Permissibility and the Aggregation of Risks

JAMES R. KIRKPATRICK
University of Oxford

Tom Dougherty has recently argued that non-consequentialists cannot accommodate our judgements about acceptable levels of risk-imposition. More specifically, he argues that the following two intuitively plausible claims are inconsistent: (i) that it is impermissible to provide small benefits to many people rather than saving the life of someone else, and (ii) that it is permissible to expose someone to a negligible risk of death in order to otherwise provide this person with a small benefit. Abandoning either principle has significant consequences: rejecting (i) requires rejecting an important argument against consequentialism and consequentialist approaches to beneficence; rejecting (ii) requires radically rethinking the way we live our lives, as we routinely expose individuals to negligible risks of death. This article shows that Dougherty's argument relies on a scope ambiguity involving permissibility. Once this ambiguity is resolved, Dougherty's argument fails.

I. INTRODUCTION

Common-sense morality is perfectly clear that, when faced with the choice to cure a billion people's headaches or to save another person's life, one ought to save the life. Such judgements motivate non-consequentialist approaches to beneficence which reject the claim that small benefits to sufficiently many people may aggregate to form a total amount of welfare greater than substantial benefits to a few.[1] If common-sense morality is correct, then consequentialist explanations of beneficence should be rejected.

Common-sense morality is also perfectly clear that exposing someone to a negligible risk of death in order to cure their headache is permissible. Such judgements underpin a general permissive attitude towards risk-imposition according to which we may impose minute risks of harm to provide small potential benefits.[2] Given the ubiquity of these seemingly permissible risk-impositions, if common-sense

[1] See, for example, T. M. Scanlon, *What We Owe To Each Other* (Cambridge, MA, 1998).

[2] See, for example, J. J. Thomson, *Rights, Restitution, and Risk* (Cambridge, MA, 1986); J. Lenman, 'Contractualism and Risk Imposition', *Politics, Philosophy and Economics* 7 (2008), pp. 99–122. For more recent attempts to reconcile non-consequentialism with the intuitive permissibility of acceptable levels of risk-imposition by defending versions of *ex ante* contractualism, see R. Kumar, 'Risking and Wronging', *Philosophy and Public Affairs* 43 (2015), pp. 27–51; J. Frick, 'Contractualism and Social Risk', *Philosophy and Public Affairs* 43 (2015), pp. 175–223. Considering whether such versions of contractualism adequately reconcile non-consequentialist judgements with permissive risk imposition is outside the scope of this article, since Kumar and Frick

© Cambridge University Press 2017
Utilitas Vol. 30, No. 1, March 2018
doi:10.1017/S0953820817000036

CrossMark
https://doi.org/10.1017/S0953820817000036 Published online by Cambridge University Press

morality is incorrect on this point, then it is hard to justify many of our everyday activities.

In a recent paper, Tom Dougherty examines the plausibility of two principles which explain these aspects of common-sense morality:[3]

> *Many–Few*: You may not provide small benefits to many people rather than save the life of someone else. (All else being equal.)
>
> *Risk Tolerance*: You may expose someone to a negligible risk of death in order to otherwise provide this person with a small benefit. (All else being equal.)

However, despite their initial plausibility, Dougherty argues that Many–Few and Risk Tolerance are jointly inconsistent. Moreover, Dougherty claims that his argument is formulated in terms with which non-consequentialists should have no complaint.[4] Consequently, if Dougherty's argument is correct, then it highlights a perplexing conflict between our intuitions concerning the aggregation of benefits and our intuitions concerning risk-imposition: either the non-consequentialist judgements explained by Many–Few are incorrect or our everyday attitudes towards permissible levels of risk-imposition are unjustified.[5]

The aim of this article is to show that Dougherty's argument fails to establish the inconsistency of Many–Few and Risk Tolerance. More specifically, it will argue that Dougherty's argument trades on a scope ambiguity involving permissibility. Once this ambiguity is resolved, the purported inconsistency dissolves. In support of this conclusion, this

reject a central principle to be introduced shortly in favour of alternative principles. Instead, this article aims to show that a natural way of capturing the above aspects of common-sense morality are compatible in their current form, rather than to entertain alternative principles.

This article treats the risk of some event occurring as the probability of that event occurring multiplied by the potential harm. It also adopts an epistemic understanding of probability, although nothing hinges on whether we have a frequentist or objective understanding of probability. For a survey of different characterizations of 'risk', see M. Hayenhjelm and J. Wolff, 'The Moral Problem of Risk Impositions: A Survey of the Literature', *European Journal of Philosophy* 20 (2012), pp. 26–51.

[3] T. Dougherty, 'Aggregation, Beneficence and Chance', *Journal of Ethics and Social Philosophy* 7 (2013), pp. 1–19.

[4] For a similar argument for affirming an analogue of Risk Tolerance and rejecting an analogue of Many–Few which presupposes a consequentialist approach to the ethics of risk-imposition and aggregation of benefits, see A. Norcross, 'Great Harms from Small Benefits Grow: How Death Can Be Outweighed by Headaches', *Analysis* 58 (1998), pp. 152–8. However, Norcross's argument is dialectically ineffective against non-consequentialists, since they will simply reject the consequentialist assumptions underpinning his argument. See Dougherty 'Aggregation, Beneficence and Chance', pp. 3–4.

[5] This article follows Dougherty in assuming that our common-sense moral judgements are to be explained by Many–Few and Risk Tolerance. Consequently, if either principle is false, then so are the judgements which they explain.

https://doi.org/10.1017/S0953820817000036 Published online by Cambridge University Press

article will also seek to establish some general principles concerning permissible risk-imposition.

II. THE REPETITION ARGUMENT

Dougherty gives an argument that Many–Few and Risk Tolerance are inconsistent which he calls the *Repetition Argument*. His argument is so-called because it examines the likely effects of repeating risky actions. This section will outline the Repetition Argument.

Dougherty begins by asking us to consider the following scenario:

> *Poison*: A large number of people each have swallowed poison. If nothing is done, each victim will have a headache and then die. Fortunately, you can bring each victim one of two antidotes. Taking the first antidote, Reliable Antidote, is certain to save the victim's life, but will do nothing for their headache. Taking the second antidote, Chancy Antidote, is extremely likely to cure the victim's headache and save their life, but there is a one-in-a-billion chance that it will do nothing.[6]

Dougherty then reasons as follows. By Risk Tolerance, you are permitted to bring a victim Chancy Antidote. If one is unconvinced, the size of the risk that Chancy Antidote may do nothing can be altered until one is satisfied that Risk Tolerance entails that it is permissible to bring a victim Chancy Antidote. Consequently, Risk Tolerance entails (1):

(1) You may bring Chancy Antidote to a victim in Poison once.

Since you may bring Chancy Antidote to a victim in Poison once, Dougherty argues that you may bring Chancy Antidote to any number of people in a similar situation: 'If you may bring the Chancy Antidote to [one victim], then you may also bring the Chancy Antidote to [that victim] and then bring it to [another victim]'.[7] Consequently, Dougherty argues that Risk Tolerance entails (2):

(2) You may bring Chancy Antidote to any sequence of people in Poison.

Given (2), by repeatedly administering Chancy Antidote, it follows that for every billion people to whom you bring Chancy Antidote, we would expect that you would fail to save someone's life while everyone else's

[6] Dougherty, 'Aggregation, Beneficence and Chance', p. 5.
[7] Dougherty, 'Aggregation, Beneficence, and Chance', p. 5.

https://doi.org/10.1017/S0953820817000036 Published online by Cambridge University Press

lives are saved and their headaches cured. As a result, (2) has the problematic consequence of entailing (3):[8]

(3) You may cure the headaches of 999,999,999 people rather than saving the life of someone else.

The trouble is that Risk Tolerance entails (1) and Many–Few contradicts (3). Consequently, Dougherty concludes that these principles are inconsistent.

Dougherty's argument is deeply worrying. For if we reject Many–Few, then we must call into question our judgement that we may act to benefit the few even when the aggregate benefits to a larger group of individuals is greater. In turn, revising this judgement calls into question a powerful argument against consequentialist approaches to ethics. And if we reject Risk Tolerance, then we are forced to revise our judgements about everyday risky actions. In particular, we would have to avoid any beneficial action that would expose potential beneficiaries to a small risk of serious harm. Revising our actions in this way would be severely limiting.

Dougherty attempts to resolve this problem by searching for alternative principles that capture our intuitions about risk-imposition and aggregation. By his own lights, he is 'unsuccessful in discovering any such principles'.[9] He then throws down the gauntlet for others.

III. A SCOPE AMBIGUITY IN PREMISE (2)

This section contends that Dougherty's problem need not be resolved by searching for alternative principles to Many–Few or Risk Tolerance. Rather, the solution is that the Repetition Argument trades on an ambiguity, namely, that (2) is ambiguous. Once (2) is disambiguated, the purported inconsistency between Many–Few and Risk Tolerance dissolves: the Repetition Argument provides no reason to think that Many–Few and Risk Tolerance are not both true.

Let us make this argument clearer by adopting the following formalization to illuminate the ambiguity of (2). Suppose that there is a finite number of people $N = 1, 2, 3, \ldots, n$, each of whom is a victim in Poison. Let $p_1, p_2, \ldots, p_n$ be, respectively, the sentences 'Person 1 is administered Chancy Antidote', 'Person 2 is administered Chancy Antidote', and so on. Let $A\varphi$ be the sentence 'You bring it about that φ' and let $\blacklozenge\varphi$ be the sentence 'It is permissible that φ'.[10] Then we may

[8] Dougherty's original argument involves eight premises. Since this discussion does not turn on those details of Dougherty's argument, for simplicity, they are suppressed.

[9] Dougherty, 'Aggregation, Beneficence and Chance', p. 18.

[10] For ease of exposition, I assume that 'It is permissible that φ' is equivalent to 'It may be that φ', where 'permissible' and 'may' have a deontic flavour, rather than epistemic or

https://doi.org/10.1017/S0953820817000036 Published online by Cambridge University Press

formulate (2) in two ways, depending on the scope of 'may'. First, we can assign 'may' a narrow scope reading, such that you may bring it about that person 1 is administered Chancy Antidote, that you may bring it about that person 2 is administered Chancy Antidote, and so on:

(2_{NARROW}) $\Diamond Ap_1 \,\&\, \Diamond Ap_2 \,\&\, \ldots \,\&\, \Diamond Ap_n$.

The second reading assigns 'may' a wide scope reading, such that you may bring it about that persons 1 to n are administered Chancy Antidote all at once:

(2_{WIDE}) $\Diamond(Ap_1 \,\&\, Ap_2 \,\&\, \ldots \,\&\, Ap_n)$.

On the one hand, (2_{WIDE}) has the problematic consequence of entailing (3) and thus conflicting with Many–Few. For if you are permitted to give Chancy Antidote to the first person, to the second person, and so on, then, as Dougherty argues, you are permitted to cure the headaches of 999,999,999 people rather than saving someone's life. Given Many–Few, you are not permitted to cure the headaches of 999,999,999 people rather than saving someone's life. Consequently, if (2) is read as (2_{WIDE}), it generates a seriously problematic consequence. On the assumption that Many–Few is true, this is good reason for thinking that (2_{WIDE}) is false.

Furthermore, there is independent motivation for rejecting (2_{WIDE}) because it conflicts with the following highly plausible principle:

> *Constraint on Permissible Risk Imposition*: If the risk of serious harm from an action is sufficiently high (and is greater than the risk of serious harm from performing any relevant alternative to that action), then it is impermissible to perform that action.

If you give everyone Chancy Antidote, then it is near certain that someone will die, even though it is not the case that there is someone for whom it is near certain that they will die. Exposure to near certain death runs a sufficiently high risk of serious harm. Consequently, given Constraint on Permissible Risk Imposition, it is impermissible to give everyone Chancy Antidote. So, the fact that (2_{WIDE}) conflicts with Constraint on Permissible Risk Imposition, as well as Many–Few, is good independent reason for thinking that (2_{WIDE}) is false.

On the other hand, (2_{NARROW}) does not have this problematic consequence, since (2_{NARROW}) does not entail (2_{WIDE}) and so does

metaphysical. Nothing substantial hinges on this assumption, and we may read $\Diamond\varphi$ as 'It may be that φ' if we so wish. Greek letters are metavariables ranging over sentences; quotation marks containing such sentences should be read as quasi-quotations.

https://doi.org/10.1017/S0953820817000036 Published online by Cambridge University Press

not entail (3). This is because deontic permissions do not necessarily agglomerate; that is, deontic permissions do not satisfy the following principle:

Agglomeration: If $(\blacklozenge\varphi \ \& \ \blacklozenge\psi)$ then $\blacklozenge(\varphi \ \& \ \psi)$.

In other words, given the falsity of Agglomeration, if it is permissible that φ and it is permissible that ψ, it does not follow that it is permissible that φ and ψ. Such failures of Agglomeration occur in single-agent cases when two independently permissible actions, each of which affect only one and the same individual, are collectively impermissible. Consider the following case:[11]

> *Insulin*: A is suffering from insulin shock. She needs insulin, otherwise she will die. B has an insulin shot which B may permissibly give to A and C has an insulin shot which C may permissibly give to A. However, if both B and C give insulin shots to A, A will become hyperglycaemic and die.

In Insulin, it is permissible for B to give an insulin shot to A and it is permissible for C to give an insulin shot to A, since either action will save A from insulin shock. But it is impermissible for both B and C to give A each of their insulin, since this conjunctive action will cause A to become hyperglycaemic and die. Consequently, Insulin is a case where agglomeration fails.

Failures of Agglomeration also occur in multi-agent cases when two or more independently permissible actions, each of which affect different individuals, are collectively impermissible. Consider the following case:

> *Lifeboat*: Ten people are in dangerous water and will each probably drown if you do not save them. No individual has a special claim to be rescued. There is a small, but non-zero, chance that another lifeboat will arrive to save them, but this is no certainty. Your lifeboat has a capacity of ten people including yourself. If the boat is overcapacity, it will disintegrate and everyone will certainly die.

In Lifeboat, for each individual, it is permissible to save that individual from drowning. But it is impermissible to save every individual, since doing so will cause the certain death of everyone including yourself, even though there is no individual for whom it is impermissible for you to save. Consequently, Lifeboat is a case where Agglomeration fails.

Given the general failure of Agglomeration, it should be no surprise that you may give Chancy Antidote to person 1, and you may give

[11] Thanks to Joseph Bowen for suggesting Insulin and Lifeboat.

https://doi.org/10.1017/S0953820817000036 Published online by Cambridge University Press

Chancy Antidote to person 2, and so on, even though it does not follow that you may give Chancy Antidote to person 1, and to person 2, and so on. Consequently, if (2) is read as (2_{NARROW}), it neither entails (2_{WIDE}) nor (3). That is, it does not follow from (2_{NARROW}) that you may cure the headaches of 999,999,999 people rather than save the life of someone else.

Risk Tolerance provides good reason for thinking that the narrow reading of (2) is true. We can also provide independent reason to think that (2_{NARROW}) is true. For the permissibility of exposing someone to a risk does not depend on whether you have previously exposed other people to this risk. This point is captured by the following highly plausible principle endorsed by Dougherty:[12]

> *Only Effects Matter*: When only one person is affected by a risky action, the permissibility of the action depends on the possible or actual effects of *that action* for *that person*.

Given Only Effects Matter, the permissibility of giving Chancy Antidote to an individual in Poison depends upon the possible or actual effects of *that action* for *that individual*. Consider an arbitrary individual in Poison. Since only that individual is affected if we give *her* Chancy Antidote, and the risk of serious harm to her by giving her Chancy Antidote is sufficiently low, then it seems permissible to give her Chancy Antidote. Consequently, Only Effects Matter provides independent reason for thinking that the narrow reading is true.

We are now in a position to present the solution to Dougherty's problem. Dougherty argues that Risk Tolerance and Many–Few are inconsistent because Risk Tolerance entails (2), (2) entails (3) and (3) contradicts Many–Few. However, (2) is ambiguous between (2_{NARROW}) and (2_{WIDE}). Once (2) is disambiguated, we see that the Repetition Argument must implicitly assume the following instance of Agglomeration which takes (2_{NARROW}) as its antecedent and (2_{WIDE}) as its consequent:

> (2_{BRIDGE}) If $(\blacklozenge Ap_1 \ \& \ \blacklozenge Ap_2 \ \& \ldots \& \ \blacklozenge Ap_n)$, then $\blacklozenge(Ap_1 \ \& \ Ap_2 \ \& \ldots \& \ Ap_n)$.

However, proponents of Many–Few and Risk Tolerance need not accept (2_{BRIDGE}), since Risk Tolerance entails that its antecedent is true and Many–Few entails that its consequent is false, and a conditional with a true antecedent and a false consequent is itself false. Moreover, the rejection of (2_{BRIDGE}) is strongly supported by independent arguments for the truth of (2_{NARROW}), the falsity of (2_{WIDE}) and the invalidity of

[12] Dougherty, 'Aggregation, Beneficence and Chance', p. 10.

https://doi.org/10.1017/S0953820817000036 Published online by Cambridge University Press

Agglomeration. Consequently, since the Repetition Argument relies on an appeal to a false instance of the invalid principle of Agglomeration, there is no argument from Risk Tolerance to something that contradicts Many–Few: you may bring Chancy Antidote to a victim in Poison once, but this does not mean that you may bring Chancy Antidote to person 1 and to person 2, and so on. Therefore, the Repetition Argument fails: Many–Few and Risk Tolerance are not inconsistent.

IV. WHY AGGLOMERATION FAILS AND THE LIMITS OF ACCEPTABLE AGGLOMERATION

The previous section argued that the Repetition Argument provides no reason for believing that either Many–Few or Risk Tolerance is false. Advocates of both principles may accept (2_{NARROW}), because it is entailed by Risk Tolerance, and reject (2_{WIDE}), because it contradicts Many–Few, without worrying about inconsistency because Agglomeration fails. However, it would be good to have a general, independent rationale for why Agglomeration fails, especially in the case of permissions to expose people to small risks. A satisfying answer to this question would explain why we should accept the narrow and not the wide reading of Dougherty's second premise.

We do not need anything other than the two principles Constraint on Permissible Risk Imposition and Only Effects Matter to explain the failure of Agglomeration in cases involving risk imposition.[13] To see this, let q_1 and q_2 be two propositions differing at most in the individual whom they are about. Let us assume that the risk of serious harm caused from bringing about q_1 is sufficiently low, and the risk of serious harm caused from bringing about q_2 is sufficiently low, while the risk of serious harm caused from bringing about both q_1 and q_2 is sufficiently high. This may be true even if the risks of the individual events are independent.[14] Given Only Effects Matter, it follows that:

(4) $\blacklozenge Aq_1 \,\&\, \blacklozenge Aq_2$.

[13] Thanks to an anonymous reviewer for pointing out that Constraint on Permissible Risk Imposition and Only Effects Matter are insufficient to explain the failure of agglomeration in other cases, such as when performing both of two harmless actions breaks a promise but performing either (but not both) of them does not. The reviewer suggests the following more general principle to capture such cases: performing two (or more) actions, each of which is permissible, may constitute performing a third, more complex action that is impermissible. While I agree with the reviewer's observation and suggestion, we should focus on Constraint on Permissible Risk Imposition in our current setting, since it better explains *why* the more general principle holds in cases involving agglomeration of risk impositions.

[14] For example, if I have a 1/4 chance of causing serious harm by firing a gun at someone once, I almost double my chances of causing serious harm by firing my gun at two individuals.

https://doi.org/10.1017/S0953820817000036 Published online by Cambridge University Press

Given the highly plausible principle that action is closed under conjunction (that is, if one brings it about that φ and one brings it about that ψ, then one brings it about that both φ and ψ):

(5) If $\blacklozenge(Aq_1 \,\&\, Aq_2)$, then $\blacklozenge A(q_1 \,\&\, q_2)$.

But, by Constraint on Permissible Risk Imposition, it is impermissible to bring about q_1 and q_2, since the risk of serious harm caused by bringing about both q_1 and q_2 is sufficiently high:

(6) $\neg\blacklozenge A(q_1 \,\&\, q_2)$.

By *modus tollens*, it is impermissible to bring about q_1 and to bring about q_2:

(7) $\neg\blacklozenge(Aq_1 \,\&\, Aq_2)$.

The Repetition Argument requires that something like the following substitution instance of Agglomeration is true:

(8) If $(\blacklozenge Aq_1 \,\&\, \blacklozenge Aq_2)$, then $\blacklozenge(Aq_1 \,\&\, Aq_2)$.

However, it follows from (4) and (7) that (8) is false, since it is a conditional with a true antecedent and a false consequent. Since (4) and (7) are entailed by Only Effects Matter and Constraint on Permissible Risk Imposition, respectively, only these principles are required to show how agglomeration fails in the above case. By appealing to these principles, we can make similar remarks about the failure of Agglomeration in the Repetition Argument.

It would also be good to discuss the limits of acceptable agglomeration. More specifically, it seems that the above picture of risk-imposition is committed to the following seemingly problematic claim:[15]

(*) There is some number k, where $1 \leq k < n$, such that one is permitted to give Chancy Antidote to k people sequentially but one is not permitted to give Chancy Antidote to $k{+}1$ people.

This claim is seemingly problematic because it implies the following two claims:

(9) $\blacklozenge(Ap_1 \,\&\, \ldots \,\&\, Ap_k)$.
(10) $\neg\blacklozenge(Ap_1 \,\&\, \ldots \,\&\, Ap_{k+1})$.

Suppose that an agent has already performed the set of actions $(Ap_1 \,\&\, \ldots \,\&\, Ap_k)$, and is now confronted with a scenario where they could

[15] Thanks to Joseph Bowen and an anonymous reviewer for both raising this objection.

https://doi.org/10.1017/S0953820817000036 Published online by Cambridge University Press

perform Ap_{k+1}. If we accept the above claims, then the agent has not yet violated a requirement. The question is whether she would be violating a requirement by performing Ap_{k+1}. The answer is: yes and no. Yes; after all, the agent will be doing something she is not permitted to do, namely, perform the set of actions $(Ap_1 \,\&\, \ldots \,\&\, Ap_{k+1})$. But also no; for although that set of actions is impermissible, it does not follow that performing Ap_{k+1} is impermissible nor that any other particular action is impermissible. In fact, the agent will be doing something that she is permitted to do, namely, performing the action Ap_{k+1}. This should not strike us as unusual. Recall Lifeboat; if I save all 10 individuals, it is clear that I have done something impermissible since we will all die, even though there is no individual for whom it is impermissible to save. Therefore, (*) is not an undesired feature of any theory of risk-imposition.

V. SOME OBJECTIONS

To summarize, this article has argued that once the wide-scope and narrow-scope readings of (2) have been distinguished, the Repetition Argument gives no reason to think that Many–Few and Risk Tolerance are inconsistent. However, the picture of risk that falls out may strike many as unnatural. This section aims to address some objections that might be raised to it.

One may object that my solution embodies a mistaken way of thinking about risk.[16] To see the purported mistake, suppose I am deciding whether to ride my bicycle without my helmet. I judge that there is a small risk of serious harm if I were not to wear my helmet. On the present view of risk, it is permissible for me to ride my bicycle without a helmet *just this once*, since the convenience of not wearing it is worth the risk. However, it also follows from the present view that I am not permitted to continually take these small risks, since it would become likely that one time I suffer serious harm and this is not worth the small inconvenience.

Dougherty objects that this view of risk does not appreciate the 'importance of the independence of the risks in question':[17]

> The correct way to think about the risks is to look at the possible or actual effects of each action independently, and assess the permissibility of each action on these grounds. The permissibility of exposing someone to a risk does not depend on whether you have previously exposed others to this risk.

[16] For an objection along these lines, see Dougherty, 'Aggregation, Beneficence and Chance', pp. 8–11. Thanks to Tom Dougherty for encouraging me to engage with this objection explicitly.

[17] Dougherty, 'Aggregation, Beneficence and Chance', p. 9.

https://doi.org/10.1017/S0953820817000036 Published online by Cambridge University Press

Dougherty argues that this thought can be captured by Only Effects Matter. The problem may then be stated as follows. My argument holds that (i) it is permissible to take a risk *just this once* and (ii) it is impermissible to repeatedly take risks. But Only Effects Matter shows that any account of risk that entails (i) and (ii) is mistaken. As a result, my solution opposes Only Effects Matter in spirit, if not in word. Once we appreciate that the correct way to think about risks is by assessing the permissibility of actions on the grounds of the effects of each action independently, we should reject my argument above.

Moreover, Dougherty claims that Only Effects Matter entails the following iterative premise, the antecedent of which is (1) and the consequent is (2):

> *Iterative Premise*: If you may bring the Chancy Antidote to a victim in Poison once, then you may bring the Chancy Antidote to any sequence of people in Poison.

But, if Only Effects Matter entails Iterative Premise, this provides another argument for the purported inconsistency between Many–Few and Risk Tolerance.

However, appealing to Only Effects Matter is not sufficient for rejecting my argument. Let us begin with the second objection first. Before we can agree with Dougherty that Only Effects Matter entails the iterative premise, we must disambiguate between a narrow-scope and a wide-scope reading of its consequent:[18]

($\text{IP}_{\text{NARROW}}$) If $\blacklozenge Ap_i$, then $(\blacklozenge Ap_1 \,\&\, \blacklozenge Ap_2 \,\&\, \ldots \,\&\, \blacklozenge Ap_n)$, for any i.

(IP_{WIDE}) If $\blacklozenge Ap_i$, then $\blacklozenge(Ap_1 \,\&\, Ap_2 \,\&\, \ldots \,\&\, Ap_n)$, for any i.

It should be clear that Only Effects Matter entails only ($\text{IP}_{\text{NARROW}}$), and not ($\text{IP}_{\text{WIDE}}$). Moreover, the consequent of ($\text{IP}_{\text{NARROW}}$) is my ($2_{\text{NARROW}}$), and I have argued that only (2_{WIDE}), and not (2_{NARROW}), leads to a contradiction with Many–Few. Consequently, one cannot appeal to Only Effects Matter in defence of Dougherty's argument.

Second, we can show that (i) it is permissible to take a risk *just this once* and (ii) it is impermissible to repeatedly take risks do not jointly conflict with Only Effects Matter.

It is easy to show that (i) is compatible with Only Effects Matter, for when we judge that it is permissible to bring the Chancy Antidote in

[18] Thanks to an anonymous reviewer for suggesting that I make explicit the quantificational structure of the disambiguations of Iterative Principle as follows:

($\text{IP}_{\text{NARROW}}'$) $\forall i\,(\blacklozenge Ap_i \rightarrow (\blacklozenge Ap_1 \,\&\, \blacklozenge Ap_2 \,\&\, \ldots \,\&\, \blacklozenge Ap_n))$.

(IP_{WIDE}') $\forall i(\blacklozenge Ap_i \rightarrow \blacklozenge(Ap_1 \,\&\, Ap_2 \,\&\, \ldots \,\&\, Ap_n))$.

https://doi.org/10.1017/S0953820817000036 Published online by Cambridge University Press

Poison once, we are only considering the effects of *that action* on *that person*.

To show that (ii) is also compatible with Only Effects Matter, suppose that one brings the Chancy Antidote to any sequence of people in Poison:

(11) $Ap_1 \,\&\, Ap_2 \,\&\, \ldots \,\&\, Ap_n$.

Then, given the closure of action under conjunction, it follows from (11) that (12):

(12) $A(p_1 \,\&\, p_2 \,\&\, \ldots \,\&\, p_n)$.

When we contemplate the permissibility of (12), it might be tempting to appeal to Only Effects Matter. However, such an appeal is illegitimate, since Only Effects Matter only governs the permissibility of risky actions that affect *one person* and the action in (12) affects more than one person. Consequently, we align our judgements about the permissibility of (12) with (ii), concluding that (12) is impermissible, without fear of Only Effects Matter legislating against this judgement.

Moreover, we can appeal to the following multi-agent version of Only Effects Matter to support the intuitive verdict that (ii):

> *Only Effects Matter 2*: When more than one person is affected by a risky action, the permissibility of the action depends on the possible or actual effects of *that action* for *those people*.

Such a principle supports the intuitive verdict (ii). For if there is sufficiently large risk that an action may harm *some person whomever that might be*, then that action is impermissible. Appealing to Only Effects Matter 2, the multi-agent version of Only Effects Matter, then, provides further support for the compatibility of (i) and (ii). Consequently, Dougherty's objections to the present view of risk-imposition fall short.

Lastly, it is worth considering a version of Only Effects Matter that abstracts away from the permissibility of an action to violations of requirements more generally:[19]

> *Generalized Only Effects Matter*: When only one person is affected by a risky action, whether the agent violates some requirement in virtue of performing this action depends on the possible or actual effects of that action for that person.

[19] Thanks to an anonymous reviewer for suggesting that I consider Generalized Only Effects Matter explicitly.

https://doi.org/10.1017/S0953820817000036 Published online by Cambridge University Press

The picture of risk sketched here is not incompatible with Generalized Only Effects Matter. When assessing whether to perform a risky action that affects only one agent, we need only look at the possible or actual effects of that action for that person. However, when a risky action affects more than one agent, such as the conjunctive actions involved in Poison and Lifeboat, we must broaden our normative horizons and consider the possible or actual effects of that action for all agents involved. Sometimes a conjunctive action composed of a number of independent risky actions makes it near certain that some agent will suffer serious harm, even though it is not certain that any specific agent will suffer serious harm. In such cases, we are not permitted to perform the conjunctive action, even though we are permitted to perform each of its constituent actions.[20]

james.kirkpatrick@philosophy.ox.ac.uk

[20] I should like to thank Joseph Bowen, Karamvir Chadha, Matthew Clark, Daisy Dixon, Tom Dougherty, Jeff McMahan, Jessica Testro, and two anonymous referees for helpful comments and discussions on previous versions of this article. I should also like to thank the AHRC for financial support.

https://doi.org/10.1017/S0953820817000036 Published online by Cambridge University Press