

## Medical Error and Moral Luck

**ABSTRACT.** Medical error is the third-leading cause of death in the United States, but there has been little work done on the associated conceptual and normative questions. What *is* medical error? Is all medical error *bad*? The first section of this paper surveys the dominant conception of medical error—promulgated by the Institute of Medicine—and tries to understand whether error necessarily eventuates in adverse events. The second section challenges an asymmetry in the way that we think about error: For example, the received view would allow that undertesting could comprise medical error, whereas overtesting cannot. The third section considers the concept of moral luck and how it bears on our ascriptions of medical error.

### INTRODUCTION

**T**his special issue on ethics and error in medicine reinvigorates a conversation that has been substantially dormant for twenty years. The papers in this issue elaborate and update that conversation in significant ways, particularly with regard to vulnerable populations and the epistemology of medical error. But this first paper is largely conceptual, laying out the motivation for caring about medical error in the first place, exploring what medical error is, and proposing a moral framework to help us think about it. This paper therefore sets up those that follow—while, at the same time, remaining largely neutral about the substantive views advanced by those authors. The papers are therefore complementary and form a useful whole that can be explored on various levels. A new book (Allhoff and Borden forthcoming) considers additional avenues, while at the same time recognizing the centrality of the papers featured in this issue.

### MEDICAL ERROR

Two decades ago, the Institute of Medicine released a report, *To Err Is Human: Building a Safer Health System* (Kohn, Corrifan, and Donaldson

2000).<sup>1</sup> The report noted that at least 44,000—and perhaps as many as 98,000—Americans died each year as a result of medical errors (American Hospital Association 1999, 1; Kohn, Corrigan, and Donaldson 2000, 1). At the time, this would have been the eighth-leading cause of death, comprising more deaths than other high-visibility causes, such as motor vehicle accidents (43,000), breast cancer (42,000), or AIDS (17,000) (Kohn, Corrigan, and Donaldson 2000, 1; Centers for Disease Control and Prevention 1999, 6). In the intervening years, medical error has moved up the charts, now comprising the third-leading cause of deaths in the United States (Makary and Daniel 2016; Allen and Pierce 2016). The overall number of deaths from medical error in 2016 was over 250,000, which was 9.7% of deaths nationwide (Makary and Daniel 2016), over double what had been reported in 1999.<sup>2</sup>

But what is lost in these statistics is how medical error is *defined*; obviously the numbers have to be predicated on some antecedent definition, lest it be unclear what actually counts as a death due to medical error. The Institute of Medicine defines error “as the failure of a planned action to be completed as intended or the use of a wrong plan to achieve an aim” (4). The report goes on to note that “errors depend on two kinds of failures: either the correct action does not proceed as intended (an error of execution) or the original intended action is not correct (an error of planning)” (4). It further emphasizes that error “can happen in all stages of the process of care, from diagnosis, to treatment, to preventive care” (4).

The report then goes on to build a taxonomy of medical error that tracks these three categories. Diagnostic error includes error or delay in diagnosis, failure to employ indicated tests, use of outmoded tests or therapy, and failure to act on results of monitoring or testing. Treatment error includes error in the performance of an operation, procedure, or test; error in administering the treatment; error in the dose or method of using a drug; avoidable delay in treatment or in responding to an abnormal test; and inappropriate (not indicated) care. Preventive error includes failure to provide prophylactic treatment and inadequate monitoring or follow-up of treatment. A fourth catchall category, “other,” includes failure of communication, equipment failure, and other system failure (36).

The next section will problematize these categories, but let me make some further comments on the Institute of Medicine’s approach in the remainder of this section. First, the report acknowledges that not all errors result in actual harm (2). This is an important observation, and one worth lingering on. Say, for example, a healthcare provider<sup>3</sup> prescribes

penicillin to a patient with a known penicillin allergy, but the patient never takes the penicillin. Maybe she never fills the prescription; maybe she loses it after having it filled; maybe she feels better by the time she gets home and elects to not use it at all; or some other change occurs. The analysis here is that the provider *erred*, but that error did not eventuate in harm to the patient. We can easily contrast that patient—and indeed we will—to another who both received the errant prescription *and* took it, suffering anaphylaxis. What contrasts these patients is not the error, but whether there was an “adverse event” (i.e., bad outcome) following the error (Kohn, Corrigan, and Donaldson 2000, 4).<sup>4</sup> What we ultimately care about is adverse events, not errors per se. Or, to put it another way, the principal reason that we care about errors is that they tend to lead to adverse events. But if a provider committed an isolated error that did not lead to such an adverse event—and setting aside any propensity for future recurrence—it just seems like something closer to good luck than to pernicious error.

Second, one of the major contributions of the report is to pivot from blaming *individuals* to blaming *institutions*. They put it this way:

The common initial reaction when an error occurs is to find and blame someone. However, even apparently single events or errors are due most often to the convergence of multiple contributing factors. Blaming an individual does not change these factors, and the same error is likely to recur. Preventing errors and improving safety for patients require a systems approach in order to modify the conditions that contribute to errors. People working in health care are among the most educated and dedicated workforce in any industry. The problem is not bad people; the problem is that the system needs to be made safer. (36)

This is important. We can think of some archetypical medical error, such as when a surgeon amputates the wrong leg. It should probably be the case that no surgeon, independently, exercises sole judgment over which appendage to amputate. So if the wrong leg does get amputated, it would not just be the surgeon who is implicated: The medical records could be incorrect, thus implicating any number of providers; the records should have been verified, thus implicating hospital oversight; the patient’s family could have noticed the wrong markings while the patient was being prepped for surgery; or, frankly, a huge range of other possibilities could have contributed to the error.

Even think of some particularly nefarious errors, like those effected by a drunk surgeon. Here it is perfectly reasonable to blame that surgeon, but

even in cases like this, there is plenty of blame to go around. How was the *hospital* configured such that a drunk surgeon got access to an operating room? Think of some fairly simple remediation, like putting a breathalyzer on the operating rooms' doors. Would that be a good idea? Did the surgeon have to sign a form, declaring lack of compromised abilities? If not, would the implementation of such a form promote better outcomes? These are just examples, but the simple point is that, even in the limited number of cases where we might attribute blame to “bad apples,” there are a range of institutional safeguards that could be in play. Thinking through error invites us to countenance all of these—some of whose costs could exceed the associated benefits—at myriad institutional levels. In fact, this sort of “systems-based thinking” is a manifest strength of the report, particularly insofar as it broadens the conversation beyond individual practitioners.<sup>5</sup>

Finally, it goes beyond present purposes to indicate, specifically, how we should deal with medical error.<sup>6</sup> But, in recognition that this discussion does not take place in a vacuum, let me indicate broad thinking on how to mitigate error. The Institute of Medicine proposes a four-tier approach that can be usefully summarized in this regard. First, it proposes a national focus to create leadership, research, tools, and protocols to enhance the knowledge base about safety (69–85). Second, it recommends identifying and learning from errors by developing a nationwide public mandatory reporting system and by encouraging healthcare organizations and practitioners to develop and participate in voluntary reporting systems (86–108). Third, it aims to raise performance standards and expectations for improvements in safety through the actions of oversight organizations, professional groups, and group purchasers of healthcare (109–131). Fourth, it suggests implementing safety systems in healthcare organizations to ensure safe practices at the delivery level (132–154).

#### ASYMMETRIC ERRORS

In order to motivate the discussion of medical error and moral luck that will come in subsequent sections, it is worth highlighting how the Institute of Medicine's conception of error is, in some key regards, asymmetric. In that regard, let me start with an example, use that example to motivate a contrast with fairly standard thinking in moral philosophy, and then return to the issue more directly. As the example, consider the diagnostic error described as “failure to employ indicated tests” (Kohn, Corrigan, and Donaldson 2000, 36). Suppose a healthcare provider is looking at potentially cancerous tissue for which a biopsy is the indicated procedure.

Further suppose that, for whatever reason, the provider fails to order the biopsy. Per above, we are going to call this an error whether or not it leads to an adverse event (i.e., patient harm)—simply because the test should have been ordered. If the failure to order it were to lead to the development of melanoma, that would be both an error and an adverse event. But even if the tissue were benign—but could have been malignant—there would have been a harm-free error.

On the other hand, though, suppose the provider orders the biopsy even when it is not indicated. Here we can imagine a patient coming to the clinic, fearing skin cancer. The patient has a mole on her forearm, but it is shallow, small in circumference, lacks irregular edges, there is no family history, and so on. So, after a visual exam, the provider deems skin cancer to be a low risk and the biopsy not to be medically indicated. Suppose, though, that the patient presses, persistently asking for the biopsy. Despite the procedure not being medically indicated, we can imagine the provider relenting for any number of reasons. Maybe the provider dislikes conflict. Maybe the patient is particularly strident or annoying. Maybe the provider's husband died of skin cancer last year. Maybe the provider gets to bill for the extra procedure, as well as follow-up care. It is not to say that any of these are good reasons, but it would hardly be surprising for any of them to drive clinical care.

So the provider does order the biopsy. Here is the important question, though: Is that a medical error? Can overtesting in general comprise error? On the Institute of Medicine's report, the answer is pretty clearly no. The reason is that the types of error that it countenances in this regard trade on a "*failure* to employ indicated tests." There simply is no prong on their analysis for "ordering non-indicated tests." There are two ways we could respond to this. The first is to draw a distinction between 'error' and something else, like 'mistake.' On this approach 'error' is a term of art, as defined by the report. And if we say overtesting is not—or *cannot*—be error, we are just tracking that definition. 'Mistake' would just pick out when providers get something wrong, whereas 'error' would be a subset of 'mistake,' and it would be circumscribed by the list in the report. So, to continue this example, overtesting and undertesting would, *ex hypothesi*, both be instances of mistakes, but only *undertesting* would be an instance of error.

There is nothing incoherent about this, but it seems curious. Specifically, it introduces a distinction between 'error' and 'mistake' that just does not track our ordinary usage of the terms—rather, we treat them as co-

extensive. For example, the Oxford English Dictionary defines ‘error’ as “[s]omething incorrectly done through ignorance or inadvertence; a *mistake*, e.g., in calculation, [judgment], speech, writing, action, etc.” (OED Online 2019, emphasis added). Conversely, it defines ‘mistake’ as “misapprehension, misunderstanding; *error*, misjudgment” (OED Online 2019, emphasis added).<sup>7</sup> So it seems fair to say that the report is, at best, being idiosyncratic. The second way to go, though, would be to say that the report’s conception is implausible: Both overtesting and undertesting are forms of error (i.e., not just mistakes), and we are being supplied with an analysis that is not just idiosyncratic, but instead downright silly. Is this just terminological, or is there anything that hangs on it, theoretically?

The biggest thing that hangs on it is the numbers. If we expand ‘error’ beyond the report’s more limited definition, there would obviously be more error, probably a lot more.<sup>8</sup> In other words, to stay with our example, we would have to figure out all the instances of overtesting, then add those back into the number of errors—whereas we have previously been excluding those. Suffice it to say that there is a lot of overtesting,<sup>9</sup> so the data is going to be affected by the definitions. There might even be optics problems as well: Raising the number of things that count as error could affect the public’s perception of the profession, which could have effects on those seeking treatment. That said, this limited definition should not necessarily be seen as a self-serving metric; maybe there are other theoretical virtues that count in favor of it.

Top among prospective candidates is the idea that overtesting and undertesting are differentially situated with regards to adverse events. So let us suppose that overtesting would rarely lead to an adverse event. We can suppose that needless x-rays, for example, would lead to increased exposure to radiation, and that additional quantum of radiation could be the difference-maker between whether the patient develops cancer or not (Harvard Women’s Health Watch 2018). Or we could suppose that needless blood draws could lead to injuries from the phlebotomist, such as nerve damage (Owen and Johnson 2017). But it is probably safe to assume that these interventions are *less* likely to eventuate in patient harms than, for example, failure to order an indicated biopsy for skin cancer. To be sure, there are many problems with overtesting, but they probably have more to do with healthcare economics than with adverse events (Golemon 2019).

However, even if this is true—and it probably is—why not talk in terms of ‘adverse events caused by error’ instead of redefining ‘error’ to track propensity for adverse events? Remember from the outset that the Institute

report *already* draws a distinction between error and adverse event: Not all error (even on its own circumscribed definition) eventuates in adverse events. In other words, we have at least the following categories from the Institute of Medicine taxonomy:

1. Error that leads to adverse events.
2. Error that does not lead to adverse events.
3. Non-error that leads to adverse events.
4. Non-error that does not lead to adverse events.

Set aside (4), which is not so interesting. If a healthcare provider fails to order an indicated biopsy and the patient dies of skin cancer, that is (1). If a provider fails to order an indicated biopsy and the patient does *not* die of skin cancer, that is (2). Suppose, for example, the patient *actually had* skin cancer, but inadvertently lost the affected arm in a chainsaw accident, before the cancer spread. Or suppose that, given the presentation of the skin defect, the patient would *usually* have skin cancer, but just happened not to in this case. Again, (2). But if a provider ordered a non-indicated x-ray and the patient developed cancer because of radiation exposure (i.e., [3]), that would not be error under the Institute of Medicine's definition.

This just does not make much sense. What we ultimately care about is if healthcare providers make errors—in the ordinary usage of the word—that eventuate in adverse events. *Those* are the things we should be counting, if we want to measure the significance of medical error. In this regard, overtesting is just as bad as undertesting (an Institute of Medicine error), at least conceptually; the only difference is that overtesting would be *less likely* to lead to adverse events than undertesting, as discussed above. But we need not sort that out just by defining 'error' to exclude overtesting altogether. Rather, we would just say something more straightforward, like that overtesting is not as bad as undertesting, in general, because less tends to hang on it this error than others.

Note, however, that we also want (2) cognized as medical error—which the Institute of Medicine report does. Conceptually, whether an error leads to an adverse event or not has nothing to do with whether it is an *error* in the first place. Of course errors that lead to adverse events are *worse* than errors that do not lead to adverse events, but, again, nothing hangs on that for our conception of error. Rather, we just say that some errors are worse than others, either in isolation or as statistical classes. We would still care about specific errors that did not lead to adverse events because, statistically, they might tend to. So again consider the provider

who failed to order an indicated biopsy, yet the patient lost the affected arm to a chainsaw accident. The provider is not vindicated simply because skin cancer never developed. Part of the explanation there would be that, if this error were committed multiple times, skin cancer *would* develop in many of them; there are not so many chainsaw accidents, for example, such that the provider is usually off the hook.

In summary, the central point is that the Institute of Medicine approach rules out certain sorts of error that, at least intuitively, it should not. The Institute of Medicine enumerates several classes of error, and if a putative error does not fall within one of those delineated classes, it is not an error, at least by their lights. This is problematic because the list is underinclusive, or at least it is underinclusive as against commonsense definitions of ‘error.’ Second, the Institute of Medicine definition problematizes error in the sense that it separates error from adverse event. And, insofar as that then makes it conceptually possible to have errors that do not lead to adverse events, it is not immediately clear that we are counting the things we should be counting. Instead, we most proximately care about errors that *do* lead to adverse events—or at least that *tend* to. The report does not commit a conceptual mistake, but rather simply admits of more errors than we necessarily care about. That can be remedied, though, just by focusing on the relevant subset (e.g., errors with adverse events, rather than all errors). By contrast, the conceptual mistake that bars certain errors from counting as errors at all is more problematic. It is not clear what this has going for it—nor does the report wrestle with the implications—and it ignores a relevant source of error. Moving forward, I shall therefore revert to the more common usage of ‘error,’ under which an error is any sort of mistake that healthcare providers make. Those errors are not constrained to an enumerative list, but rather could be virtually anything. I will also assume that an error does not require an adverse event—in this regard agreeing with the Institute of Medicine—and will now pivot from this conceptual interrogation into its normative underpinnings.

#### MORAL LUCK

In this undertaking, it will be useful to review the doctrine of moral luck. Much has been written on this topic,<sup>10</sup> which originates with Immanuel Kant:

A good will is not good because of what it effects or accomplishes, because of its fitness to attain some proposed end, but only because of its volition, that is, it is good in itself. [...] Even if, by a special disfavor of fortune or

by [mother nature], this will should wholly lack the capacity to carry out its purpose—if with its greatest efforts it should yet achieve nothing and only the good will were left [...] then, like a jewel, it would still shine by itself, as something that has its full worth in itself. Usefulness or fruitlessness can neither add anything to this worth nor take anything away from it. (Kant 1998, 50)<sup>11</sup>

The contemporary conversation owes to Thomas Nagel (1979) and Bernard Williams (1981), though Nagel is more accessible.<sup>12</sup> Specifically, he notes that “[w]here a significant aspect of what someone does depends on factors beyond his control, yet we continue to treat him as an object of moral judgment, it can be called moral luck” (32).

One reason this becomes a compelling problem in moral philosophy is that we typically “seem to be committed to the general principle that we are morally assessable only to the extent that what we are assessed for depends on factors under our control” (Nelkin 2019). Nevertheless, we routinely assess moral agents *despite* the existence of factors outside their control. For example, consider this thought experiment by Jeremy Waldron:

Two drivers, named Fate and Fortune, were on a city street one morning in their automobiles. Both were driving at or near the speed limit, Fortune a little ahead of Fate. As they passed through a shopping district, each took his eyes off the road, turning his head for a moment to look at the bargains advertised in a storefront window. (The last day of a sale was proclaimed, with 25 percent off a pair of men’s shoes). (Waldron 1995, 387)

From a moral perspective, it seems that Fate and Fortune are similarly situated. But further suppose the following:

In Fortune’s case, this momentary distraction passed without event. The road was straight, the traffic in front of him was proceeding smoothly [...] and he completed his journey without incident. Fate, however, was not so fortunate. Distracted by the bargain advertised in the shoe store, he failed to notice that the traffic ahead of him had slowed down. His car ploughed into a motorcycle ridden by [...] Hurt. Hurt was flung from his motorcycle and gravely injured. (1995, 387)

And now it suddenly looks like they are differently situated, specifically with regards to the injury Fate causes to Hurt. The issue is that injury has nothing to do with anything under the control of Fortune or Fate. Rather, what sets them apart is simply whether there was a slowed motorcycle in front of one of them, which is something completely outside of their control.<sup>13</sup>

So the problem is that we either have to give up on the control principle—thus morally assessing people for features outside their control—or else we have to retreat to moral parity between Fortune and Fate. The latter may sound more initially plausible, but it quickly runs aground. Specifically—and this is why I chose Waldron’s formulation—consider tort law. If the purpose of tort law is to promote something like rectificatory justice (i.e., restoration of the *status quo ante*), then what?<sup>14</sup> In Fortune’s case, there is no injury, and so no rectification is needed. In Fate’s there is, and it hardly makes any sense to foist the burden of that injury onto Hurt. In one case, we need to collect damages, to ensure that Hurt does not bear the costs of Fate’s negligence. In the other, there is nothing for Fortune to pay for. So parity—e.g., between carelessness and damages—is asymmetrical. Waldron proposes to resolve this in a creative way, namely charging a tax to *all drivers*, sort of like a nationalized insurance scheme (1995, 389).<sup>15</sup> That would equalize tort liability across both drivers, but the moral intuition lingers that Fortune is differently situated in virtue of *causing harm* that Fate did not cause. To credit that intuition, though, is to relent on the control principle that we otherwise take to be morally perspicuous.

Let us now pivot and approach all this through the lens of medical error. Specifically, we can readily acknowledge that our moral assessments of error can be colored by exogenous environmental factors (e.g., the presence of motorcycles)—as opposed to endogenous agential factors (e.g., the driver’s negligence)—owing to moral luck. To use the language developed in previous sections, suppose that a healthcare provider errs, but further suppose that whether that error results in an adverse event is wholly predicated on luck. Return to the case where the provider should have ordered a biopsy, but failed to. Because this is an instance of undertesting, it is properly regarded as medical error (i.e., under the Institute of Medicine classification). But whether we “care” about the error ultimately depends on non-agential features of the case, like if the patient has an errant chainsaw accident, rendering the failure to resect the cancer—which would only have been revealed through the neglected biopsy—medically irrelevant.

Again, I only mean irrelevant in the instant case. If a healthcare provider systematically fails to order indicated tests, at least some of those omissions will eventuate in adverse events. Therefore, we can identify a critical locus of moral evaluation, namely, whether actions (or omissions), when common, *tend to* lead to adverse events. Or even suppose “tend to” is too strong: Maybe the testing practices of some provider increase the incidence of

some particular adverse event from one to two percent. Here it might not be accurate to say that anything “tends to” happen (i.e., since two percent is still something you would rationally bet against). But it would be accurate to say something else, like that the probability of some adverse event is higher given non-testing than it would have been under testing. And this fact alone gives us moral purchase on the failure to order the test.<sup>16</sup>

Of course we can say the same thing in the Fate/Fortune case: We do not morally exculpate Fate simply because no motorcycle was in front of him. Rather, we say that he “got lucky,” that there easily *could have been* a motorcycle, and that, if he is routinely negligent, he will eventually hit a motorcycle—or something else, like a pedestrian. The associated probabilities between negligence causing a traffic accident and an unordered biopsy failing to prevent a death by melanoma could obviously be different. But, structurally, the cases are isomorphic. At least they seem to be: The potential difference-maker is that we might be inclined to categorize the Fate/Fortune case as trading on actions (e.g., hitting motorcyclists), whereas the unordered biopsy looks closer to an omission. But without wading too deeply into the act/omission literature, we can make a few ready observations.<sup>17</sup> Chief among them is that this is probably a mischaracterization of Fate/Fortune. The principal problem with Fate here is that he *failed to exercise due care* and that failure *caused* the injury.

It might seem odd to talk about omissions being causal. It is not, however. For example, suppose “Barry promises to water Alice’s plant, doesn’t water it, and that the plant then dries up and dies. Barry’s not watering the plant—his omitting to water the plant—caused its death” (McGrath 2005, 125). The principle is also widely endorsed in tort law.<sup>18</sup> The point is simply that moral luck can attach to both acts and omissions; it would be an impoverished account if it only attached to the former. All that moral luck need be committed to is that we seem to morally assess agents, at least in part, on factors beyond their control. That basic claim does not differentiate between acts and omissions, nor need it. Insofar as some of our primary examples (e.g., undertesting) are necessarily omissions, this is an important result.

Thus, whether we are considering acts or omissions, our ascriptions of medical error will likely be imbued with moral luck. If a provider errs in a way that does not eventuate in an adverse event, it could just be an instance of the provider getting lucky. This is not necessarily the same as the *patient* getting lucky, because various unrelated adverse events to

the patient's health might screen off the provider's error: The chainsaw accident superseding the melanoma would be such an example. But other more quotidian examples would make the same point, such as the provider failing to diagnose prostate cancer—where prostate cancer generally afflicts older men who might well die of something else before the prostate cancer becomes lethal. Because the other affliction precludes the provider's error from eventuating in an adverse event, that error remains—or at least could remain—hidden. And yet it still seems like we would want to say that an oncologist who systematically underdiagnoses prostate cancer is a *bad* oncologist, even if nothing ever seems to hang on it. Or, to put it another way, it seems fraught to say that she is a bad oncologist *only if* her underdiagnoses ultimately kill the patient.

We have yet to explore another way in which moral luck maps onto medical error. Suppose that the provider orders an unindicated test which, by sheer luck, happens to reveal a lethal pathology. For example, suppose the provider should not have ordered an MRI, but that the MRI ultimately reveals an unindicated but malignant tumor. Again, under the Institute of Medicine approach, ordering unindicated tests cannot be error at all, but set that aside for reasons previously discussed. Rather, supposing that it was an errant order, it would equally seem odd to praise the provider for discovering the tumor. The errant order is a straightforward cause of the discovery, but not all causes are morally laudable. For example, if I get in a bar fight and punch out someone's infected tooth, it hardly seems I deserve any moral credit, despite the fact that the punch might have made the recipient better off. Instead, I am morally blameworthy for getting into a fight, regardless of the upshot. In the same regard, the provider who misorders a test is also blameworthy, for failing to follow evidence-based practices. There is nothing problematic with saying that, sometimes, bad decisions eventuate in good outcomes; any of us can probably imagine myriad cases of practical deliberation where we have erred, yet gotten lucky. The point is simply that we should not conflate the conceptual analysis with the moral one.

#### FURTHER DIRECTIONS

The literature on medical error—particularly its ethical dimensions—has been largely underdeveloped, a shortcoming this issue and future work (Allhoff and Borden forthcoming) aim to remediate. This paper diagnoses the extent of medical error, and it updates the conversation in light of the two decades that have passed since the Institute of Medicine

report (Kohn, Corrigan, and Donaldson 2000). The paper also explores conceptual issues relating to the definition of ‘error,’ particularly as relates to (1) the relationship between errors and adverse events; (2) asymmetries in how certain sorts of error—for example, under- and over-testing—are commonly conceived; and (3) argues for a broader conception of error than is usually proffered (e.g., by allowing that overtesting can still comprise error). The last section then considers ways in which moral luck bears on our assessments of error, noting that our moral categorization of error—or lack thereof—often trades on features beyond practioners’ control. Insofar as this is true, it raises problems for commonsense conceptions of morality.

Beyond that, this paper sets up a broader dialogue on ethics and error. Some of that dialogue continues in the papers to follow. For example, one of those papers considers ways in which medical error particularly affects vulnerable populations (Golemon 2019). Others discuss additional epistemic issues relating to error (Peña-Guzmán and Reynolds 2019), as well as the broader historical context (Larsen 2019). Beyond these papers, new work makes additional contributions (Allhoff and Borden forthcoming), and so we hope that a debate that has been largely moribund for two decades emerges reinvigorated, with further lanes of inquiry identified, as well as substantive contributions to those inquiries.

#### ACKNOWLEDGEMENTS

Parts of this paper were written while I was a Fellow in the Center for Law and the Biosciences at Stanford University and a Fulbright Specialist in the Faculty of Political Science at the University of Iceland; I thank those institutions for their support. I also thank Sandra L. Borden, Rebecca Kukla, and Luke Golemon both for helpful conversations and for comments on earlier drafts.

#### NOTES

1. An executive summary recapitulates the key findings. Available at <http://www.nationalacademies.org/hmd/~media/Files/Report%20Files/1999/To-Err-is-Human/To%20Err%20is%20Human%201999%20%20report%20brief.pdf>.
2. Some of the discrepancy might be due to underreporting on the earlier study (Makary and Daniel 2016). So we need not conclude that the actual rates of death due to medical error doubled, though there is at least plausible evidence for an increase. It seems safe to assume that the Institute of Medicine’s stated goal of “a 50 percent reduction in errors over five years” (i.e., from 1999–2004) was not realized.

3. For the purposes of this paper, I will use “healthcare provider” when there might be myriad functional roles (e.g., physician, physician’s assistant, nurse) that could err in relevant ways. If the instance under discussion requires a more narrow functional role (e.g., surgeon, as a subset of physician), then I will be more specific. This approach is conceptually motivated, but it is also meant to acknowledge the significance of functional roles beyond physicians, and to serve as a corrective against the fact that bioethics, as a discipline, has overemphasized physicians and underemphasized a broader and more diverse (e.g., by gender, by race) healthcare profession.
4. I will follow the literature and use ‘adverse event,’ but the less technical—and more colloquial—‘bad outcome’ can be used interchangeably.
5. See also Garrett and McNulty (forthcoming).
6. See Allhoff and Borden (forthcoming) for more discussion.
7. The semicolon makes the parsing interesting, apparently linking misapprehension/misunderstanding and error/misjudgment separately. Other definitions of ‘mistake’ in the OED make it look cognitive (e.g., relating to judgment), as opposed to action-oriented, whereas the entry on ‘error’ seems to conceive of ‘mistake’ more broadly (e.g., by enumerating a list that includes “calculation, speech, writing, *action*, etc” [emphasis added]). For present purposes, nothing hangs on this.
8. Because the statistics on medical error track the report’s definition, it would be hard to be more informative about how much more error there would be on a different definition.
9. See, e.g., Golemon (2019), for a useful survey of the literature.
10. For seminal works, see Statman (1993). For an overview, see Nelkin (2019).
11. Quoted in Nelkin (2019).
12. Williams’s position is difficult to understand; even he acknowledges that his original article “may have encouraged” some misunderstandings (1993, 251).
13. Nagel (1979) uses a similar—if less descriptive—example. Nagel also classifies four different kinds of moral luck: resultant luck, circumstantial luck, constitutive luck, and causal luck. This is an instance of resultant luck. Circumstantial luck trades on the circumstances in which we find ourselves: Some Nazis might be blameworthy for atrocities, but had they been differently situated (e.g., born in Argentina, instead of Germany), they might not have participated in those atrocities at all. Constitutive luck recognizes that moral agents differ in their dispositions and traits. For example, some of us are braver or more cowardly than others, and those traits—over which we may have little control—can affect how we are morally evaluated. (Of course others, like Aristotle, might be more sanguine about the axes of control we

can exert in this regard. See Aristotle [2009, 1152a33–35].) Finally, Nagel proposes causal luck as acknowledging that antecedent circumstances may determine who we are; some commentators have proposed that this category is redundant given constitutive and circumstantial luck (Nelkin 2019). For present purposes, our examples will focus on resultant luck.

14. See, e.g., Owen (1995, 1). See also Aristotle (2009, 1130a13–1131a9).
15. Waldron, originally from New Zealand, favorably comments on a similar scheme there.
16. This could be complicated by any number of considerations, but it suffices for present purposes. One complicating set of features would just be that ordering any test carries with it some risks; not just from, e.g., radiation, but even for more quotidian reasons, like that the patient could suffer an accident (e.g., vehicular, falling) in transit for the test. And so it *could* be the case that adequate testing *increases* the risks of (some) adverse events. To think this through, we would need more empirical data: How many people die a year from, e.g., melanoma? How many of those deaths owe to failure to order a biopsy? How many deaths—or other injuries—are incidental to the biopsy itself (e.g., travel to/from the test site)? We could drill down into an ever more detailed analysis. But insofar as the investigation here is structural—as opposed to empirical—we can put these complications aside.
17. See, for example, Rachels (1975), Kagan (1989), Quinn (1989) and Kamm (2007).
18. See, for example, Hart and Honoré (1959, 59) and Wright (1985). But see *Osterland v. Hill* 160 N.E. 301 Mass. 1928, which held that the defendant was not civilly liable for failing to rescue a drowning swimmer, even though it would have been easy for the defendant to do so. Regardless, it is uncontroversial that omissions can be causes: The issue is simply which omissions comprise breach of a duty. See Husak (1980).

#### REFERENCES

- Allen, Marshall, and Olga Pierce. 2016. “Medical Errors Are No. 3 Cause of U.S. Deaths, Researchers Say.” May 3. *National Public Radio*. <https://www.npr.org/sections/health-shots/2016/05/03/476636183/death-certificates-undercount-toll-of-medical-errors>
- Allhoff, Fritz, and Sandra L. Borden, eds. Forthcoming. *Ethics and Error in Medicine*. New York: Routledge.
- American Hospital Association. 1999. *Hospital Statistics*. Chicago: Health Forum.
- Aristotle. 2009. *Nicomachean Ethics*. Translated by David Ross. Oxford: Oxford University Press.

- Centers for Disease Control and Prevention, National Center for Health Statistics. 1999. "Births and Deaths: Preliminary Data for 1998." *National Vital Statistics Reports* 47 (25).
- Centers for Disease Control and Prevention, National Center for Health Statistics. 2001. "Deaths: Final Data for 1999." *National Vital Statistics Reports* 49 (8).
- Garrett, Jeremy, and Leslie McNulty. Forthcoming. "A Just Culture Approach to Medical Error." In *Ethics and Error in Medicine*, edited by Fritz Allhoff and Sandra L. Borden. New York: Routledge.
- Golemon, Luke. 2019. "Medical Over-Testing and Racial Distrust." *Kennedy Institute of Ethics Journal* 29 (3): 273–303
- Hart, H.L.A., and Tony Honore, 1959. *Causation in the Law*. Oxford: Oxford University Press.
- Harvard Women's Health Watch. 2018. "Radiation Risk from Medical Imaging." April 2. *Harvard Health Publishing*. <https://www.health.harvard.edu/cancer/radiation-risk-from-medical-imaging>
- Husak, Douglas. 1980. "Omissions, Causation, and Liability." *Philosophical Quarterly* 30 (121): 318–326.
- Kagan, Shelly. 1989. *The Limits of Morality*. Oxford: Oxford University Press.
- Kamm, Frances. 2007. *Intricate Ethics: Rights, Responsibilities, and Permissible Harm*. Oxford: Oxford University Press.
- Kant, Immanuel. 1998. *Groundwork of the Metaphysics of Morals*. Translated by Mary Gregor. Cambridge: Cambridge University Press.
- Kohn, Linda, Janet Corrifan, and Molla Donaldson, eds. 2000. *To Err Is Human: Building a Safer Health System*. Committee on Quality of Health Care in America, Institute of Medicine. Washington: National Academy Press.
- Larsen, Rasmus. 2019. "Psychopathy Treatment and the Stigma of Yesterday's Research." *Kennedy Institute of Ethics Journal* 29 (3): 243–272
- Makary, Martin, and Michael Daniel. 2016. "Medical Error—The Third Leading Cause of Death in the U.S." *BMJ* 353: i2139.
- McGrath, Sarah. 2005. "Causation by Omission: A Dilemma." *Philosophical Studies* 123: 125–48.
- Nagel, Thomas. 1979. *Mortal Questions*. New York: Cambridge University Press.
- Nelkin, Dana. 2019. "Moral Luck." *Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/entries/moral-luck/>
- OED Online. March 2019. Oxford University Press. <https://public.oed.com/blog/new-words-in-the-oed-march-2019/>
- Owen, David. 1995. *Philosophical Foundations of Tort Law*. Oxford: Clarendon Press.

- Owen, Steven, and Jason Johnson. 2017. "Radial Nerve Damage after Venipuncture." *Journal of Hand Microsurgery* 9 (1): 43–44.
- Peña-Guzmán, David, and Joel Michael Reynolds. 2019. "The Harm of Ableism: Medical Error and Epistemic Injustice." *Kennedy Institute of Ethics Journal* 29 (3): 205–242.
- Quinn, Warren. 1989. "Actions, Intentions, and Consequences: The Doctrine of Doing and Allowing." *Philosophical Review* 98 (3): 355–82.
- Rachels, James. 1975. "Active and Passive Euthanasia." *New England Journal of Medicine* 292: 78–86.
- Statman, Daniel, ed. 1993. *Moral Luck*. Albany: State University of New York Press.
- Waldron, Jeremy. 1995. "Moments of Carelessness and Massive Loss." In *Philosophical Foundations of Tort Law*, edited by David Owen, 387–408. Oxford: Oxford University Press.
- Williams, Bernard. 1981. *Moral Luck*. Cambridge: Cambridge University Press.
- . 1993. "Postscript." In *Moral Luck*, edited by Daniel Statman, 251–58. Albany: State University of New York Press.
- Wright, Richard. 1985. "Causation in Tort Law." *California Law Review* 73 (6): 1735–1828.