Law and Philosophy (2026) 45: 73–93
https://doi.org/10.1007/s10982-024-09513-4


ADAM SLAVNY

# VARIETIES OF CONSEQUENTIALISM AND DEONTOLOGY IN THEORIES OF TORT LAW

(Accepted 1 August 2024)

ABSTRACT. An important contribution of Gregory Keating's *Reasonableness and Risk* is to mount objections to law and economics whilst articulating an alternative deontological vision of tort law to that offered by corrective justice theorists. In this paper I offer two reservations about Keating's account. One is that some forms of consequentialism can accommodate at least some aspects of his deontological theory, so we must be careful to distinguish between claims that strike at the heart of consequentialism and those that are only a problem for wealth maximisation. The second is that there are yet other varieties of deontology besides Kantianism and Keating's alternative to it, and while I accept some of Keating's claims, I think we should reject others. I will make the case for each of these reservations via discussion of three core features of Keating's account: the harm/benefit asymmetry, the commitment to objective rather than subjective valuation, and his justification of strict liability.

## I. INTRODUCTION

One of the principal debates within contemporary moral philosophy is that between consequentialists and deontologists. Consequentialism is the view that when determining the permissibility of some act (or rule, law, institution *etc.*) all that matters is the value of the state of affairs it brings about. Deontology has different manifestations, but all have a commitment to non-consequentialism,[1] or the view that the value of the state of affairs brought about by some act (or

[1] I will use the term 'deontology' as that is the term Keating uses to describe his own theory, although I generally prefer 'non-consequentialism' as this is more neutral between the various views that reject consequentialism and can be thought of as broadly deontological.

rule, law, institution *etc.*) is not the only thing that matters in determining its permissibility.[2]

Private law theory has been occupied with its own version of this debate. On the consequentialist side are those within the law and economics movement who view tort as a mechanism for maximising wealth. This is a form of consequentialism which has affinity with its cousin theory, utilitarianism.[3] On the deontological side are those, especially Kantians, who understand private law as implementing corrective justice, which is concerned exclusively with the correlative rights and duties that connect plaintiff to defendant rather than any effect the law has on overall wealth.[4]

Consequentialism vs. deontology is not the same debate as wealth maximisation vs. Kantian corrective justice. There are varieties of consequentialism beyond wealth maximisation and varieties of deontology beyond Kantianism. An important contribution of Gregory Keating's *Reasonableness and Risk* is to mount robust objections to law and economics whilst articulating an alternative deontological vision of tort law to that offered by the corrective justice theorists. I concur with Keating on many points, including his acceptance of the normative heterogeneity of tort law, his embrace of strict liability, his appeal to fairness and distributive concerns in relation to liability, and his belief that compensation schemes are not opposed to the normative structure of tort law but are in some ways an extension of it.

This broad agreement comes with two caveats that I explore in this essay. One is that there are forms of consequentialism that can accommodate at least some aspects of Keating's deontological account. We must be careful, therefore, to distinguish between claims that strike at the heart of consequentialism and so are crucial to

[2] This is a fairly traditional way of drawing the divide. Some consequentialists claim that most deontological commitments can be 'consequentialized,' see for example Jamie Dreier, "In Defense of Consequentializing", in Mark Tummons (ed.), *Oxford Studies in Normative Ethics*, (Oxford: Oxford University Press, 2011, Vol. 1)) pp. 97–119. We will return to the relevance of more sophisticated versions of consequentialism later.

[3] Richard Posner argues that wealth maximisation should be distinguished from utilitarianism, and that the former offers a firmer basis for a normative theory of law. See "Utilitarianism, Economics, and Legal Theory", *The Journal of Legal Studies* 8(1) (1979): pp. 103–40. Though I think both theories are vulnerable to the most important objections, the two should not be conflated.

[4] The most prominent examples are Ernest Weinrib and Arthur Ripstein. See Ernest Weinrib, *The Idea of Private Law* (Cambridge, Mass: Harvard University Press, 1995), Arthur Ripstein, *Force and Freedom* (Cambridge, Mass: Harvard University Press, 2009) and Arthur Ripstein, *Private Wrongs* (Cambridge, Mass: Harvard University Press, 2016).

deontological theories of tort law, and those that are a problem only for wealth maximisation. The second is that there are yet other varieties of deontology besides Kantianism and Keating's alternative to it, and while I accept some of Keating's claims, I think we should reject others. I will make the case for each of these reservations via discussion of three core features of Keating's account: the harm/benefit asymmetry, the commitment to objective rather than subjective valuation, and his justification of strict liability.

## II. THE HARM/BENEFIT ASYMMETRY

Keating argues that according to cost-benefit analysis, a mainstay of the economic approach, "harm has no special significance, and its avoidance has no special priority. Harm is just one possible cost in a calculus of cost and benefit, and costs and benefits are minuses and plusses on the same scale."[5] This symmetry between harms and benefits "is at odds with our ordinary intuitions and our law. In both morality and law, our obligations to avoid harming others are stronger than our obligations to benefit them."[6] As a result, cost-benefit analysis fails to capture some of the most deeply ingrained aspects of our legal and moral practices.

By contrast, Keating defends an asymmetrical view of harms and benefits. This asymmetry is tightly connected to his conceptualisation of what harms and benefits are: "Harms impair essential conditions of human agency. Physical harms–death, disability, disease, and the like–rob us of normal and foundational powers of action. Physical harm comes close to being unconditionally bad."[7] The power of benefits to enhance lives, on the other hand, is "much more contingent than the badness of physical harm."[8] Whether some benefit enhances a person's life depends on that person's plans and projects. For example, mathematical brilliance is of great use to one dedicated to a career in the sciences, but little use to an aspiring artist. Harms are not contingent in the same way because they impair our ability to function as agents generally. Physical injury, a standard basis for tort claims, impairs a person's life regardless of

[5] Gregory Keating, *Reasonableness and Risk* (Oxford: Oxford University Press, 2022) p. 10. Hereafter Keating, *Reasonableness and Risk*.

[6] *Reasonableness and Risk*, 10.

[7] *Reasonableness and Risk*, 10.

[8] *Reasonableness and Risk*, 10.

what their plans and projects happen to be. Keating argues that his view explains why the law prioritises physical injury and recognises only exceptional duties not to cause economic loss and emotional distress.

Keating also thinks of harm as a condition of impairment rather than a state of being worse off compared to some antecedent or alternative state.[9] When a person suffers a broken leg, for example, they are not harmed because they are worse off than they were previously, or worse off than they would now be had the injury not occurred. Rather, they are harmed because having a broken leg is a condition of impairment. An important implication of this is that harms can be identified non-comparatively, by attending to the intrinsic features of the relevant condition rather than comparing it to some prior or alternative state.

Already there is a great deal here to discuss. There are at least five distinct features of the harm/benefit asymmetry, as Keating understands it, that are worth separating:

(1) *Doing and Allowing*: Our obligations not to harm others are more stringent than our obligations to benefit others.
(2) *Incommensurability of Harms and Benefits*: Harms and benefits cannot be measured along a common quantitative scale.
(3) *Harm as Agential Impairment*: Harm impairs the essential conditions for the exercise of human agency.
(4) *Contingency of Benefits*: The value of a benefit to a person is generally contingent on that person's particular plans and projects, while the same is not true, or less true, of the disvalue of harm.
(5) *Non-comparativism*: Harm consists in a person being in a condition of impairment, not any comparison between that person's state and some alterative state.

As Keating emphasises, *Doing and Allowing* is reflected in the law's distinction between acts and omissions, and it is often accepted by deontologists and rejected by consequentialists. As such, he rightly considers it central to a deontological theory of tort law. However, it does not entail claims (2)–(5). These claims are not fundamental to

[9] *Reasonableness and Risk*, 12. The non-comparative view of harm has been prominently defended by Seana Shiffrin. See Shiffrin, "Wrongful Life, Procreative Responsibility, and the Significance of Harm", 5 (1999) *Legal Theory* p. 117, and "Harm and its Moral Significance", 18 (2012) *Legal Theory*: p. 357, at pp. 367–8.

deontology in the same way that *Doing and Allowing* is, and some consequentialists might accept them.

Consequentialist positions have at least two components: a theory of value and a view about how this value ought to be promoted. The theory of value tells us what it is that makes states of affairs valuable or disvaluable. Utilitarianism, for example, takes *utility* as its theory of value, perhaps understood as pleasure or preference-satisfaction. One state of affairs is more valuable than another if and only if it contains more utility. *Maximisation* is the manner in which utilitarians think value ought to be promoted. It tells us to bring about that state of affairs which is best in terms of utility, typically the one that contains the greatest aggregate utility.

Utilitarianism is a relatively simple form of consequentialism because it recognises only one value: utility. More sophisticated versions could accommodate some of Keating's points insofar as they are claims about the value of states of affairs, especially *Harm as Agential Impairment* and *Contingency of Benefits*. A consequentialist could accept the asymmetry of value and treat harms as more disvaluable than benefits are valuable.[10] There are consequentialisms that adopt more complicated theories of value in similar ways. For example, negative utilitarianism asks us not to maximise happiness but to minimise avoidable suffering.[11] Negative utilitarians might give weighted or lexical priority to the avoidance of suffering over the promotion of happiness on the basis of *Harm as Agential Impairment* or *Contingency of Benefits*. Such utilitarians can sensibly say we ought always to maximise value (whether through act or omission) thus denying *Doing and Allowing*, whilst maintaining that the value of states of affairs is governed by the asymmetry of harms and benefits. This should encourage us to be clearer about the proper targets of a deontological theory of tort law. Some features of the harm/benefit asymmetry count against the law and economics approach but not consequentialism more broadly. This also serves to emphasise the relative importance of *Doing and Allowing*, as this is the one principle within the harm/benefit asymmetry that bites against all, or at least most, consequentialist theories.

[10] Thomas Hurka, "Asymmetries in Value", *Nous* 44(2) (2010): pp. 199–223.

[11] For discussion, see R. N. Smart, "Negative Utilitarianism", *Mind* 67 (1958): pp. 542–3.

Perhaps Keating will accept this. His primary target, after all, is wealth maximisation in particular rather than consequentialism in general. But there is another reason why it is important to show that *Doing and Allowing* is separable from claims (2)–(5), which is that we have independent reason to reject claims (2)–(5). The remainder of this section will elaborate on these reasons.

Begin by noting that harm-avoidance is a type of benefit. For example, if I push someone into a river, I harm them, and if I pull them out of a river, I benefit them. The outcome that I cause and prevent respectively is the same (drowning), but in the first case I make the person worse off, whereas in the second I merely fail to make them better off. There are different ways of caching out the difference between harming and failing to benefit, but any plausible way of doing this will have to recognise harm-prevention as a type of benefit.

This fact casts doubt on the *Incommensurability of Harms and Benefits* and the *Contingency of Benefits*. If it is possible to measure the harm of drowning on some metric, it must also be possible to measure the benefit of an averted drowning on the same metric, as the impairment of agency in both cases is the same. In cases of benefit by harm-prevention, either both the harm and benefit are measurable, or neither are–there is no asymmetry. These cases show that the assumption underpinning the *Incommensurability of Harms and Benefits*, that harms and benefits are qualitatively different, is not always true. For the same reason, if the disvalue of a harm is not contingent on a person's plans and projects, then neither is the value of preventing that harm. Again, the impairment of agency in both cases is the same.

A possible response is to deny that the harms are really the same. Perhaps the person who drowns as a result of being pushed in a river suffers a greater harm, or at least a qualitatively distinct harm, than the person who is not rescued. There are general reasons to be suspicious of this move. It is important to distinguish harm as an element of a wrong from other elements such as fault and causation conditions. Inferring from *Doing and Allowing* that harm that is caused is worse than harm that is not prevented tends to blur this distinction. Also, we must remember that *Doing and Allowing* applies to innocently caused harm, not just culpable harm. It is hard to believe that innocently caused harm is worse *as harm* than non-prevented harm. More pertinently for our purposes, this way of thinking about harm does not rescue Keating's account. This is because, even if it points to some asymmetry between harms and benefits, it is not the asymmetry that Keating is after. His asymmetry depends on the value of benefits being much more contingent than the disvalue of harms, and this is still threatened by the fact that harm-prevention is a type of benefit.

This argument creates interpretive as well as normative problems for Keating's view. For example, the common law generally imposes no liability for failing to rescue others from harm, or for the infliction of harm by third parties. In these cases, the benefit in question is harm-prevention, yet there is usually no obligation to provide it. Some features of the asymmetry such as *Doing and Allowing* are well placed to explain these rules.[12] Others, such as the *Incommensurability of Harms and Benefits* and the *Contingency of Benefits*, are in tension with them.

Of course, not all cases of failure to benefit involve failure to prevent harm. Still, we can make similar points about the *Contingency of Benefits* and perhaps the *Incommensurability of Harms and Benefits* with reference to the value of money, the usual remedy for successful tort plaintiffs. Money is an instrumental good, a means of obtaining things, having experiences, pursuing goals and so forth. The instrumentality of money makes it an appropriate remedy for harm precisely because it is agency-enhancing and can be used for myriad purposes. In Rawlsian terms, money is a primary good, desirable for all citizens regardless of their conception of the good. It is common to make exception here for those who spurn wealth such as hermits, ascetics and the like, but even this is misleading as it treats money as though it can buy only tangible things. The ascetic can always use money to spread the word of asceticism rather than buying personal possessions. We are owed an explanation, then, of why the benefit of money is contingent in a way that harms such as physical injuries are not, given its thoroughly instrumental value.

[12] Perhaps in combination with some other considerations. It is doubtful that *Doing and Allowing* can explain the harshness of the common law rule on its own. *Doing and Allowing* holds only that obligations not to harm are more stringent than obligations to benefit, not that the latter obligations do not exist. Nevertheless, the doctrine might help to explain the common law position in combination with, say, the worry that liability for omissions will deter beneficence.

*Harm as Agential Impairment* also faces problems. It cannot explain why severe reductions in wellbeing that do not impair agency are harmful. Physical pain *can* hobble agency, but its status as a harm does not depend on this. A person who suffers chronic pain because of an accident but pursues their projects with as much or even more vigour than before is harmed despite the lack of any impairment to their agency. A possible response is that such pain is a cost rather than a harm. Keating uses the term 'cost' to describe a broader range of negative outcomes that may affect a person's wellbeing without impairing agency.[13] *Recall, though, that Keating's conception of harm is meant to track outcomes that are worthy of compensation, so even if we accept the distinction, surely severe pain that doesn't impair agency must fall on the harm side of the divide.*

It is not just that conditions that are intuitively harmful cannot be so identified on Keating's view. It is also that the view prevents us employing the concept of harm in some important ways. For example, it seems to imply that forced benefits are really harms because they deny a person's power to govern their life in accordance with their own will.[14] This risks making two familiar concepts incoherent: paternalism and self-sacrifice. Take paternalism first. Paternalism can be understood as improving a person's situation for their own good without their consent. When paternalism is wrong, it is a distinctive wrong partly because it cannot be explained by appeal to harm. It is wrong even if it succeeds in benefiting the paternalized person. If we accept the implication that forced benefits are really harms, it is hard to see how paternalism can be distinct from harm-based wrongs.

Similarly, self-sacrifice can be understood as consensually harming oneself, or accepting some harm, to further a goal one cares about more than the avoidance of that harm.[15] *Harm as Agential Impairment* suggests that this is incoherent. Since sacrifice is often an exercise of will rather than an obstruction of it, it cannot also be a harm. Any reduction in wellbeing one suffers in making the sacrifice is not enough: where there is no impairment of agency, there is no harm.

[13] *Reasonableness and Risk*, 15.

[14] *Reasonableness and Risk*, 11.

[15] Victor Tadros uses this argument against the harm principle in the debate about the proper scope of criminalisation. See Victor Tadros, "Harm, Sovereignty, and Prohibition", *Legal Theory*, 17(1) (2011): pp. 35–65, at p. 53.

Yet self-sacrifice seems conceptually coherent and a fairly common practice. Keating could respond by saying again that self-sacrifice involves cost but not harm, but it mischaracterises the significance of self-sacrifice to say that the person who undergoes it is not harmed and may even be benefitted by the furtherance of their will.

As before, these observations present interpretive as well as normative challenges. Compensation for pain and suffering as a separate item of damages is hard to explain if we maintain a strict adherence to *Harm as Agential Impairment*. We would expect damages to be more exclusively keyed to agency-impairment rather than wellbeing. Also, while forced benefits will often give rise to liability in tort (a surgeon who removes a mole without consent during an unrelated surgery commits a battery, for example), it is distorting to suggest that such cases necessarily involve harm by virtue of lacking consent.

Another response based on Keating's characterisation of his view is that harm involves being in a *condition* of impairment–it is not enough that a single event undermines a person's agency. However, it is hard to draw a principled line here. Surely it makes more sense to say that the surgeon causes an isolated harm–removing one choice–while ongoing physical injuries cause harm more systematically. It seems *ad hoc* to argue that the distinction between systematic and isolated interferences with agency marks the difference between harm and non-harm rather than merely a difference of scale. If this is right, this view of harm still does not allow us to accurately capture the wrong of paternalism or the harm of self-sacrifice.

So far, we have cast some doubt on claims (2)–(4) of Keating's harm/benefit asymmetry. This leaves (5): *Non-comparativism*. *Non-comparativism* does not follow from *Doing and Allowing* simply because the debate between comparative and non-comparative accounts of harm is orthogonal to the debate about *Doing and Allowing*. Accepting the latter doctrine does not commit one to measuring or identifying harms comparatively or non-comparatively. This is just as well as *Non-comparativism* runs into serious problems. Take Matthew Hanser's example of the Nobel prize-winner who suffers a minor stroke, leaving them less able than they were before but still with well above average intelligence.[16] *Non-comparativism* implies, counterintuitively, that the stroke does not harm them because it does not cause them to be in a non-comparatively bad state. If anything, we are compelled to say that it benefits them by causing them to be in the non-comparatively good state of having above average intelligence. Or consider Victor Tadros's example of giving a small amount of food to a starving person which still leaves them in a sickly condition.[17] Non-comparativism implies, again counterintuitively, that the food does not benefit them because it causes them to be in the non-comparatively bad state of being sickly with hunger. If anything, we are compelled to say that it harms them by causing them to be in this state. We cannot capture the sense in which the prize-winner is harmed, or the starving person is benefitted, without comparing their current state to the one they were in before the relevant event, or the one they would now be in had it not occurred.

Again, these observations present both normative and interpretive problems. Comparative judgments are essential for reaching accurate conclusions about actionable damage. For example, if a severely hearing-impaired plaintiff suffers a negligent blow to the head which improves their hearing, leaving them in a slightly impaired state, this impairment is not compensable even though it is a harmful condition. Similarly, if a plaintiff has a valid claim for compensation, it is no legal response that they are very well off and therefore left in a non-comparatively good state even after the injury. Of course, it is open to defenders of non-comparativism to find ways of explaining these results (the literature on harm is full of counterexamples and bug-fixes) but the point is that the challenge is a systematic one. Appeal to comparative judgement seems baked in to our everyday–and legal–assessment of harm.

Let's now leave this discussion and turn to another core feature of Keating's deontological account, the commitment to objective valuation, where we will draw similar conclusions.

[16] See Matthew Hanser, "The Metaphysics of Harm", *Philosophy and Phenomenological Research* 77 (2008): p. 432 at p. 431.

[17] Victor Tadros, "What Might Have Been", in John Oberdiek (ed.), *Philosophical Foundations of the Law of Torts* (Oxford: Oxford University Press, 2014), pp. 177–8.

## III. OBJECTIVE AND SUBJECTIVE VALUATION

Keating defends a conception of autonomy at odds with that adopted by his main foil, the economic theory.[18] According to the latter, autonomy is understood as consumer sovereignty, which holds that we respect people's freedom when we respect their preferences. In the realm of accidents and personal injury, this conception of autonomy would favour the personalised protections of contract. Contract mechanisms would allow individuals greater scope to choose the extent of the protections they wish to enjoy against unintended harm. Keating's deontological view, by contrast, defines autonomy as freedom to pursue one's conception of the good under favourable conditions. This leans towards the inalienable safeguards of tort, which permit less individual choice about protection from harm but guarantee certain rights without requiring people to make contractual arrangements.

At the heart of this disagreement is a distinction between two different methods of valuing harms and benefits. According to a subjective valuation, what is good or bad for a person depends on that person's assessment of what is good or bad for them. According to an objective valuation, something may be good or bad for a person irrespective of their preferences or personal assessment. It is important to note that there are two separate questions regarding the distinction between subjective and objective valuation. One is the ethical question of which is the correct account of wellbeing or what makes a person's life go better. The other is the legal or political question of which should be implemented within a coercive institutional framework like tort law. Here we are interested in the latter question. When the law needs to apply one of these methods of valuation, which does and should it apply? Keating comes down firmly in favour of objective valuation.

However, it is not clear that there is any deep inconsistency between subjective valuation and the best deontological understanding of freedom. To illustrate this, let's consider a theory that puts individual preferences front and centre but is appropriately regarded as deontological. Ronald Dworkin's account of distributive justice fits the bill. It aims to consult rather than usurp[19] individual

[18] *Reasonableness and Risk*, 150.

[19] For this terminology, see Tom Parr, "How to Identify Disadvantage: Taking the Envy Test Seriously", *Political Studies* 66(2) (2018): pp. 306–322.

preferences while embodying a form of equality that is in deep conflict with consequentialist views such as utilitarianism and wealth maximisation.[20]

Dworkin imagines a hypothetical auction in which people are given equal purchasing power to bid for resources that suit their particular plans and projects. From this equal starting point, individuals may end up with more or less as a result of their responsible choices, some investing wisely and others not, some indulging in expensive tastes and others not. To the extent that these inequalities result from responsible choices–or 'option luck' to use Dworkin's term–this is acceptable. However, Dworkin also recognises that many disadvantages–disability, ill-health, unmarketable talents and other setbacks–are not rooted in responsible choices and nevertheless have a significant impact on people's life prospects. They represent bad 'brute luck.' Hence, he supplements the auction with a hypothetical insurance scheme in which people are informed of the prevalence of various disadvantages but not whether they will end up suffering any of them. They can then choose from behind this thin veil of ignorance (again from a position of equal purchasing power) whether to insure against any of these potential disadvantages based on their subjective valuations.

This sketch of Dworkin's view of distributive justice shows how the core tension between deontological and utilitarian or economic theories need not be identified as the latter's emphasis on subjective valuation. Dworkin's argument centralises something akin to 'consumer sovereignty' without a commitment to maximising consequentialism. It does this by integrating an emphasis on individual preference and choice into a framework of equality of resources. We ought to distinguish, then, between views that implement subjective valuation on the one hand, and the choice between distributional mechanisms that focus on maximisation and those that favour equality on the other.

As with the harm/benefit asymmetry, the fact that objective valuation is not required by a deontological account of tort law

[20] Ronald Dworkin, *Sovereign Virtue: The Theory and Practice of Equality* (Harvard: Harvard University Press, 2002), especially Ch. 2.

becomes important given that there is independent reason to question Keating's reliance on it. Keating considers objective valuation the proper metric for assessing rules of conduct in tort law. For example, he notes that the economic commitment to subjective valuation conflicts with the average reasonable person doctrine in negligence. The reasonable person of negligence law is "both standardized and idealized,"[21] employing objective criteria to determine reasonable care based on substantive claims about what goods and opportunities make life better.[22] Keating adds that this rejection of subjective valuation has a "deep normative justification."

To support this claim, he observes that it is not permissible to impose risks on others for objectively worthless or low-value ends. Teenagers who linger on train tracks and leap to safety at the last moment may subjectively value the thrill of this activity, but this does not justify the risks it imposes on the train's operators or passengers. However, there are two possible explanations for this judgment. One is that favoured by Keating: the activity is not objectively valuable enough to justify the risks. An alternative explanation is the Dworkinian one that, as a matter of distributive fairness, those who have resource-intensive or high-risk preferences cannot externalise the costs of these preferences onto others. This is a variation of Dworkin's expensive tastes objection to equality of welfare. Why reject such dangerous tastes by appeal to an objective account of value rather than a Dworkinian account of responsibility?

Keating also argues that subjective valuation "licenses deeply nonreciprocal risk impositions and puts the reasonable at the mercy of the unreasonable."[23] This is because under a subjective valuation, those who "moderate their self-interest and rein in their egos" confer the benefit of this restraint on others while being made to bear the brunt of others' lack of restraint towards them. There are a few problems with this argument. One is the alignment of objective reasonableness with self-restraint. This assumes a particular reason for failing to adhere to the average reasonable person standard: self-restraint. However, unreasonableness consists in falling below the standard *tout court*. It does not require any specific reason or failing as the cause of falling below the standard. Many people fall below it for

[21] *Reasonableness and Risk*, 167.
[22] *Reasonableness and Risk*, 168.
[23] *Reasonableness and Risk*, 172.

reasons they are not morally responsible for, such as possessing lesser than average capacities.

A second worry is that self-restraint is not strictly a form of subjective valuation. The person who is willing to impose a risk on another for some end does not necessarily value that end more than someone who is not willing to do so. Those who respect others do not for that reason value their ends less than those who do not. A person's valuation of an end is better determined by the costs that they themselves would be willing to bear for the sake of that end, not the costs they are willing to impose on others.

A final problem is the one we encountered above. An advocate of subjective valuation need not hold that merely valuing risk-imposing activities to a high degree makes them permissible or worthy of consideration from a legal perspective. This would be a type of expensive taste and is one of the reasons Dworkin thinks distributive equality is best realised by allowing people to make their own choices about how to spend their resources against a baseline of equal purchasing power rather than trying to achieve equality of welfare by equalising preference-satisfaction. Taking all these points together, it seems a subjective account of valuation can yield the same result as that which Keating claims for the objective account.

Keating has other criticisms of subjective valuation. He says that implementing it would lead to haphazard and unpredictable risks and would erode trust and co-operation. In the long term, this might undermine the capacity of the law to guide conduct effectively. We can even add to Keating's argument here. Perhaps it is impossible to measure a person's subjective valuation of an end, and the degree of risk they are willing to impose on others is the best available proxy, so the average reasonable person standard is far more practicable.

I regard these as practical concerns appealing to the consequences of enshrining subjective valuation within the law. I am not sure they contribute to a 'deep normative justification' of objective valuation, especially if this is understood in deontological terms. If subjective valuation is justified in principle, this means we will be forced to make trade-offs between values such as practicability and stability on the one hand and respect for individual autonomy on the other. Defenders of subjective valuation need not deny this. But equally, we cannot assume from the outset that these conflicts will always come out in favour of the former values. Even when they do, we should be cognisant of the distinction between normative justification and practical implementation here. Dworkin's model of distributive justice, which I have been using as an illustration, is an exercise in 'ideal theory.'[24] Actual distributive justice involves the state intervening to mimic the outcomes of these hypothetical markets, which no doubt involves its share of practical compromise, epistemic uncertainty and balancing with other important values. Similarly, proponents of subjective and objective forms of valuation might well converge on something like the average reasonable person standard in practice. So, if Keating's defence of objective valuation ultimately rests on these considerations, this may point more to a practical justification of the average reasonable person standard (which proponents of subjective valuation might also accept) than to a deep normative one.

Let me conclude this section by stating that I think the relevance of subjective valuation to tort law remains an open and complex question. In some cases, it is plausible that subjective valuation should be rejected. Suppose a negligent actor destroys someone's arts studio, but the owner has ceased to value it due to crippling self-doubt about their artistic ability. Clearly this fact has no normative relevance to whether the negligent actor should pay compensation. I doubt this intuition is due only to the need to maintain a stable regime of property rights, or the difficulty of interrogating a person's values. It is due also to an underlying conviction that the victim's property is in some sense a benefit that should be protected despite their current preferences.

However, in other cases the matter is not so clear. In *McFarlane v Tayside Health Board*,[25] for instance, the House of Lords ruled that where a doctor's negligence in performing a sterilisation procedure resulted in the plaintiff having a baby that she would not otherwise have had, the plaintiff could not sue for the cost of bringing up the child. This decision was reached, in part, because the judges believed that the birth of a healthy child was a benefit that offset the financial

[24] Rawls's famous and often criticised statement of the importance of ideal theory can be found in *A Theory of Justice* (Cambridge, MA: Harvard University Press, 1971, Revised Ed.), sec. 2,8. For a more recent defence of the role of ideals within moral and political philosophy, see David Estlund, *Utopophobia: On the Limits (If Any) of Political Philosophy* (Princeton: Princeton University Press, 2019).

[25] [2000] 2 AC 59.

burden of upbringing.[26] The question whether it was acceptable for the court to impose on the parents a judgment about their wellbeing that might be in deep conflict with their subjective valuation was not addressed. It is arguable that, even if the court was right in its objective assessment of the value of parenthood, it was entirely inappropriate to impose this on plaintiffs who may have rejected it in favour of their own reasonable conception of wellbeing.[27] These observations suggest that framing the debate between two general paradigms of subjective and objective valuation is something of a simplification. Perhaps subjective valuation is sometimes but not always warranted, depending on other circumstances such as the foreseeability of subjective harm or the costs subjective valuation imposes on others, including the wrongdoer.[28]

## IV. STRICT LIABILITY

There is a view commonly attributed to advocates of the law and economics movement which holds that, contrary to appearances, tort law is not really a law of wrongs. Liability is in fact a price attached to certain activity which aims to facilitate the future maximisation of wealth. The language of wrongdoing when applied to tort law is thus systematically in error. The conduct the law considers tortious is not wrong or impermissible; it must merely be paid for in the form of incentive-generating liability.

Many non-instrumentalist tort theorists reject this view,[29] and in doing so articulate a multifaceted conception of wrongdoing that can explain the seemingly heterogenous forms of liability found in tort law. Keating takes the same strategy, although instead of minimising or rejecting strict liability because it does not fit into a fault-based paradigm, he explains it in terms of a distinct form of wrongdoing. Keating divides strict liability into two types: sovereignty-based and harm-based. Harm-based strict liability represents the wrong of harming and failing to compensate. Keating applies this analysis to

[26] See Lord Steyn at p. 82 and Lord Millet at p. 114.

[27] I focus here on this particular argument. I acknowledge, of course, that the conclusion in McFarlane might be justified by other reasons. I also say the claimants 'may' have rejected this view because, once the child was born, they accepted it as an integral part of their family.

[28] For more discussion, see Adam Slavny, "On Being Wronged and Being Wrong", Philosophy, Politics and Economics 16(1) (2017): 3–24.

[29] For a recent example, see John CP Goldberg and Benjamin C Zipursky, *Recognizing Wrongs* (Harvard: Harvard University Press, 2000).

the well-known case of *Vincent v Lake Erie Transportation Co*. In this case, although it was permissible for the defendant to tie a steamship to an exposed dock in a storm because they were acting out of necessity, they nevertheless incurred liability to compensate for the damage.[30] This result cannot be explained by positing the breach of a prior duty as the appropriate ground for liability. Instead, Keating argues that the relevant wrong was not the act of tying the ship to the dock, but the failure to compensate when this resulted in damage.

By including strict-liability within the fold of a wrongs-based conception of tort law, Keating rejects one of the core commitments of the economic approach and reinforces the idea that a deontological theory will treat tort as a law of wrongs through and through. There are a few problems with Keating's wrongs-based analysis, though. One is that appealing to it to explain and justify harm-based strict liability seems redundant. If it is wrong to harm without offering compensation in certain cases, this must be because there is a pre-existing duty to offer compensation. If there were no such duty, it would be difficult to see why withholding compensation would be wrong. But then why not appeal directly to this duty to explain and justify liability? The addition of the wrong of harming and failing to compensate adds little to this explanation. If the defendant in *Vincent* had stepped forward to offer compensation, they would not have committed any wrong. Yet their duty to compensate would be the same as that which they incurred by committing the wrong. If this is so, why not appeal directly to the duty to explain liability in all cases, including those in which the defendant withholds?

Another reason to think that the wrong is explanatorily redundant in cases of harm-based strict liability is that the same wrong is

[30] John Goldberg disputes that necessity justifies or even excuses the violation of property rights. "Tort Law's Missing Excuses", in Andrew Dyson, James Goudkamp, and Frederick Wilmot-Smith (Eds.), *Defences in Tort* (Oxford: Hart, 2015), p. 54. On Goldberg's view, one may act to defend one's property from trespass within the bounds of reasonableness, although this may not extend to ejecting a trespasser as this may be disproportionate. As John Oberdiek argues, besides interpretive concerns, the central problem with this move is that it puts the right to exclude above the right to safety and bodily autonomy. See the "The Trouble with Trespass", in Leslie Green and Brian Leiter (Eds.), *Oxford Studies in Philosophy of Law*, Volume V (Oxford: Oxford University Press, forthcoming 2024). More than this, the claim that necessity does not justify trespass implies that I have a duty to allow myself to freeze rather than enter your property, but if I do enter (illegally) you have a duty not to kick me out if doing so is disproportionate. This is an odd combination of claims, as it leaves the law-abiding citizen to the mercy of the elements whilst protecting the trespasser from losing the fruits of their wrongdoing.

committed in many other circumstances that generate liability. Most people who have committed torts do not offer compensation outright. They must be forced to pay it either through legal action or threats of legal action. Despite its broad scope, Keating does not rely on the wrong of failing to pay one's due in explaining other forms of liability, especially fault liability. But why not? Refusing to compensate for one's negligent conduct is, if anything, a greater wrong than refusing to compensate for one's justified but harmful conduct. True, a defendant's conduct in the pre-trial process can have implications for their position if the matter comes before a court, but initial refusal to compensate does not usually have the significance that Keating attributes to it in cases of harm-based strict liability.[31]

Of course, Keating does not merely assert that defendants have duties to compensate in these cases. These duties are themselves justified with reference to a theory of enterprise liability, which holds roughly that it is fair to make those who undertake risky endeavours internalise the costs they impose on others. I will not scrutinise this theory of liability here. The point is that, if this theory is correct, it is better to explain the liability it generates by appealing to it directly. Failing to do this not only invites the charge of redundancy, it also risks obscuring the grounds for liability. In negligence cases, a defendant is liable if *and because* they committed a wrong. It is the fact that they did something they were not legally permitted to do that explains their duty to pay. This explanation is not available in cases like *Vincent* where the defendant acted permissibly. If, in these cases, we appeal to the distinct wrong of harming without compensating, we obscure the fact that this wrong is not a necessary ground for liability because the defendant would owe a duty to compensate whether they committed it or not.

Another question about Keating's account of strict liability concerns his treatment of sovereignty-based liability. Keating recognises that torts such as conversion, trespass and battery can be committed innocently, or at least with evidence-relative justification. That is to say, it is perfectly possible for a person to be liable for one of these torts despite the fact that, on any reasonable appraisal of the evidence available to that person at the time, their conduct was per-

[31] For example, under the Civil Procedure Rules of England and Wales, a defendant can make a 'part 36' offer to settle. A claimant who rejects the offer but fails to obtain a higher amount in damages can be liable to pay the defendant's costs, see CPR 36.

missible. Nevertheless, strict liability for such 'in-fact-rights-violations' is justified. There are two parts to Keating's argument for this conclusion. The first is that, even if the agent could not reasonably have been expected to avoid the act in question, it is still a rights violation and rights are worth having. Doing away with them whenever the behaviour of others is evidence-relative justified would diminish the power these rights grant us over our persons and property.

This first point does not by itself settle whether we should recognise 'in-fact-rights-violations.' This is because, whilst recognising such violations does indeed enhance our power over our persons and property, it also imposes costs on others for conduct that was entirely justified from their perspective, and this is also undesirable. What we have here are two competing normative considerations, and we require a further argument to determine which should win out. We cannot simply assert that avoiding diminishment in our powers of control must take priority.

This is where the second facet of Keating's argument comes into play. He says that "requiring someone whose rights were violated through neither fault nor action on their part to suffer the unrepaired consequences of that wrong is surely more unjust than holding the party who innocently wronged them responsible for having done so."[32] One question about this rationale is whether there is any deeper justification for it. Why is it more unjust to let the consequences lie where they fall? It cannot always be because the victim is not at fault, as neither is the agent of the harm in many strict liability cases. It cannot always be because the agent acted while the victim did not, as most accidents result from the intersection of a chain of actions performed by both parties. It cannot always be because the agent could more easily have avoided the harm, as this might not be true if the agent's action is evidence-relative justified, and the victim should have foreseen that the harm was likely to occur. Perhaps the key word here is 'occasioned,' and the explanation is rooted in the moral significance of causation. It is the agent and not the victim who causes the harm, and this makes it fair that, as between the two of them, the agent should be responsible for the consequences. My own view is that this claim is plausible up to a point, but one

[32] *Reasonableness and Risk*, 245.

problem with it as an interpretation of Keating's fairness argument is that it relies on a claim about the nature and moral significance of causation which is probably more controversial than the initial appeal to fairness it was supposed to explain.

Let me make a final point about the scope of this rationale for sovereignty-based strict liability. If it is more unjust to force the victim of an 'in-fact-rights-violation' to bear it than holding the one who caused it responsible, why not apply this reasoning much more generally? There are many fact-relative wrongs that do not generate liability. This is true of the majority of non-negligent harming. Careful driving, good medical practice, safe working environments, and many other types of reasonable activities can unfortunately result in harm. When they do, a fact-relative wrong is usually committed, which is to say that, had the injurer known the results beforehand, it would have been wrong for them to proceed. Their action is evidence-relative justified but fact-relative wrong. Many of these harms do not generate sovereignty-based liability, though some may be covered by strict liability regimes. If the justification for sovereignty-based liability is as Keating states it, it is hard to see why. Surely it is more unjust that the victim of an innocent road accident, or a reasonable medical mistake or indeed any form of non-negligent harm, be left to bear the burden than to hold the person who caused the accident responsible. In other words, if this rationale is valid, it seems less a justification for a subcategory of strict liability and more a catch-all justification for replacing all forms of tort liability with strict liability.

## V. CONCLUSION

Deontologists reject consequentialism, and certain features of Keating's account of tort law, such as *Doing and Allowing* and the identification of wrongdoing (independently of consequences) as a key ground for compensatory liability, fit naturally within a deontological framework. The connection between deontology and other elements of Keating's account is less obvious, however. For example, we saw that consequentialists could adopt a more sophisticated theory of value than that offered by wealth maximisation. More importantly, in other areas, there are deontological alternatives to Keating's view which are more plausible, as the discussion of the harm/benefit asymmetry and strict liability suggests. So while Keating's contribution is very welcome insofar as it moves us in the direction of a sound deontological theory of tort law, I'm not sure it gets us all the way.



*School of Law*
*University of Warwick, Coventry, CV47AL, UK*
*E-mail: a.slavny@warwick.ac.uk*