



Espionage and The Harming of Innocents

Lars Christie¹

Accepted: 28 February 2024
© The Author(s) 2024

Abstract

In her latest book *Spying Through a Glass Darkly: The Ethics of Espionage and Counter-Intelligence*, Cécile Fabre suggests that the deception of third parties during an infiltration operation can be justified as a foreseen but unintended side effect. In this essay, I criticize this view. Such deception, I argue, is better justified paternalistically as a means of preventing third parties from becoming wrongful threats. In the second part of the article, I show that Fabre ignores an important moral complication in deception operations where agents intentionally allow others to be harmed as a means of protecting a secret. I argue that intentionally allowing harm to others as a means is a particularly problematic mode of agency which must be addressed in a normative account of espionage.

Keywords Espionage · Just war · Defensive harm · Deception · Means principle

1 Introduction

Espionage, unlike other harmful foreign policies such as war and economic sanctions, is not regulated in international law. In fact, it is hard to imagine how espionage ever can be so regulated. For law to be effective, parties must either trust each other to abide by the law or be able to verify each other's compliance. Given that either alternative is unlikely in the case of espionage, the prospects for future international regulation seem dim. Moreover, the ethics of espionage has not received anywhere near the same attention from contemporary moral or legal philosophers as the morality of war and sanctions. Cécile Fabre's book *Spying Through a Glass Darkly: The*

✉ Lars Christie
Lars.christie@inn.no

¹ Department of Law, Philosophy and International Studies, Inland Norway University of Applied Sciences, Lillehammer, Norway

Ethics of Espionage and Counter-Intelligence is a welcome contribution to fill this normative lacuna.

Fabre notes that a common way of framing the ethics of espionage is by appeal to just war theory. Fabre explicitly rejects this approach, pointing out that espionage harms people in “less obvious and more complex ways” than war does.¹ A normative account of espionage must be sensitive to these complexities and we should therefore “resist the temptation of mechanistically applying the constitutive principles of just war theory”, according to Fabre.² I agree. In this article, I aim to use this sensible counsel to criticize some of Fabre’s own arguments.

The article is divided in two parts. In the first part, I point out that Fabre fails to resist the above-mentioned temptation when she applies a constitutive principle of just war theory, the doctrine of double effect, to analyze a case of infiltration. This leads her to ignore other and more plausible justifications of infiltration operations, or so I argue. In the second part I show that Fabre’s appeal to the distinction between doing and allowing harm obscures important moral complexities involved in certain deception operations where agents intentionally allow others to be harmed.

2 Harming Innocents as Part of an Infiltration Operation

Consider the following case offered by Fabre:

Infiltration

Green’s services have identified a high-level Blue commander who, they have good reason to believe, is preparing an attack on Green’s soil. They encourage Asset to insert himself in the commander’s live-in entourage, including his innocent family members, with a view to spying on him, thereby also gathering information about his family. The commander’s relatives do not contribute to his unjust ends.³

Here is how Fabre explains the case:

Generally, one is sometimes justified in harming innocent civilians in pursuit of a just cause, as a foreseen but unintended side effect of otherwise-just missions (subject to the requirements of necessity, effectiveness, and proportionality.) By implication, Asset is sometimes justified in deceiving the commander’s relatives as a foreseen but unintended side effect of the justified deception of responsible agents, thereby (justifiably) infringing their right not to be deceived.⁴

¹ Cécile Fabre, *Spying Through a Glass Darkly: The Ethics of Espionage and Counter-Intelligence* (Oxford University Press, 2022) 25.

² *ibid.* 25.

³ *ibid.* 210.

⁴ *ibid.* 210.

I believe this is an instance where Fabre falls to the temptation of applying a constitutive principle of just war theory – the doctrine of double effect. But this is a mistake. The relatives are not, contrary to what Fabre’s suggestion implies, causally uninvolved bystanders who foreseeably but regrettably will be harmed, akin to collateral victims in war. The relatives are latent threats. It would both be dangerous for Asset and could threaten the entire operation if the relatives realized his true identity. To prevent this latent threat from materializing, Asset must deceive the relatives. Such harm is not best described as a foreseen side effect which can be justified by appeal to the doctrine of double effect. It is an intended effect of Asset’s behavior.

Once we realize that the relatives are latent threats, we are better able to appreciate the complexities of the case. To begin with, the relatives are merely latent *possible* threats. Even if they found out Asset’s true identity, we cannot be sure that they would threaten Asset by informing the commander.

Consider an example. Suppose that Rebecca is clueless about her relative’s role as commander of a terrorist group. Rebecca may be justified in believing that Asset is up to no good if she finds out Asset is lying and she may then decide to inform her relative (i.e. the commander) about her discovery. In this case, Rebecca would become a threat to Asset. However, there is also a chance that Asset would be able to explain himself in such a way that Rebecca would realize the justification for his mission and refrain from informing her relative about Asset’s true identity.

Harming Rebecca because she might become a threatener is presumably harder to justify than harming someone who is currently posing a threat. But aside from raising the justificatory burden, this difference also affects the complexity of the justification. The justification for harming an actual threatener is grounded in the interest and rights of the victim. The justification for harming Rebecca has an additional paternalistic grounding: it serves Rebecca’s own interest of avoiding a situation where she would be liable to significant defensive harm.⁵

If Rebecca significantly contributes to a threat of wrongful harm, Asset could permissibly impose defensive harm on her. According to the majority view in the philosophy of self-defense, this is so even if the person is epistemically justified in believing that the threat she contributes to is morally justified.⁶ Posing or contributing to a threat of wrongful harm in this way is a form of bad moral luck, because you become liable to defensive harm through no fault of your own. Asset can justify deceiving Rebecca to prevent her from suffering the bad moral luck of becoming an

⁵ Kimberly Brownlee suggests that we also have paternalistic reasons to impose defensive harm on (actual) threateners. Brownlee argues that part of the justification for imposing defensive harm on threateners is grounded in their moral interest in not committing serious wrongdoing. In the case of potential attackers, I have drawn attention to the prudential interest potential threateners have in not becoming liable to defensive harm. These two paternalistic justifications may be combined, as potential threateners may also have a moral interest in not becoming involved in a wrongful threat. See Kimberly Brownlee, “Acting defensively for the sake of our attacker”, *Journal of Moral Philosophy*, 16(2) (2019), 105-130.

⁶ See Jeff McMahan, “The basis of moral liability to defensive killing”, *Philosophical Issues*, 15(1) (2005), 386–405, Jonathan Quong, “Killing in self-defense”, *Ethics*, 119(3) (2009), 507-537 and Helen Frowe, *Defensive Killing: An Essay on War and Self-Defence* (Oxford University Press, 2014).

innocent threatener, thus prevent a situation where Asset would be justified in imposing significant bodily harm on her.⁷

We can imagine Asset explaining this to Rebecca after the operation is over if Rebecca asks for a justification for why she was deliberately deceived. Asset could explain to Rebecca that she, amid confusion and alarm, would likely mistakenly directly or indirectly threaten Asset if she found out he was a spy. Even if Rebecca would not be blameworthy if this happened, Asset would then be justified in imposing significant defensive harm on Rebecca to thwart her threatening actions. Asset would moreover have to explain that infiltrating the commander's family was less harmful than other available strategies for achieving the goal of the mission. If the scenario laid out by Asset is reasonable, and Rebecca realizes this, she has reason to welcome having been intentionally deceived to minimize the risk of becoming a wrongful threat to Asset.⁸

In the discussion so far, I have made several simplifying assumptions. Plausibly, Asset would face additional uncertainties to those mentioned above. In the case where there is a chance that Rebecca would innocently contribute to a threat of harm, there could also be a chance that she would do so culpably. Even if Rebecca is innocent with respect to the commander's unjust enterprise, and even if she understood the true nature of Asset's mission, she may seize on the opportunity to blackmail Asset. Under those circumstances, Rebecca would be culpable for contributing to an unjust threat to Asset. Asset may well face a situation where he must entertain at least three distinct possible outcomes if a relative finds out his true identity: The relative may pose no threat at all, she may pose one innocently, or she may pose one culpably. Additionally, it would be uncertain whether the relative herself would directly threaten Asset, or whether she would threaten him in more indirect ways. I lack the space to explore the complexities these additional uncertainties introduce. These complications do not, however, undermine the point that Asset would have a paternalistic justification for intentionally deceiving Rebecca and the other relatives.

So far, I have assumed that Asset's deception of the relatives merely prevents them from running the risk of becoming threats. But deceiving the relatives could also produce a net benefit to Asset, if their belief in Asset's cover increases the likelihood that the commander will believe Asset's cover. If Asset derives a benefit from the deception of the relatives, a question arises as to whether he is not merely preventing them from becoming threats but also using the relatives as a means.

It is, however, not clear from the mere fact that Asset derives a benefit from deceiving the relatives, that he is *using* the relatives. As long as he has another independent justification for deceiving the relatives (such as the paternalistic justification suggested above), and as long as this is a condition for his deception, the further positive effect this has on Asset's ability to deceive the commander is merely a positive side effect. But insofar as he cultivates a relationship with third parties in order to achieve a benefit above and beyond what is necessary to neutralize their potential as a threat,

⁷ If Rebecca realized Asset's true identity, it might also jeopardize Asset's mission. In that case Asset would have a further justification for harming Rebecca beyond preventing her from threatening Asset and protecting Asset's rights.

⁸ I am grateful to Kasper Lippert-Rasmussen for a helpful discussion on this point.

he uses those third parties as means. The question is whether this admittedly more problematic approach can nevertheless be justifiable.

Discussing a variant of *Infiltration* above, where Asset deceives local third parties in order to get information about legitimate targets in a terrorist organization, Fabre writes: “Asset is not justified in intentionally deceiving the locals if they neither directly contribute to Blue’s wrongful ends nor are under a duty to help Asset thwart those ends.”⁹ This seems to me unduly restrictive. Fabre’s own proposal that the deception of the innocent relatives is a case of side-effect harm accepts that one can justifiably infringe the rights of individuals who are neither liable nor have a duty to carry this harm. The justification for this seems to be a form of lesser-evil reasoning. Presumably the same lesser-evil reasoning can be used to explain why intentionally infringing the right not to be deceived may be permissible, provided enough is at stake.

3 Allowing Harm to Innocents as Part of a Deception Operation

Certain intelligence operations require that we refrain from preventing wrongful harm in order to keep a secret. Fabre offers the famous Enigma case from WW2 as an example of this type of operation.

Enigma

In order to hide from the Germans that their code-breakers had successfully broken the Enigma machine, WWII British authorities had to allow the German Navy to sink a number of Allied ships, at the cost of thousands of lives. By preserving the secret that the code was broken, Britain and its Allies gained a decisive advantage at sea.

Fabre asks us to compare this case with operation Crossbow.

Crossbow

By June 1944, Germany’s long-range rockets (the V1 and V2) had become operational and were used to attack a number of strategic sites, in particular Central London. The Allies misled Germany’s high command into shortening the rockets’ range. Redirecting the rockets this way resulted in central London being relatively untouched though at the expense of South London.

Comparing the two cases leads Fabre to conclude the following: “On the assumption (which I share) that other things being equal, harming is pro tanto harder to justify than allowing harm to happen, Operation Crossbow stood in greater need of justification than the measures taken to protect the Enigma secret.”¹⁰

⁹ Cécile Fabre, *Spying Through a Glass Darkly*, 102.

¹⁰ *ibid.* 110.

I agree with Fabre's view that *Crossbow* was a case where British authorities causally contributed to the harm suffered by the inhabitants of South London, in contrast to *Enigma*, where British authorities allowed harm to be inflicted on Allied ships. But contrary to what Fabre writes, I do not believe *Enigma* stands in lesser need of justification than *Crossbow* merely in virtue of being a case of allowing rather than doing harm.

To see the problematic nature of *Enigma* clearer, it is helpful to compare it with a similar hypothetical case.

Rescue

Allied ships are about to be attacked in two different locations. British authorities have a single aircraft carrier in the immediate surrounding area which can thwart one of the attacks. If it heads to the first location, it will save two thousand lives. If it heads to the second, it will save a higher number of lives.

Assume that the number of lives that would be saved and allowed to die are the same between the *Rescue* and *Enigma*. Some feature of *Enigma* makes the plan adopted by the British authorities harder to justify than the plan adopted in *Rescue*. This feature also makes it clear why *Enigma* is also importantly unlike *Crossbow*. The feature I have in mind may be easier to see by considering a structurally similar case introduced by Warren Quinn:

Guinea Pig

A doctor withholds lifesaving medication from Jim in order to observe the fatal progression of his deadly but painless disease. In this way, she will gain the knowledge necessary to save five other patients with the same disease.¹¹

Quinn points out that *Guinea Pig* contains morally objectionable features which distinguishes it from cases like *Rescue*. In *Guinea Pig*, Jim serves the doctor's goal by becoming a guinea pig. Quinn argues it is especially wrong to involve others in one's plan in such a way that harm foreseeably comes to them, precisely because they are involved in that way. Quinn refers to this as Harmful Direct Agency and contrasts it with Harmful Indirect Agency, where harm comes to a victim without the victim playing a role in the agent's plan (as in the case of *Rescue* and *Crossbow*). Justifying the former mode of agency is harder than the latter. The basic thought behind this idea is that people have a right against being forced to serve a purpose in the plans of others, even if the impersonally best end is achieved by this. As Quinn puts it: "People

¹¹ From Warren S. Quinn, "Actions, intentions, and consequences: The doctrine of double effect", *Philosophy and Public Affairs*, 18(4) (1989), 334-351, 336. I have changed details of the case. In Quinn's original case, he does not stipulate that the disease itself is painless. Quinn's point is that the doctor does not intend for the patient to suffer harm. My version makes this point clearer.

have a strong prima facie right not to be sacrificed in strategic roles over which they have no say.”¹²

The British authorities’ plan to protect the secret that the Enigma code was broken [hereafter the “Enigma secret”] involved *not* taking defensive measures to prevent the German navy from sinking Allied ships. Just like the doctor could easily cure Jim of the disease, the British authorities could easily have engaged in defensive measures against Germany’s attacks, significantly reducing their effectiveness. And just like Jim is being made to serve the doctors’ goal in *Guinea Pig* by not being treated, the sailors are made to serve the goal of the British allies by not being protected. Both are cases where Harmful Direct Agency is involved.

To contrast Jim’s role with other cases where innocent parties are harmed or allowed to suffer harm, Quinn imagines a variation of *Guinea Pig* where Jim has the medicine. In this scenario, the doctor’s plan would only succeed if Jim cooperates in playing the role of guinea pig. We can likewise imagine Allied ships having gained access to the Enigma code by other channels (for instance if other Allied nations had also broken the code and passed it on to them) and had means to significantly reduce casualties themselves by evasive maneuvering and other emergency procedures. The British plan of keeping the code secret would only succeed if the Allied ships cooperated in appearing to be unprepared for the attack. This contrasts with cases such as *Rescue*, where the victims that are not saved make no contribution to the British authorities’ plans. If the first group of Allied ships in *Rescue* were able to engage in defensive measures, it would make no difference to the plan of saving the other group of Allied ships.

On the assumption that the disparity between the harm allowed and the harm averted was sufficiently great in *Enigma* British authorities acted permissibly in carrying out their plan, unlike the doctor in *Guinea Pig*. But this does not alter the fact that the British authorities displayed Harmful Direct Agency in carrying out their plan in *Enigma*. This undermines Fabre’s conclusion that *Enigma*, merely in virtue of being a case of allowing harm, is in lesser need of justification than *Crossbow*.

A critic may point to putative disanalogies between *Guinea Pig* and *Enigma*. Firstly, it may be argued that while Jim has a positive right to medical treatment in *Guinea Pig* (as Quinn in fact assumes) the crew on the Allied ships do not have the same right of rescue in *Enigma*.

There are two ways to respond to this. The first is to reject that there is a disanalogy between the cases with respect to the right of rescue. The second is to argue that the presence or absence of a prior right to rescue is irrelevant to the constraint against Direct Harmful Agency.

How we reject the disanalogy depends on how we ground Jim’s right of rescue in *Guinea Pig*. If we ground the right of rescue in the doctor’s role-based duty to assist those in medical need, we can merely restrict the analogy between *Enigma* and *Guinea Pig* to the sinking of ships with British civilian crew members (the convoys with Allied ships consisted of both merchant and naval ships). Plausibly the British authorities have an equally strong (if not stronger) role-based duty to save British

¹² *ibid.* 350.

citizens in mortal danger when it is within their capacity to do so, as the doctor has vis-à-vis her patients.

Assuming instead that Jim's right of rescue in *Guinea Pig* is an instance of a universal right of rescue at a reasonable cost, then the crew on the Allied ships would likewise have this right.¹³ It may be tempting to restrict a right to rescue to cases where it does not lead to an unreasonable cost on third parties. Since rescuing the ships comes at the cost of losing a decisive advantage at sea later in the war (which would foreseeably lead to a greater loss of life) this may be thought to count as an unreasonable cost in the relevant sense. If this argument is successful in undermining the claim that the British authorities have a duty to rescue the Allied ships, it will also succeed in undermining the argument that the doctor has a duty to treat Jim. The argument would, if effective, deny that it could be right to rescue someone even if that means one will be unable to save a greater number later, denying the clear intuitive verdict in *Guinea Pig*.

One can also deny that victims must have a previous independent right of rescue for Harmful Direct Agency to be morally problematic. Consider the following case, originally offered by Dana Nelkin and Sam Rickless, where the agent has no role-based duty to the patient.

Rachel is about to save the life of a stranger to whom she owes no duty of positive service, but then changes her mind when she sees that the stranger's organs can be used to save the lives of five others.¹⁴

I agree with Nelkin and Rickless that this case does not seem significantly less bad than *Guinea Pig*. The stranger's right not to be pressed into serving Rachel's goals does not seem much different from Jim's right not to be pressed into service of the doctor's goals in *Guinea Pig*.¹⁵

Here is another possible disanalogy between *Guinea Pig* and *Enigma*. In *Guinea Pig* the threat is developing inside Jim's body, in *Enigma* the threat to the sailors is coming from outside. Intentionally allowing the threat to develop inside Jim necessarily entails involving Jim in the plan to save the five. Intentionally allowing a German threat to advance does not similarly entail involving the sailors in the plan to keep the *Enigma* secret. The fact that the sailors most likely will suffer lethal harm may be considered as a mere foreseeable consequence of the British authority's plan.

¹³ In fact, it may put the crew on the Allied ships on stronger moral grounds than Jim insofar as one believes that the duty to protect someone from a harmful rights violation is more demanding than the duty to protect someone from a threat of harm alone. Fabre believes that we have both a general duty of assistance and a narrower duty to protect someone from a rights violation. Fabre does not explicitly contrast these duties, but only with respect to the latter does she hold that wrongfully breaching it makes the agent liable to defensive harm (Spying Through a Glass Darkly, 27-30).

¹⁴ Dana K. Nelkin and Sam C. Rickless, "Three cheers for double effect", *Philosophy and Phenomenological Research*, 89(1) (2014), 125-158. I've slightly changed the case to keep it as similar to *Guinea Pig* as possible. In the original case Rachel sees that the stranger's organs can be used to save two of her friends. I have changed the number of beneficiaries and removed the friendship relation between Rachel and the beneficiaries of the action for the case to resemble *Guinea Pig* more closely.

¹⁵ Jonathan Parry makes the same point with respect to unreasonably costly rescue. See, "The scope of the means principle" in *Journal of Moral Philosophy*, 20(5-6) (2022) 439-460.

But unless this is a foreseen effect of a role the sailors play in the British authorities' plan it does not carry the objectionable features which characterizes the harm Jim suffers.

This objection fails because it relies on an implausibly narrow description of the British authorities' plan. Their plan is not merely to allow the German threat to advance. They need the German threat to advance *as if it has not been discovered*. This appearance is created by intentionally failing to notify the navy ships escorting the convoys and by not sending any military reinforcements. By playing the role of remaining unprotected and unprepared in the face of an imminent German attack, the sailors contribute towards the end of the British authorities. And precisely because they are made to play this role, they are exposed to a significant increase in their risk of being harmed.

Let's consider a final possible disanalogy between *Guinea Pig* and *Enigma*. In *Guinea Pig*, the doctor exploits an opportunity to secure a good (medical knowledge) that only exists because of Jim. The same is not true in *Enigma*. The British authorities do not exploit an opportunity that exists only because of the Allied ships to secure a good (preservation of the Enigma secret). While both Jim and the Allied sailors are intentionally involved in a plan that harms them, Jim's presence is being exploited in a sense that the sailors' presence is not. Allowing harm to a person to secure goods that could not be obtained in that person's absence, it may be argued, is an additional exploitative feature of *Guinea Pig*.¹⁶

It is not clear how the counterfactual scenario should be spelled out to test whether *Enigma* was a case of exploitation. Are we to imagine a counterfactual scenario where the British authorities had the Enigma secret, but where there were no Allied convoys? It does not seem clear that in the absence of Allied convoys the British authorities could have kept the Enigma secret at no cost. If there had been no convoys, the supplies for the Allied war effort would have been transported by some other route which the Germans would presumably have attacked as well. In this scenario, the British authorities would equally have had to intentionally allow a number of German wrongful attacks to be carried out uninterrupted to preserve the Enigma secret.

The conclusion we should draw from this is that it is wrongheaded to suggest that a situation where a benefit is secured must be compared to a counterfactual situation where the victim is absent – *and no one else is there instead* – to determine whether it is a case of exploitation.¹⁷ Firstly there are many cases where there is no plausible counterfactual scenario with only the victims absent and nothing in their place. It seems unduly restrictive to deny that victims could be exploited in such cases, just because we cannot establish that their *presence* is being exploited. Moreover, when

¹⁶ Several philosophers have advanced the idea that harmful agency that aims to derive a benefit from a victim's presence is more wrongful than harmful agency that aims to eliminate a problem posed by a victim's presence (for instance if they are an obstacle). See Warren S. Quinn, "Actions, intentions, and consequences", Jonahntan Quong, "Killing in self-defense" and Helen Frowe, *Defensive Killing: An Essay on War and Self-Defence*.

¹⁷ See S. Matthew Liao and Christian Barry for a similar line of criticism. S. Matthew Liao and Christian Barry, "A critique of some recent victim-centered theories of nonconsequentialism", *Law and Philosophy*, 39(5) (2020), 503-526.

the relevant counterfactual scenario is one where someone else would take the victim's place, this does not make exploitation less wrongful. To see this, consider an amended version of *Guinea Pig*, where we stipulate that if Jim was not there, Jon would be there, suffering from the exact same disease. In this version, the doctor is not taking advantage of Jim's presence to obtain a benefit he would not have in his absence, because in that case Jon would present the doctor with the same opportunity. Since the intuition is that the doctor acts just as wrongly in this case, the explanation for the impermissibility verdict in the original *Guinea Pig* case cannot be that the doctor obtains a benefit from Jim that he would not have in Jim's absence. Rather, it is because the doctor is forcing Jim to serve a purpose in his plan to save the others, where harm comes to Jim precisely because of this role.¹⁸ And as I've argued, this is analogous to the type of agency displayed by the British authorities when they intentionally allowed the Allied ships to be attacked as a means of sustaining the German belief in the intactness of the Enigma code.

4 Conclusion

Fabre has written an important book on a topic to which contemporary philosophers have not given sufficient normative attention. I sympathize with her warning against the mechanical application of doctrines familiar from other areas of philosophy and have pointed to some cases where I believe Fabre could follow her own counsel more consistently. Firstly, I believe Fabre is mistaken in suggesting that we should understand the deception of third parties during infiltration operations as a side effect which can be justified according to the doctrine of double effect. It is better understood, I argue, as a case of intentionally deceiving a third party to prevent him from getting caught up in a complex operation where he might end up posing a wrongful threat.

Secondly, I have argued that Fabre overlooks an important moral complication in covert operations where agents intentionally allow others to be harmed as a means of protecting a secret. In such operations, of which I have argued *Enigma* is an example, agents intentionally involve third parties in their plan in a way which is harmful. In such operations we should not invoke the doctrine of doing and allowing to conclude that they stand in lesser need of justification than operations where agents foreseeably cause harm to the victim.

In no manner does my critical response detract from Fabre's truly impressive achievement. To the contrary, I hope it encourages others to engage with Fabre's book which no doubt will set the agenda for the debate in the ethics of espionage and counter-intelligence for many years to come.

Acknowledgements An early version of this paper was presented at a workshop on Cecile Fabre's book *Spying Through a Glass Darkly: The Ethics of Espionage and Counter-Intelligence* at King's College London, November 2022. Many thanks to Massimo Renzo for organizing the workshop, to the participants for a stimulating discussion and to Cecile Fabre in particular for her detailed response. I have also benefited

¹⁸ For a criticism of the moral relevance of this notion of exploitation, see Ketan Ramakrishnan, "Treating people as Tools", *Philosophy & Public Affairs*, 44(2) (2016), 133-165.

greatly from presenting a version of this paper to the philosophy colloquium at Inland Norway University of Applied Sciences. Finally, I wish to thank Kasper Lippert-Rasmussen for lengthy and helpful discussions on the topics addressed in the paper.

Funding Open access funding provided by Inland Norway University Of Applied Sciences

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.