

## Against the Odds: the Case for a Modal Understanding of Due Care\*

By Jeff Helmreich and Duncan Pritchard

How do you know if an injury was inflicted negligently? How do jurors, in particular, determine whether the speeding ambulance driver or the dynamiting construction worker took what is known as “due care” toward those they injured, so that harming them – while tragic and regrettable – fell short of negligence?

The standard answer in contemporary legal doctrine involves probability. Specifically, we determine whether the probability that some injurious activity would harm anyone was low enough, in light of both the severity of the harm risked and the burdens of preventing it.<sup>1</sup> As the immortal Hand formula puts it: was  $BA > pL$  – did the burden of preventing harm outweigh the probability of harm, or loss, factoring in its severity?<sup>2</sup> If so, then the injurer exercised due care and is not liable.

For at least the past 100 years, this probabilistic standard has worked well enough for the familiar types of cases in which someone, like a small business or service professional, must choose a course of action that may harm others, but is otherwise too valuable to ban altogether.<sup>3</sup> What, for example, will be the express truck driver’s permitted speed, or the ambulance siren’s maximum volume, or the prescribed regimen of a medical treatment? In such cases, it is

---

\* Forthcoming in *Truth and Trials: Dilemmas at the Intersection of Epistemology and Philosophy of Law*, (eds.) Z. Hoskins & J. Robson, (Routledge).

<sup>1</sup> For the canonical expression of this standard, if not the first, see *United States v. Carroll Towing Co.*, 159 F 2d 169, 173 (2d Cir. 1947).

<sup>2</sup> *Id* at 173.

<sup>3</sup> Gregory Keating, “Must the Hand Formula not be Named?” *University of Pennsylvania Law Review Online* Vol. 163 (2015), 367-75, 369.

reasonable to assume that choosing conduct with a sufficiently low probability of harm constitutes a reasonable effort to prevent it (even if one could, in principle, try even harder.) Probability, in such cases, seems to track care.

That, however, does not apply to an increasingly common type of injury generating lawsuits today. With a growing number of injurious activities, the populations and products involved are so numerous – and the rate of injury so predictable – that even the lowest probability of injury is sure to materialize, and the injurer can reliably foresee it.<sup>4</sup> A large chain construction company whose crews' helmets fail once in every hundred falls, or something along those lines, will surely lead workers to harm over many years and hundreds of jobs. Often this result is expected by the corporation, which proceeds anyway. In such cases, the low probability of injury does not in itself show that the agent took care of some kind not to inflict it; the same agent might well have intended the harm or been entirely indifferent to it.<sup>5</sup> The probabilities would be the same.

That, of course, is not to deny that it is virtually impossible for such large-scale actors – “aggregative agents,” we might call them – to stay in business without injuring and even killing people on rare occasions. Such is the familiar, and perhaps worthwhile, price of having ambulances, airlines, hospitals and mass-market medicine, for example, among many other goods we take for granted. Nor is it to condemn these “aggregative injurers” as negligent. It is,

---

<sup>4</sup> The first and only in-depth analysis of this phenomenon, to our knowledge, is Kenneth W. Simons, “Statistical Knowledge Deconstructed,” *Boston University Law Review* 1 (2012): 1-87, which shows that, ethically speaking, such injurers are not easy to distinguish from individuals who knowingly harm particular people (even if the law's differential treatment of them is ultimately defensible, as Simons argues (esp. pp. 62-66)).

<sup>5</sup> It *can* show, however, that the corporation took care not to inflict *more* injuries, and indeed, that its cost-benefit calculation could have been sufficiently careful to make liability unjustified. See Simons, *supra*, pp. 62-66. Our question, in contrast, is whether it took sufficient care not to cause injury at all, as is arguably required of individuals (see *infra*, I(a)), on pain of moral negligence.

rather, to raise the question of how we might epistemically capture the care they take not to inflict harm, if such there be, given the inadequacy of the traditional probabilistic test in this context. What, in other words, demonstrates whether aggregative agents exercise due care to avoid injuring the people they are statistically certain to injure?

We proceed by first identifying more precisely why a probabilistic standard fails to capture due care in these large-scale contexts. We then explore ways that classical standard might be preserved anyway, even when applied to aggregative agents, and explain why we do not find them successful. Instead, we argue, the best understanding of due care, at least as it extends to such large-scale cases, is non-probabilistic. In particular, we argue the aggregative agent's riskiness to others, which must be low on pain of negligence, should be understood *modally*, not probabilistically. With a modal understanding of risk, we argue, legal standards can capture what it is to try to reduce one's harmfulness in the face of a virtual guarantee that the risk will materialize. A modal understanding has the further advantage of matching our pre-aggregative understandings of due care, as well, or so we will argue.

## I. Why Not the Probabilistic Standard?

Due care is defined as the care that a reasonably prudent person (or party) would take to avoid injuring someone.<sup>6</sup> Whatever else such care amounts to, we assume that, at a minimum, it is incompatible with either intending to inflict the harm in question or absolute indifference to whether one inflicts it, as manifest in having no disposition whatever to prevent it, even if one

---

<sup>6</sup> NEW YORK PATTERN JURY INSTRUCTIONS-CML 2:10 (3d ed. 1999).

could do so without cost. Both acts – intentional and indifferent infliction of harm – clash with taking due care to prevent it, or so we assume. Yet aggregative agents can be guilty of just such malice or indifference and still meet the probabilistic standard for due care.

Drawing on the infamous Ford Pinto case,<sup>7</sup> we might consider a bus company, Breakout, that uses brakes that function well except in highly unusual weather conditions – when the temperature, air pressure and humidity converge at just such a combination – at which point they fail suddenly and completely, injuring and even killing drivers and passengers. But weather statistics show that this danger can affect at most one in 10,000 rides in a given year. The cost of a safer system, however, is prohibitive – it would slash the company’s profits nearly in half, leading either to bankruptcy or a transportation shortage in its main client city. We can assume that the probability of harm, in light of this cost, is just low enough to satisfy the classical probabilistic formula for due care, but high enough that in any two years, it is reasonable to expect at least one bus have an injurious accident. Is the bus company exercising due care?

Here is one possibility that, if realized, could suggest otherwise: suppose it turns out that Breakout has just acquired another company, a car service, for which a safer system is financially manageable. It’s safer than the buses. Careful market research reveals, moreover, that one source of consumer interest in the car service is precisely that it is not prone to occasional brake failure or similar defects, like those of the buses. Breakout thereby deduces that maintaining its current bus service, with its occasional accidents, is actually profitable, more so than without the occasional accidents. So Breakout deliberately continues to market, promote and produce its bus services, intending – indeed hoping for – the injurious accidents once in a

---

<sup>7</sup> Grimshaw v. Ford Motor Company, 174 Cal. Rptr., 348 (Cal. Ct. App. 1981). The decision for Ford, unlike here, involved the risk of a fuel tank catching fire, and the probability of harm, though perhaps reasonably avoidable by a way of a precautionary modification, was actually much lower than in the case at hand, as we’re constructing it.

while. It's all part of the plan, as is quietly expressed whenever the question of whether to continue the bus service comes up in board meetings.

In short, Breakout is intentionally inflicting the harm its buses are statistically certain to cause.

We can also consider a different company, Flakeout, a bus company with the same size and brakes hazard as Flakeout, but which has not acquired another company. What Flakeout does have, instead, is a new study by its team of researchers promising a means by which the (occasionally) faulty brakes could be replaced with safer ones at no cost (we can assume nothing, yet, about whether the alternative will actually work). Nevertheless, Flakeout's board decides not even to consider the recommendations; they do not even give the researchers a hearing. They're comfortable in their current business plan, with the injuries at their low, if expected, level, and would rather not bother with exploring hypotheticals. The CEOs have other things to do with their time.

In short, Flakeout is indifferent to whether it is harming people.

We propose that Flakeout and especially Breakout are plainly not exercising due care to avoid the dangerous accidents and injuries their businesses will inevitably inflict. They are, in other words, negligent or worse. And yet, such agents by hypothesis meet probabilistic standards of due care. For that reason, we propose, such standards are unreliable. They simply fail to capture whether injurers at this level of aggregation were taking sufficient care – indeed any care at all – not to cause harm to others.

This conclusion may seem less than dramatic in light of the fact that most aggregative agents are not Breakout or Flakeout; they are neither willful nor absolutely indifferent about the

harm they stand to inflict. There is, however, a large range of conduct between reasonable care, on the one hand, and outright willfulness or indifference, on the other. Our worry is that a standard that cannot differentiate between these extremes – at least for aggregative agents – surely fails at the more challenging, if more common, borderline cases that raise genuine questions for negligence law.<sup>8</sup> In these sorts of cases, we ask not whether the injurers acted with malice, but rather with too much acceptance, or too little resistance, vis a vis the harms they stood to inflict on others. Did they do enough to prevent it? That is the question that a typical negligence juror must answer. For the reasons we just laid out, a standard based entirely on the probability of harm will not answer it.

#### A. Cost-Benefit Tradeoff

One way of defending a probabilistic standard, even in these large-scale contexts, is to revise slightly the notion of due care at work here, specifically the canonical (and, we submit, commonsensical) association of it with effort, of some kind, not to cause harm. Instead, it might be proposed, due care simply involves striking the right balance between the disvalue of inflicting the harm and the burden shouldered to prevent it, as reflected in equations like the Hand formula.<sup>9</sup> Due care is, in short, a cost-benefit tradeoff. One decides just how much is worth sacrificing to avoid just how much harm to others, and then proceeds accordingly, even if it involves knowingly injuring a certain amount of people. The right tradeoff, in light of the burden and benefits involved, constitutes due care, or non-negligence, on this reading.

---

<sup>8</sup> As Benjamin Zipursky points out in his “Reasonableness in and out of Negligence Law,” *University of Pennsylvania Law Review* Vol. 163 (2015), p. 2154.

<sup>9</sup> For a version of this interpretation, see Barbara Fried, “Facing up to Risk,” *Journal of Legal Analysis* 10 (2018), pp. 1-24.

In support of this understanding of due care, it might be noted that parties in any context – aggregative or individual, repeat-actors or one-off risk imposers – are not required to devote all their energies to ensure the safety of others. Rather, the amount of effort they expend, on pain of negligence, varies with the harm they might inflict and the benefits of the action that may inflict it.<sup>10</sup> A professional driver can clearly pose less danger if she sticks to 10 miles per hour, but that is obviously not required, even as permitting a higher speed virtually guarantees more injuries in the long run.

This tolerance of harmful conduct – built into our ethics of care – appears to suggest that due care consists less in trying not to injure anyone, at all, than in making sure the injuries one stands to inflict reflect the right tradeoff, the right balance of harms against benefits overall. In this way, the aggregate cases merely make manifest a proportional balance that constitutes all exercises of due care. The probabilistic standard, then, is perfectly appropriate for capturing the amount of harm a reasonable party stands ready to inflict, given the costs and benefits involved. And that is just what due care amounts to, on the proposed (revised) understanding.

True, such a standard can be met by agents like Breakout, who mean to inflict the harm that such a tradeoff permits. But to object on those grounds would arguably be an overreach. After all, any permissible infliction of harm could, in principle, be performed by someone who privately embraces or intends it. The joke about sadists tending towards dentistry is not, after all, an indictment of the dental profession; some jobs just permit or at times require hurting people, so why not use it as an outlet for a perverse desire to inflict pain, as long as it is done within ethically permissible limits anyway? Breakout may privately hope for the harms it inflicts, for

---

<sup>10</sup> This point is arguably essential even to our non-legal, purely moral, notion of the riskiness of certain conduct. See Judith Jarvis Thomson, *The Realm of Rights* (1990), pp. 243-46.

whatever perverse or profiteering reason, but we need not scrutinize intentions or aims once the right conduct has been followed, or so the cost-benefit care theorists could reasonably insist.

Of course, this proposal asks that we ignore the plain meaning of the notion of due care, specifically thoutmatic notions like “care,” which seem to connote trying. But that would be an acceptable adjustment if cost-benefit balancing captured what is most importantly involved in due care.

There are, however, reasons for doubt. Note that on this understanding of due care, even the individual agent – the driver, the surgeon, or the service-person settling on a “standardly” safe procedure – is tasked with a mere cost-benefit tradeoff. That is all the care she owes others, as far as ensuring their safety on pain of negligence. Although this end may move her to expend effort aimed at not harming others, such effort is not strictly necessary for the purpose. In principle, the same cost-benefit balance might be struck by an alternative, less effortful, means.

Consider, for example, a surgeon who finds she can reduce her occasional slip-ups – and the harmful nicks and bad calls they engender – by taking a new kind of energy pill that eliminates fatigue and inattention altogether. She becomes, instead, unwaveringly focused and alert. The only downside is that, as a side effect of the pills, she will blank out, essentially falling asleep, for around two seconds per week. Yet she calculates that any harm resulting from such a blackout – which could uncontrollably overtake her at the operating table – is statistically offset by the overall reduction in harmful slipups during surgery. Taking the pill strikes the right balance overall between harms and benefits, exactly as the Hand formula would dictate for the activity of surgery, let’s say. Still, the surgeon would uncontroversially be found negligent if she took the pills, and consequently blacked out during surgery, injuring a patient. It wouldn’t matter that the right balance was struck overall in the decision to take the pills in the first place. She



could not claim to her victim that she was exercising due care to avoid harming her, anymore than a pilot who crashed a plane during a cat nap he took to increase his overall alertness.

In other words, individuals are not permitted to simply refrain from trying not to harm others even if it's only for a proportion of their time and it leads to greater gains in overall safety. Still, even when their actions strike the proper cost-benefit balance, they would not count as exercising due care if they decided to forgo all effort for a subset of people, places, or time periods, even if no particular class of people were singled out for this inattention.<sup>11</sup> As John Gardner has argued, the duty of care, at least at the individual level, is essentially a duty to try.<sup>12</sup>

Legal practice has reflected this understanding. A ship captain who proceeded on course despite learning that a crew member was missing, and believed overboard, was found to have breached a duty to the missing seaman (when his family sued) despite establishing that there was no reasonable means to save him.<sup>13</sup> Nevertheless, the court held, he and his crew should have made some effort.<sup>14</sup> They should have stopped course and searched, no matter how futile they could reasonably expect it to be. Failure to do so breached the duty of care.

In other words, due care, as a general standard, cannot be reduced to a simple cost-benefit tradeoff. At the individual level, at least, some effort must be directed against harming people altogether, even as the amount or extent of effort reflects a balancing of the costs and benefits involved.

---

<sup>11</sup> For an explanation of why such conduct is worse when it imposes a concentrated risk on a small number of people, see Simons, "Statistical Knowledge Deconstructed," *supra* note 4, esp. pp. 36-43, 84-85.

<sup>12</sup> John Gardner, "Obligations and Outcomes in the Law of Torts," in P. Cane and J. Gardner, eds., *Relating to Responsibility: Essays for Tony Honore* (Hart Publishing, 2001), p. 10.

<sup>13</sup> *Gardner v. Nat'l Bulk Carriers, Inc.*, 310 F.2d 284, 287 (4th Cir. 1962). Thanks to Ken Simons for bringing this case to our attention.

<sup>14</sup> *Id.*

## B. A Hybrid Standard

Still, it may be proposed that things are different at the aggregate level. For corporations, public agencies and others affecting many thousands of people over a long period of time, where injury is statistically certain no matter how low the odds, it might be suggested that trying not to inflict harm is simply untenable. To test this hypothesis, we would need to attend more closely to what is involved in the effort that seems necessary for due care, at least at the individual level, and see if it can be replicated by aggregative agents.

To make the comparison as close as possible, we should imagine an individual agent who – like his aggregative counterparts – is statistically certain to cause injury in the long run, and he knows it.

For example, consider a truck driver who follows the prescribed standard speed, based on the weight of his load and the power of the brakes, among other factors, thus keeping the probability of harm negligibly low, say two per cent. In other words, he drives carefully, by any reasonable estimate, at all times. Yet he plans to do it for life; he's in it for the "long haul," as he likes to put it. And he knows that over 10 years of daily driving, even if his carefulness and competence never falters, his seemingly low 2% chance of injury is virtually guaranteed to materialize. He will, inevitably, injure someone sooner or later. In what sense, then, can we say that he is, nevertheless, exercising due care?

Without attempting a full account of trying, we can observe that with such agents, if they exercise reasonable care, they are characterized by the following disposition, which could be called the *Effortful Stance*:

*A risk-imposing agent is in the effortful stance if and only if she is disposed to ensure that she either reasonably expects not to harm anyone, or is taking steps to remove whatever factors make that expectation unreasonable.*

Thus when driving under ordinary conditions, the 2% probability of harm makes the trucker reasonably confident that, at this very moment, he will injure nobody. But when a pedestrian suddenly winds up in his path, failing to budge as he heads straight for her, he acts to restore the earlier equilibrium: slamming on the brakes, swerving, honking and otherwise manipulating the situation to bring down the odds to their earlier level. Despite the statistical certainty of harming someone over the long haul, then, we can nevertheless grant that Long Hauler is exercising due care, in part because he is always in the effortful stance, reasonably expecting to be safe or else poised to remove any factors that suddenly challenge that expectation. I say “in part” because taking steps to remove dangers is not necessarily sufficient for “due care”; at least since *Vaughn v. Menlove*,<sup>15</sup> we legally require that people take the *right* kinds of steps – actions that really *do* make it unreasonable to expect to harm someone.

The problem for aggregative agents, it seems, is that they cannot take the effortful stance, or so a hybridist about due care could reasonably claim. That is because, first, they cannot reasonably expect to avoid injuring anyone. The owner of Long Hauler’s trucking service company, unlike Long Hauler himself, knows that as soon as she takes the action of prescribing a

---

<sup>15</sup> 132 Eng. Rep. 490 (C.P. 1837). A landowner was found negligent for building a highly combustible hayrick near his neighbor’s property – which consequently burned when a spark from a nearby railroad set the rick on fire – despite his having taken great pains to prevent this occurrence, at least from his own (flawed) point of view. Taking strenuous care, in other words, is insufficient; non-negligence requires *reasonable* care.

certain speed for the truck or ambulance drivers and unleashes her fleet on the world for the coming year, she has acted in a way that is statistically certain to inflict some injury. In proceeding with the business plan anyway, the corporate decisionmakers – though hopefully not intending the harm they stand to inflict – do reasonably expect it. And even if they did not, they cannot take steps to remove factors that seem to raise the probability of causing harm in some particular instance. They are, after all, aggregative agents – acting in a context necessarily removed from the individual cases that may implicate or depart from the statistical “normal” they try to determine. They will not be in the driver’s seat, so to speak, poised to stop or swerve as needed.

Or so these theorists would argue. As a result, on their view, aggregative agents — the large corporation or the massive public agency or the national government – cannot adopt the effortful stance that due care may require at the individual level. They must, instead, settle for cost-benefit calculations and hope they strike the right balance. Due care, then, should be understood differently at the aggregate level.

Notice, though, that this proposal depends on the claim that the aggregative agent – the truck company, for example – is acting in a way it reasonably expects to cause harm. That is what we aim to dispute; indeed, we aim to show that, to the contrary, aggregative agents can adopt the effortful stance, after all. The next section is devoted to showing just how that could be.

## II. Modal Risk

So far we have argued that the probabilistic standard for due care fails with regard to aggregative agents, inasmuch as it fails to capture whether they are acting to prevent the harms

they stand to inflict, rather than fully intending or at least accepting them. This way of putting things, however, suggests that it is in principle possible for aggregative agents to try to avoid harming others, even when they know the probability of such harm is statistically certain to materialize. It suggests they can adopt the effortful stance. But that stance requires that agents act to ensure that they can reasonably expect their repeated activities *not* to harm others, or else take reasonable steps to remove whatever factors make that expectation unreasonable.

How, then, can aggregative agents – corporations, agencies, countries and the like – adopt the effortful stance, and what standard of care will capture whether they have done so?

We believe both questions can be answered by substituting a modal understanding of risk for a probabilistic one. On this proposal, an agent – including an aggregative agent – is exercising due care if and only if it satisfies the legal standard of care, or non-negligence, (re)interpreted modally, rather than probabilistically. What is a modal understanding of the standard of care? Recall that a duty of care or non-negligence, in both law and morality, is premised on the view that one should, on pain of negligence, act with a sufficiently low risk of harming others, varying with the extent of the harm and risk involved, and the burden of preventing it. Despite all we have said so far, we agree with this basic tenet. Where we disagree with the prevalent practice is in understanding that risk probabilistically, so that a low risk of harming others amounts to a low probability of doing so. That understanding, we have tried to show, generates problems distinguishing the obviously negligent from the non-negligent in aggregate cases. On the other hand, those same problems do not arise if we substitute a modal understanding of risk.

On a modal understanding, as developed by Pritchard in several recent papers,<sup>16</sup> the riskiness of some event is a function of its modal distance, or the degree to which current conditions – the actual world we live in (as contrasted with possible, different worlds) – must change for it to take place. So an agent who seeks to lower the risk of his  $\phi$ -ing, on the modal understanding, ensures that the conditions in place will block his  $\phi$ -ing, unless basic features of our world change unpredictably. On this view, when we take due care not to harm someone, and consequently make sure the risk is low, the relevant current conditions – including the ones we’re putting or keeping in place – are incompatible or at least in tension with causing such harm.

Return to the case of Breakout, the truck company with the brakes that fail under rare but recurring weather conditions. Faced with a potential scandal, the company may consider two ways of warding off the danger it now poses to its own passengers and others, Option A and Option B. On Option A, at a slight cost, Breakout could replace the brake system with one that fails in only 50% of the weather scenarios that currently bring on the danger. On Option B, at the same cost, the company can build in a warning system that detects the weather conditions immediately and slowly and safely begins to shut the bus down as they occur, with enough time for the driver to be alerted and guide the bus to safety. Both options still pose dangers in the long run: the first option for the obvious reason that it still preserves the same harmful conditions, albeit in a much lower proportion of weather scenarios. The second because it involves a man-made system that, like any other, can reasonably be expected to fail over time at least as much as any other – perhaps in the same proportion of cases that remain hazardous on Option A.

---

<sup>16</sup> See Duncan Pritchard, “Risk,” *Metaphilosophy* 46 (3):436-461 (2015); and Duncan Pritchard, “Legal Risk, Legal Evidence and the Arithmetic of Criminal Justice,” *Jurisprudence* 9:1 (2017), 108-19.

On a modal understanding of risk, only Option B actually reduces the riskiness of the harm. The conditions that would be put place clash with injury occurring at all; that is, unless something goes wrong or very differently with the world as it is. When harm does, nevertheless, result, it is reasonably characterized as a puzzle, a mystery or at least a surprise, appropriately met with the query, “What went wrong?” Or, it may be discovered, the conditions the company reasonably believed were in place – which should have prevented an accident – did not perfectly obtain in every instance, much as they tried. The shut-down system failed. In contrast, the company that merely lowers the probability of its weather-induced brake failure, so that statistically it occurs once as opposed to twice every five years, let’s say, is acting in a way that is perfectly compatible – indeed hospitable – to the occurrence of harm, however infrequent. Its occurrence comes as no surprise; it clashes with nothing in place. The probability is low, in other words, but the modal risk is high.

Similarly, a bridge building company could decrease the injuries due to falling workers by lowering the height at which they work, which has been shown to lower the overall extent of injuries, or placing nets that intercept workers as they fall. Only the latter option reduces the modal risk, we maintain, inasmuch as it is designed and intended to prevent harm altogether. True, as with any measure, things don’t always work out: occasionally, a net fails, snaps, or is improperly fastened, let’s say. Or some other, as-yet-unimagined mishap occurs. Still, if the net system is in place, and has no currently accessible flaw – so that it is reasonably expected to work – then modal risk is kept low.

Notice that on a modal view of risk, we can see how even an aggregative agent, aware that its chosen course of action will injure people, can nevertheless meet the conditions of the effortful stance and so of due care. To take that stance is, again, to be disposed to ensure that one

either reasonably expects not to injure anyone or is working to remove whatever factors make that expectation unreasonable. One seeks, in other words, to make it a more surprising, less expectable occurrence. As we have seen, this cannot be done by merely lowering its probability. An event that occurs once in 10,000 cases will still definitely occur in the aggregate; it is no surprise. But it can be done by putting in place conditions that are incompatible with the occurrence of harm in any particular scenario. In so doing, the aggregative agent makes injuring others a more distant modal possibility, one that should not be expected.

While these aggregative decisionmakers are aware their efforts will fail on occasion, like those of Long Hauler the trucker, they are at the same time actively taking steps to remove any reasonably expected danger. Although harm is bound to occur in the long run, they are taking steps that, on a first-order level, make any particular injury unreasonable to expect. Their preventive measures are, on their face, reasonably taken to be incompatible with the harm's occurrence; the better the measures, the more incompatible.

Notice this proposal can at times resemble the probabilistic view of risk. The modal distance of an event may correlate inversely with the odds of its occurrence. If an event would not occur in the ordinary course of things (in this world or nearby worlds, in modal terms), it is surely unlikely in the probabilistic sense, as well. Importantly, however, the reverse is not true. Some events with very low probability are not modally distant. A snack that will kill people with a rare gene that is impossible to detect – it is known only that one in a million have it – carries a low risk in terms of probability, but not modally. The world in which the risk materializes is exactly our own – nothing in the everyday course of things goes wrong, or changes dramatically, when the food proves fatal. A firm producing and selling such a product can be reasonably



presumed to accept this harm, seldom as it is likely to occur – and so failing to exercise due care. That is because the firm is not in any sense trying to prevent the (low) risk from materializing.

That brings us to a second and more important feature of the modal account of risk-taking: it is dynamic. An event cannot be modally distant if its occurrence becomes a reliable part of our world, our everyday state of things, even if we are unaware of how, precisely, it happens. For example, some viruses develop resistance to vaccines, despite the good reasons we have to expect the drugs to work generally. It's rare, and it's the exception, but it happens regularly and reliably – even though we don't know the precise causal mechanism by which it does. In these cases, then, the drug-resistance is not modally distant from the world in which the vaccines are introduced and, in general, effective.

The same, then, can be said of safety precautions or driving practices that reliably result in a certain amount of harm, over time. Once that figure sets in – like a law of nature – the harm is no longer modally distant, even if we cannot account for it given all we did to prevent it. If that failure reliably occurs, with a predictable pattern, then it is arguably a built-in feature of our precautionary steps, and so those steps do not increase its modal distance.

In other words, a modally distant event that nevertheless occurs must be, in some sense, a surprise given the conditions that preceded it – it must be in tension with them –including, in this case, our efforts to prevent it. We must be stumped or at least challenged by its particular occurrence, even if in some detached mode we can expect such “surprises” or “challenges” to occur. “Expect the unexpected,” we might say from a more detached, global perspective, but the particular occurrence in question remains a puzzle, a kink, something to be investigated and eventually eliminated.

Our safety efforts, then, are – on pain of negligence or recklessness – constantly in flux: as soon as we learn of a flaw in our current system – defined simply as a regular and predictable proportion of failures – we have to revise the system so that, given the new one, the event would be highly unexpected under any conditions, even if it will probably fail again (and we know it).

### III. Potential Worries

It might be worried that this proposal is little more than an elaborate directive to try to keep the probability of harm as low as possible. The difference, then, bears more explicit clarification. Consider car-racing in the mid-20th century. Both the stadiums and the car companies may be doing their utmost to keep the probability of fatal or injurious crashes as low as possible. Suppose that the best efforts available place the odds at 10% or less. In such cases, they fail to take due care, because even their best efforts are highly compatible with injury occurring. That, moreover, is not simply due to the high odds of one in 10; indeed, it would be true if the odds were one in 1,000. It is, rather, that the safety efforts of the various organizers and trainers (and drivers, perhaps) are plainly compatible with injury, nevertheless, occurring, if less often than without those efforts. There is nothing about these efforts that, even *prima facie*, rules out injury.

Similarly, a doctor who tries her hardest to save a terminally ill patient from near certain death is not taking due care as understood on the instant proposal, even though her efforts couldn't be more strenuous (and should, of course, be made). The reason is that she is not taking steps that could reasonably be seen as making this world incompatible with the patient's death.

She is merely working to make a likely possibility a little less likely. Her failure in this respect would not be considered a malfunction or misuse of efforts at due care.

This talk of “reasonably seen” and “prima facie” might, however, invite a different sort of worry. The “modal risk” avoider may be seen to face a dilemma: either modal distance is purely subjective, a function of what stumps or mystifies us, or it is objective, a matter of the actual metaphysical compatibility of a harm with the conditions of this world. If it’s subjective, goes this worry, then due care could be (absurdly) a matter of looking the other way, making sure we don’t see the possibilities that our precautionary efforts may fail. As long as we don’t see the link between our behavior and the harm we inflicted, then it was a surprise and, on the proposed “subjective” understanding, modally distant.

But if, in contrast to this absurd proposal, we take modal distance to be an objective matter, a metaphysical property of correctly ordered possibilities, then any time injury takes place we can infer that it was, in fact, modally close. For although we don’t know the underlying causal explanation of the harm, there presumably was some such causal story behind the scenes. And that casual story, barring magic or supernatural intervention, follows more or less from the conditions that obtained at the time. Making an injury modally distant, on this worry, is either too easy or impossible!

In reply, it should be noticed, first, that the same dilemma applies to the probabilistic notion of risk. If the probability of a risked harm is measured given what we know, then the less we know, the less risky it is. If, on the other hand, it is a matter of how likely an injury’s occurrence is given the conditions that actually obtain, then talk of risk or probability is absurd, inasmuch anything that did happen was, *ex ante*, bound to happen (excluding sophisticated scientific examples of indeterminacy).

In both cases, the dilemma rests on neglecting that due care is both subjective or, more precisely, agent-relative and agent-neutral. It is agent-relative in that it requires an agent to act only on what she is capable of considering at the time. The actual underlying causal mechanism between possible events need not be her responsibility to learn. On the other hand, due care is agent-neutral in that what she should aim for is not a psychological condition but a metaphysical one: doing no harm. Also agent-neutral is the standard by which this aim is assessed: is she taking reasonable measures? On this standard, deliberately ignoring or trying not to see ways in which one's behavior may harm another does not constitute due care. It is not a reasonable way to aim at the goal of causing no harm. Nor, for that matter, is putting in place conditions that only seem to us modally distant from our causing harm, without any effort to check them more closely.

Again, the present thesis is that only a modal understanding of risk-avoidance can capture the proper, objective aim of the non-negligent agent. In contrast, an agent taking steps merely to ensure that the probability of harm is low might be aiming to inflict some harm, albeit a tiny amount, or might be aiming only not to inflict culpable harm. Or, as in the Breakout-type cases, he may be trading the harm he expects to inflict against the benefits of doing so.

A different worry one might have about the present proposal is that, unlike cost-benefit analysis, it fails to capture the way due care is always a matter of degree, varying both with the value of the risky activity in question (racecar driving vs. ambulance driving) and the nature and severity of the harm risked (minor injury vs. multiple deaths). On this worry, the modal view of risk looks dichotomous: a harm is either compatible or incompatible with the world as it is, which fails to make sense of the weighing and prioritizing involved in due care.

It is worth recalling, then, that modal distance, like probability, is a matter of degree, after all. With some events, little has to change about our world for them to occur; for others, a lot more has to be different. Your winning the lottery next time, instead of whoever does win it, is rather like the world as it is and will be. In contrast, your scoring the winning homerun in the next World Series requires much more to change. Many present-day conditions would be in tension with that outcome, or incompatible with it (no offense).

If the bridge construction company, mentioned above, elected to place well-tested netting below the elevated spots, to catch falling workers before they could get hurt, the firm would make harm, at least from falls, modally distant, even if the probability of a faulty netting (perhaps they patronized a scam netting manufacturer) equaled that of harmful falling under previous conditions. If they took further steps aimed at blocking other forms of injury – introducing safety helmets, for example – they would even further increase the modal distance of injury. And if they placed rubber or cushion on the ground, they would make injury even more modally distant. Modal distance, in short, varies continuously just as probability does. On a modal understanding of risk, then, due care amounts to ensuring that the risk is *sufficiently distant*, modally speaking, in light of the benefits and burdens involved.

#### IV. Conclusion: the Account in Perspective

Having introduced and defended a modal account of due care, we should take a moment to define its limits. While we believe such an account best captures what aggregative agents, like large agencies or corporations, can do to count as non-negligent with respect to the harms they might inflict, we do not claim that this is how every person taking due care needs to understand her activity. Again, the individual trucker who acts in a way that keeps the probability of harm low still counts as exercising due care. That is, of course, not to claim that aggregative agents are held to a different standard of care, as was suggested in the hybrid proposals discussed in Part II (B). Rather, it simply follows from what we have said about the effortful stance. As long as one reasonably expects not to inflict harm, the duty to try, in any active sense, simply does not get triggered; the effortful stance is in place but largely passive. And for an individual at a given moment, a low probability is, indeed, a reason to be confident one is not about to injure someone. It is only once injury becomes foreseeable that we must act to block it, which is, in effect, to make injury modally far off, removed from this world. For aggregative agents, however, the first condition is almost never satisfied; they do foresee the injury they stand to inflict and so must, already, act to prevent it. What we have provided is a way they can be seen to be taking this sort of effort, exercising this standard of care, even when – given the quantities and replications involved – injury is inevitable.

We have also not argued that reducing modal risk, as we have defined taking due care at the aggregate level, is the morally best or required thing to do. There may be cases in which the inevitable fallibility of any modal risk-reducing efforts – the failures sure to occur in large enough productions and populations – will foreseeably exceed the damage of simply lowering the statistical odds. That is a matter for both research and further moral theorizing, inasmuch as it

turns partly on a larger question: is it always best to try not to harm, altogether, rather than sometimes making difficult harm-benefit tradeoffs, instead?

We have not argued any answer to this question. But it bears recalling that we are writing in part against the backdrop of an opposing concern, which has gained prominence recently. On this picture, it is a mark of sentimental naïvete to insist that the care we require of individuals, trying never to harm anyone, extend to the aggregate level. On this picture, the layperson in the street may be outraged by a massive corporation's engaging in a policy or product it fully expects to harm people, but that is only because such folk fail to see things from the sophisticated, aggregative point of view, from which such callousness is unavoidable.

Against this picture, our accounts lends some support to the unsophisticated outrage. Indeed, we share the intuition that there is something objectionable about simply accepting that one's activities will harm a certain number of people, and being indisposed to change this as long as the numbers are low enough. This objection could be muted if it turned out that aggregative agents don't have any other choice, and so due care – at their level – has no meaning. As we have tried to show, however, they do, and it does.