we can see

[22] For some interpretations and precisifications, see Wolf (2012), Baron (2008), Smyth (2018).

that PML will entail these minor deficits in equal consideration will repeatedly occur over a lifetime, often and systematically on the same people. Recall Octavia and Solita. It is not as though this time around, Solita cannot appeal to siblings for financial assistance, but perhaps next time she will be able to: If Solita remains an only child, this reduction in moral force in comparison to Octavia is a *standing* fact about her moral status vis-a-vis others. Meanwhile, as long as she does not abuse her moral force, Octavia will always (or at least as long as they live) have siblings on which she exerts special sui generis moral force. And while I gave the case in terms of financial assistance for simplicity, there are any number of intangible forms of assistance that Octavia has access to that Solita does not in virtue of her moral force: Emotional support, career help, help with childcare, a protection against loneliness, life advice, and so forth. Solita may be able to compensate for some or all of these things, but the point remains - she must compensate, and she can only hope to do so through her ability to form as many and as equally meaningful relationships as to generate the same kinds of Associative Reasons. And some may attempt to compensate and fail through no fault of their own, as discussed above.

PML is not a minor quirk built out of exotic thought experiments, but is a widespread and perhaps even inevitable phenomenon. But it disproportionately affects some moral patients, causing them extended and repeated lack of equal moral consideration even assuming that agents are acting exactly as they are morally required to.

As I noted early on in the paper, the existence and extent of *Associative Reasons* as common sensically understood are perfectly compatible with *EC* (as well as *ET*), given the latter's "other things being equal" clause. However, PML *as a normative problem* appears to arise when the ways in which applying this ceteris paribus clause leads to the systemic and pervasive ways in which "other things" never *are* equal. We can formulate the principle behind this intuition as follows:

> *Principle of Fixed Total Force* (PFTF)*:* While in any given moment, *equal consideration* can be trumped by other considerations, such considerations cannot generalize in such a way that different moral patients, for morally irrelevant reasons, have radically unequal global moral force spread across all agents over the course of an entire life.

The intuition behind PFTF is really the same as that behind the original, common sense *Equal Consideration*, although framed in the terminology of moral force. PFTF is the application of the egalitarian principle that people do not deserve unequal shares due to morally arbitrary factors to the good of moral consideration.[23] We might say that PML is an issue, then, of *Considerative Justice*. The difference is that PFTF is necessary to avoid ways in which the

[23] See, e.g. Temkin (1986),101, Rawls (1999), p.64.

scales can be tipped away from equal consideration across time and space, as in cases of Patient Moral Luck. And while *Equal Consideration* is compatible with *Associative Reasons*, PFTF doesn't seem to be,[24] at least once we note the empirical fact that not all moral patients have the same amount of total moral force as a result of their force-generating relationships. This is a genuine puzzle, and while it is theoretical, it isn't purely theoretical, since it suggests a change in our normative theory which may have effects on how we think of Associative Reasons and/or the general moral principles that surround moral patienthood.

## 6. Resolving the Puzzle

My main goal in this paper is to defend the existence and genuine puzzling nature of Patient Moral Luck. But it is also worth sketching some of the most promising ways to address the puzzle. As far as I can see, there are three possible kinds of response, which I'll call *Denial, Radical Revisionism,* and *Moderate Revisionism*. The last is the more interesting, so after briefly touching upon the first two, I'll discuss what I take to be the two most promising (and mutually compatible) versions of *Moderate Revisionism*, which I call the *Dinner Party Principle* and the *Principle of Moral Force Conservation*.

### *6.1 Denial and Radical Revisionism*

As its name suggests, the *Denial* response is more or less just to deny that PML is anything worth worrying about, that it is not a reason to think that our commonsense moral theory is in any trouble. For the fan of *Associative Reasons*, this would require an argument that our *Equal Consideration* intuitions don't actually support PFTF (or anything like it). On the other hand, there is also a moral theory with a rich and storied tradition which completely avoids the problem of PML: Agent-Neutral Consequentialism (ANC). PML can only get off the ground if we accept the existence of (ineliminable) Associative Reasons, and ANC doesn't accept that special relations can in themselves ground moral reasons.[25] I can't rule out these possibilities, but as they are each well-worn ground, I will focus instead on the two less revisionary accounts.

### *6.2 Moderate Revisionism*

Suppose you take PML seriously, but you aren't attracted to a purely agent-neutralist view. One strategy for reducing or eliminating PML might be to add new restrictions on the permissibility conditions for forming new relationships, given the implications of such new relationships on other moral patients' total moral force. A second strategy would involve some kind of *moral force prioritarianism,* according to which, when other things are equal, we

[24] But see section 6.4.

[25] I say "ineliminable" and "in themselves" here because proponents of ANC often try to accommodate our intuitions of partiality. This is a strategy that goes all the way back to Mill in *Utilitarianism* (Mill 1861), but has also been defended by Sidgwick (1907), Railton (1984), and Arneson (2003), among others.

give more weighty consideration to those with fewer reason-generating relationships. These are both odd sounding proposals at first blush, but, while there remains questions about how to apply them in practice, they can both be independently motivated. I'll consider each in turn.

*6.3 The Dinner Party Principle*

A first strategy for reducing or eliminating PML would be to add rules governing the permissibility of forming new (Associative Reason generating) relationships. Insofar as good friendships and relationships are built on meaningful connections, and we cannot simply *choose* to have a meaningful connection with any other given person, this may seem like a perverse sort of duty. Furthermore, one might claim that building a relationship with someone in part because of PML is itself immoral, even a form of moral fetishism. But these objections would be to misunderstand the proposal. The imperfect duty that I am proposing is not about directly choosing who to form meaningful connections with, but rather is about putting ourselves in a position to make it more likely that we will form meaningful connections with those who otherwise have less. To illustrate, consider an unrealistically simplified case:

> *Party Choice.* Shira has been invited to two dinner parties, *A* and *B*. Her prudential reasons for going to *A* and going to *B* are equally weighty - an equal amount of people she knows and likes will be at each of them, they should be equally fun, etc. However, she has good reason to believe that, of the people she doesn't know at each party, party *A* will have far more people with less meaningful relationships than party *B*, through no fault of their own.

According to this response to PML, Shira has a duty, perhaps an imperfect one, to go to party *A*, because it puts her in a position to increase the chances of reducing PML.

Of course, once we move away from the simplified and idealized case, how to meet this imperfect duty will become quite messy. However, I don't think the general idea is at all implausible. Imagine you are discussing with your partner who to invite to a dinner party, and you only have room for eight people. You have seven people on your guest list and are trying to decide on an eighth. Your partner says "Why don't we invite Miriam, your new colleague? I am sure she would love the opportunity to make some friends since she is new in town." Surely this is *a* sensible consideration in favor of inviting Miriam. On the current proposal, your partner's suggestion is not just supererogatory, but is an imperfect duty.[26] That's revisionary, but not radically so. Let's call this the *Dinner Party Principle*:

[26] The connection between imperfect duties and supererogation is conceptually complicated, but hopefully this is clear enough for present purposes. See Heyd (2019), Section 2.

> *The Dinner Party Principle.* We have a standing pro tanto duty to increase the likelihood of relationships forming which will reduce the amount of PML in the world.

I find this idea plausible and compelling, but it is quite restricted in scope. We can put people in a position to make friends, we can even try to set people up for romantic partnerships, but we can't make someone have more siblings (except in the case where we are their parents!). Perhaps, though, if we think of meaningful relationships of certain sorts as compensating for those that we may otherwise lack (consider the phenomenon of a "chosen family"), it is less restrictive than it at first seems.

*6.4 The Principle of Moral Force Conservation*

Return again to my variant of Equal Consideration, the *Principle of Fixed Total Force:*

> *Principle of Fixed Total Force* (PFTF)*:* While in any given moment, *equal consideration* can be trumped by other considerations, such considerations cannot generalize in such a way that different moral patients, for morally irrelevant reasons, have radically unequal global moral force spread across all agents over the course of an entire life.

As I said above, PFTF looks incompatible with the unequal distribution of relationships that generate Associative Reasons, which generates the normative puzzle of PML. However, while this is a natural conclusion to draw, it is not actually entailed by PFTF even when combined with a common sense view of Associative Reasons. The initially intuitive way to think about a given moral patient's total moral force is as follows: She begins with a certain strength of force that binds every rational agent merely in virtue of the fact that she is the kind of thing that has a moral status which can generate impartial reasons on anyone. As she forms relationships, the thought goes, she gains new sui generis relations of moral force which are binding on particular agents - namely, those she is in relationship with. This violates PFTF if it is unequal across patients and lives, as it likely will be.

However, there is another way of thinking about what happens when new relationships are formed. Instead of thinking of relationships as generating *new* moral force, we may think of it as restributing the fixed amount of total moral force that a given moral patient has. In a vacuum, her moral force gets distributed over all moral agents in her environment. However, as relationships form, her moral force gets clumped into particular agents, thus reducing the charge distributed over each other agent (even if very slightly). The more special relationships some patient is in, the more clumped their moral force is into these relationship partners and less onto others. This generates small increases in the weight of reasons placed on those deliberating about or between two strangers, as the moral force

each of them bears on her will be slightly different. Call this idea the *Principle of Moral Force Conservation*.

The *Principle of Moral Force Conservation* corrects for PML by adjusting the varying strength of a patient's moral force on strangers rather than attempting to eliminate or reduce the strength of moral force generated by relationships. Or put in terms of agents, it makes the strength of our *agent neutral* reasons vary in part on the basis of facts about moral patients' distribution of agent relative moral force. Such an account will have the surprising implication that, all else being equal, we should be slightly partial to a stranger with less special relationships than one with more. (This will be compatible with *Equal Consideration* as originally stated as well, since it contains the 'other things being equal' clause and this is a morally relevant respect in which things are not equal, given PML.) Of course, in the real world, this will often be unknown, so defaulting to equal consideration will be the appropriate deliberative approach. Nonetheless, when these are features that can be known, the principle implies that you should take these kinds of complications into consideration.

I am not sure whether the *Dinner Party Principle* or the *Principle of Moral Force Conservation* are true (and notice that they are compatible with each other). This would require further exploration and argumentation. However, they do show that the problem of PML doesn't in principle require giving up on unequally distributed Associative Reasons.

## 7. Taking Stock

In this paper, I've argued for a puzzle that arises from two principles of common sense moral theory. Patient Moral Luck arises once we think about the implications of applying our common sense principles across cases. Equal consideration for moral patients may plausibly be bracketed for the purposes of one off cases of partiality, such as loaning a sibling-in-need rent money, even when one would not do so for a stranger. However, the systematic application of this kind of reasoning leads to unacceptably unequal levels of total moral charge across different moral patients, as a result of factors beyond their control. I take it that this puzzle has not been explicitly noticed in the literature up until now for two reasons: First, normative theory is often excessively focused on *agents* as such: What reasons they have, what is permissible for them to do, what is obligatory, etc. Patient moral luck is only easily noticeable once we look at morality from a patient-centric perspective by adopting the notion of "moral charge". Second, that patient moral luck is a *normative problem* is really only noticeable once we look past individual cases to a systematic assessment of moral patients' total moral charge over space and time. Once we keep these two things simultaneously in mind, PML begins to come into view as a serious normative issue.

My primary aim in this paper was to elaborate on the issue of PML and to defend it as a serious normative puzzle. However, I finished, in section 6, by considering two principles that may help to reduce or even eliminate PML - the *Dinner Party Principle* and the *Principle of Moral Force Conservation*.

Patient Moral Luck is a significant problem arising out of widely accepted moral principles, applied across real world contexts. It's easy to miss this if we don't take a wide enough view of the systemic implications of our moral theorizing, instead focusing on particular cases or sets of cases. My hope is that this paper will stimulate thinking about how to address the puzzle of PML while retaining what is so intuitive about both partiality in relationships as well as equal consideration for all moral patients.

**Works Cited:**


Adler, David & Laura Adler (2018). "How Parents Kill Progress", in *Current Affairs.* Available at: https://www.currentaffairs.org/2018/06/how-parents-kill-progress (Accessed November 12, 2023)

Arneson, Richard. (2003). "Consequentialism vs. Special-Ties Partiality", in *The Monist* 86(3):382-401.

Ashwell, Lauren. (2010). "Superficial Dispositionalism", in *Australasian Journal of Philosophy*, 88(4):635-653.

Baron, Marcia. (2008). "Virtue ethics, Kantian ethics, and the 'one thought too many' objection", in Monika Betzler (ed.), *Kant's Ethics of Virtues.* De Gruyter, pp.245-278.

Broome, John. (1990). "Fairness", in Proceedings of the Aristotelian Society (Vol. 91, pp. 87-101). Aristotelian Society, Wiley.

Campbell, Kelly, Nicole Holderness, & Matt Riggs. (2015). "Friendship chemistry: An examination of underlying factors", in *The Social Science Journal*, 52(2), pp.239-247.

Choi, Shungho & Michael Fara. (2018). "Dispositions", in *Stanford Encyclopedia of Philosophy*, Available at: https://plato.stanford.edu/entries/dispositions/ (Accessed November 12, 2023).

Ciulla, Joanne B. (2005). "In Praise of Nepotism?", in *Business Ethics Quarterly*, 15(1):153-160.

Enoch, David and Andrei Marmor. (2007). "The case against moral luck", in *Law and Philosophy*, 26:405.

Enoch, David. (2019). "Playing the Hand You're Dealt: How Moral Luck is Different from Morally Significant Plain Luck", in *Midwest Studies in Philosophy*, 43(1):257-270.

Flanagan, Caitlin. (2021). "Private Schools Have Become Truly Obscene", in *The Atlantic*, April 2021. Available at: https://www.theatlantic.com/magazine/archive/2021/04/private-schools-are-indefensible/618078/ (Accessed November 13, 2023)

United States Congress Committee on Ethics. (2022). *House Ethics Manual.* Available at: https://ethics.house.gov/sites/ethics.house.gov/files/documents/Dec%202022%20House%20Ethics%20Manual%20website%20version.pdf (Accessed November 13, 2023)

Kelley, Harold & John Thibaut. (1978). *Interpersonal relations: A theory of interdependence.* New York, NY: John Wiley & Sons, Inc.

Krantz, Murray. (1987). "Physical attractiveness and popularity: A predictive study", in *Psychological Reports*, 60(3):723-726.

Mason, Andrew. (2006). *Levelling the Playing Field: the Idea of Equal Opportunity and its Place in Egalitarian Thought*. Oxford University Press.

Mill, John Stuart. (1861). *Utilitarianism.*

Nagel, Thomas. (1976). "Moral Luck", in *Aristotelian Society Supplementary Volume* 50:115-151.

Nagel, Thomas. (2007). "The value of inviolability", in Paul Bloomfield (ed.), *Morality and Self-Interest.* Oxford University Press.

Nelkin, Dana. (2019). "Moral Luck", in *Stanford Encyclopedia of Philosophy.* Available at: https://plato.stanford.edu/entries/moral-luck/ (Accessed November 13, 2023).

Nozick, Robert. (1974). *Anarchy, State, and Utopia.* Basic Books.

Portmore, Douglas. (2011). *Commonsense Consequentialism.* Oxford University Press.

Railton, Peter. (1984). "Alienation, consequentialism, and the demands of morality", in *Philosophy and Public Affairs*, 13(2):134-171.

Rawls, John. (1999). *A Theory of Justice*, revised ed. Oxford: Oxford University Press.

Reeves, Richard. (2018). *Dream Hoarders: How the American Upper Middle Class is Leaving Everyone Else in the Dust, Why That Is a Problem, and What to Do About It*. Brookings Institution Press.

Scheffler, Samuel. (1994). *Families, nations, and strangers.* Lindlay Lecture, University of Kansas, Department of Philosophy.

Segall, Shlomi. (2011). "If you're a luck egalitarian, how come you read bedtime stories to your children?", in *Critical Review of International Social and Political Philosophy*, 14(1):23-40.

Sidgwick, Henry. (1907). *The Methods of Ethics. 7th Edition.* Macmillan.

Smyth, Nick. (2018). "Integration and authority: rescuing the 'one thought too many' problem", in *Canadian Journal of Philosophy*, 48(6):812-830.

Sobel, David. (2013). "Self-Ownership and the Conflation Problem", in Mark Timmons (ed.) *Oxford Studies in Normative Ethics: Volume 3.* Oxford University Press.

Swift, Adam. (2003). *How Not to be a Hypocrite: School Choice for the Morally Perplexed Parent.* Routledge.

Temkin, Larry S. (1986). "Inequality", in *Philosophy and Public Affairs*, 15:99–121

Van Lange, P.A.M. & D. Balliet (2015). "Interdependence theory", in *Theories of Social Psychology*, 251.

Vallentyne, Peter, Hillel Steiner, & Michael Otsuka. (2005). "Why left-libertarianism is not incoherent, indeterminate, or irrelevant: A reply to Fried", in *Philosophy & Public Affairs*, 33(2):201-215.

Williams, Bernard. (1976). "Moral Luck", in *Aristotelian Society Supplementary Volume* 50:115-151.

Williams, Bernard. (1981). *Moral Luck: Philosophical Papers 1973-1980.* Oxford University Press.

Wolf, Susan. (2012). " 'One Thought Too Many': Love, Morality, and the Ordering of Commitment", in *Luck, Value, and Commitment: Themes from the Ethics of Bernard Williams*, pp.71-92.