

The teleological account of proportional surveillance

Frej Klem Thomsen, Senior Researcher, Danish Institute of Human Rights, fret@humanrights.dk

June 2019 draft

Abstract: This article analyses proportionality as a potential element of a theory of morally justified surveillance, and sets out a teleological account. It draws on conceptions in criminal justice ethics and just war theory, defines teleological proportionality in the context of surveillance, and sketches some of the central values likely to go into the consideration. It then explores some of the ways in which deontologists might want to modify the account and illustrates the difficulties of doing so. Having set out the account, however, it considers whether the proportionality condition is necessary to a theory of morally justified surveillance. The article concludes that we need and should apply only a necessity condition, but notes that proportionality considerations may retain some use in practice, as a form of coarse-grained filter applied before assessing necessity when deliberating the permissibility of potential forms of surveillance.

Surveillance is all around us in at least two ways: as practices which citizens of many countries in the 21st century are routinely subjected to on a daily basis, and as a topic of heated public and political debate. Surveillance is also many different things, from the undercover police officer shadowing a suspect, across CCTV cameras monitoring parking lots or public transport, to software logging online-behaviour. Speaking broadly, we can say, I think, that an agent engages in surveillance against a person (or persons) when she performs an observation or series of observations of that person, which are suitable for obtaining information about the person, intended to obtain information about the person, and sufficiently extensive to qualify as surveillance, as opposed to e.g. glancing at a passer-by.¹

¹ Cf. Ross Bellaby: “‘Surveillance’ can cover a wide range of activities from CCTV cameras and ‘covert surveillance’ to dataveillance and datamining. Who the individual ‘is’, where s/he is going, with whom s/he is associating or what s/he is doing all become the concern of the watchful eye.” (Bellaby, 2012, p. 105) And Kevin Macnish: “I use “surveillance” here to mean the monitoring of a competent adult or adults over a period of time without their consent.” (Macnish,

On this conception, surveillance is descriptive. However, surveillance also requires moral analysis. Not only because surveillance involves, or at least can involve, collecting intimate information about persons in a way that many find intuitively troubling, but also because we live in a world in which developing technologies allow surveillance to constantly become more ubiquitous and refined. Thus, a pressing issue in applied ethics is the provision of a complete and credible account of the conditions of morally permissible surveillance.

Given its public and political importance, it is no surprise that there is a growing academic literature on the moral analysis of surveillance, and related phenomena.² However, in order to provide a complete account, we need to answer a whole string of further questions, and as a result it is, I think, no exaggeration to say that much work remains to be done. One important further question is this: given that surveillance will in most cases involve both moral goods and bads, that is, factors that count for and against the permissibility of surveillance, when are the goods of surveillance sufficient to justify the imposition of the bads?³ Lurking beneath this question is another group of sub-questions: what goods and bads are relevant and likely to occur in practice? How does the distribution of goods and bads across persons affect permissibility? Given that some goods and bads are going to be risky, how should probabilities be factored in? And finally, how ought we to weigh and compare goods and bads to ultimately establish which direction the scales tip? Jointly, answers to these questions constitute (most of) a theory of proportionality for surveillance. The ambition of this article is to set out a particular such account of proportional surveillance, illustrate its usefulness, and argue against some alternatives.

2015, p. 530) Also: “Surveillance involves paying close and sustained attention to another person. It is distinct from casual yet focused people-watching, such as might occur at a pavement cafe, to the extent that it is sustained over time. Furthermore the design is not to pay attention to just anyone, but to pay attention to some entity (a person or group) in particular and for a particular reason. Nor does surveillance have to involve watching. It may also involve listening, as when a telephone conversation is bugged, or even smelling, as in the case of dogs trained to discover drugs, or hardware which is able to discover explosives at a distance.” (Macnish) I will not defend my definition here, partly for considerations of space, and partly because I pursue that task in a different article ([self-reference deleted](#)), but primarily because it is unlikely to be controversial in a way that affects the arguments of this article.

² Notable contributions to that literature are: (Austin, 2003; Bellaby, 2012; Doyle, 2009; Hadjimatheou, 2016; Macnish, 2014, 2015, 2016, 2017; McCloskey, 1980; Nissenbaum, 1998; Parent, 1983; Powers, 1996; Rachels, 1975; Rubel, 2007, 2011; Ryberg, 2007; Scanlon, 1975; Solove, 2002, 2006, 2007; Tavani, 1999, 2007; Tavani & Moor, 2001; Thomson, 1975; Von Hirsch, 2000)

³ We might speak somewhat more familiarly of “costs and benefits”. I prefer the terms “moral goods” and “bads” here, so as to reduce the risk of being misunderstood as building upon or arguing for a consequentialist theory of permissible surveillance.

I proceed as follows: section two below discusses how to understand proportionality, drawing on conceptions in criminal justice ethics and just war theory. Section three defines teleological proportionality in the context of surveillance, sketches the values likely to go into the consideration, and argues for a counterfactual baseline. Section four explores some of the ways in which deontologists might want to modify the account and illustrates some of the difficulties of doing so. Specifically, it reviews the idea that the set of goods should be restricted, that proportionality should require a surplus of good, and that there should be both an upper and a lower-bound on the badness of proportional surveillance. Section five then argues in favour of a wider proportionality condition, which expands the set of acts reviewed, and briefly reviews an argument for narrow proportionality. Section six summarises and concludes.

2. Two accounts of moral proportionality

Proportionality is a complex concept. At a very general level, a condition of proportionality in ethics consists of a requirement that the action at stake possess some feature to a degree that is sufficient in comparison with the degree to which the situation possesses some other feature. We can say as an **abstract schema of morally proportional action** that:

An action φ is proportional *iff* φ is P_1 enough (R) for the degree to which the situation is P_2 , where P_1 and P_2 are properties of the action and situation respectively.

In order to apply this to a specific context, i.e. an action-type and situation, we need to specify the two properties P_1 and P_2 and the sufficiency relation R between the properties, that is, what it means for some degree of P_1 -ness to be enough for some degree of P_2 -ness. There are two prominent accounts of proportionality in contemporary moral philosophy, which illustrate how the schema can be filled out.

The first, commonly applied in just war theory, holds (roughly) that an act of war, which inflicts harm upon the enemy, is morally permissible only if the strategic gain is great enough when compared to the harm inflicted.⁴ Thus, even a protagonist who meets the other conditions ordinarily held to be

⁴ Cf. Seth Lazar: “Defensive harm is proportionate when (and only when) the harm inflicted is an appropriate fit for the threat thereby averted.” (Lazar, 2012, p. 5) Thomas Hurka: “Despite their differences, the various proportionality conditions – *ad bellum* and *in bello*, simple and comparative, objective and subjective – all say a war or act in war is wrong if the relevant harm

necessary for actions of war to be morally permissible, such as fighting for a just cause (i.e. self-defence or the defence of innocent others), is not permitted to kill or harm the enemy, if this does nothing or too little to serve that cause. The rationale behind the requirement is simple: causing harm is morally bad, even when the victim of that harm is an enemy, and is permissible only if the badness of doing so is outweighed by thereby achieving a sufficiently great amount of good. In the context of just war, such good ordinarily means the prevention of future harm by bringing the war to a successful end. In summary: theories of just war theory conventionally claim that an act of war is proportional in the sense relevant to moral permissibility *iff* a sufficient balance of good over bad is produced by the act.

The second account of proportionality is commonly evoked in deontological criminal justice ethics and holds (roughly) that criminal sanction of an offender is morally permissible only if the magnitude of the wrong that the offender has done fits the severity of the applied sanction.⁵ As in the case of just war, the condition rests upon the notion that causing harm is morally bad, even when the victim of that harm is a convicted offender, and is therefore permissible only if the badness of doing so is justified by countervailing considerations.⁶ As arguably the most classical example, traditional retributivists hold that this consideration is a desert-based reason to impose suffering on the offender. (e.g. M. S. Moore, 1997) In summary then: deontological theories of criminal justice hold that a harmful criminal sanction is proportional to the crime, *iff* the sanction strikes a (deontically) appropriate balance between the magnitude of the crime in question and the severity of the response.⁷

it will cause is out of proportion to its relevant good.” (Hurka, 2005, p. 38) Jeff McMahan: “[Jus ad bellum] proportionality holds that the resort to war is impermissible if the bad effects of the war would outweigh the good.” (McMahan, 2009, p. 18)

⁵ Cf. Andrew von Hirsch and Andrew Ashworth: “...the principle of proportionality [requires] the severity of the penalty to be proportionate to the seriousness of the criminal conduct of which the defendant has been convicted.” (Von Hirsch & Ashworth, 2005, pp. 131-132) Ashworth: “Cardinal proportionality [...] requires that the penalty should not be out of proportion to the gravity of the crime involved.” (Ashworth, 2005, pp. 84-85) Anthony Duff: “Any normative theory of punishment includes some principle of proportionality, which requires that punishment must be ‘proportionate’ to what justifies it.” (Duff, 2001, p. 132) Von Hirsch: “The primary basis for deciding quanta of punishments, under [contemporary desert theory], is the principle of proportionality or ‘commensurate deserts’, requiring the severity of the penalty to be proportionate to the gravity of the defendant’s criminal conduct.” (von Hirsch, 2009, p. 118) There is a second sense of proportionality at stake in criminal justice ethics, which holds only that punishments must be ranked, so that no crime is punished less severely than an equally or more serious offence, or more severely than an equally or less serious offence. I set aside such ordinal proportionality here.

⁶ The so-called “problem of punishment”. (Boonin, 2008)

⁷ I deliberately avoid use of the phrasing that the crime *deserves* a certain severity of sanction in response. My broader formulation above is intended to capture the whole set of deontological theories that subscribe to a principle of

There are two significant differences between proportionality conceptions in just war theory and criminal justice ethics. The first is that, unlike just war theorists, criminal justice theorists who appeal to a principle of proportionality are typically adamant that this consideration *cannot* be avoiding future harm, since this would entail the potential permissibility of punishing the innocent.⁸ The perspective applied to proportionality in criminal justice ethics is strictly retrospective – it does not compare the badness of the harm inflicted upon the offender with the prospective benefits of deterring potential future offenders, but rather with the badness of the past wrong that the offender committed. The second difference is that at least some criminal justice theorists appeal to both negative and positive proportionality: the punishment must be proportional not merely in the sense of not being excessive when compared to the crime, but also in not being too lenient.

Can we transfer parts or all of either of these two accounts of proportionality to the context of surveillance? Perhaps, but not without adjustments. The justification for meeting the proportionality condition in the context of criminal justice is that the harm inflicted is a response to the sufferer's past wrongdoing. The justification for meeting the proportionality condition in the context of just war is that the harm inflicted is a means to defending innocents against an imminent or ongoing attack. The motivation for the condition in the former case is that those harmed by the agent routinely suffer years of deprivation, hardship, and isolation; in the latter, that they suffer death, mutilation, and serious injury both physical and psychological. As is hopefully obvious, these justifications and motivations do not apply to the standard situations in which we evaluate the permissibility of surveillance. Rather than transplant one of these conceptions of proportionality wholesale, we will need to adapt it to the context of surveillance.

proportionality, not merely the subset of retributivist theories that take the normative reasons at stake to be grounded in desert. For similar reasons I avoid the theoretically loaded term "punishment" in favour of the more inclusive "sanction". Note finally that the formulation is intended to capture the way proportionality in criminal justice ethics may require that the sanction be *neither* more nor less than fits the crime. We will return to that point in the below.

⁸ Just war theory conventionally employs a similar restriction in the shape of "non-combatant immunity" but accepts that there can be cases of permissibly harming non-combatants, subject to the condition of proportionality, if the harm is unintended (although most theorists accept that in extreme conditions even intended harm to innocents may be permissible). The different shape of the condition in criminal justice ethics might therefore be said to emerge because punishment cannot be unintentional, although Victor Tadros has shown how the self-defense justification that underlies much of contemporary just war theory can be applied to criminal justice ethics more broadly. (Tadros, 2011) However, retrospective criminal justice theorists could potentially avail themselves of the considerations employed in just war theory to explain the potential permissibility of harming bystanders to punishment, e.g. the offender's family, who suffer through shaming and deprivation of access to a partner/parent/etc. In that context, prospective proportionality would presumably become pertinent.

3. The teleological account of proportional surveillance

How ought we to transfer the concept of proportionality to the context of surveillance? I suggested above that a proportionality condition must fill out the abstract schema, specifying the properties and sufficiency-relation of the action-type/situation-pair. In the following I first set out a teleological account of proportional surveillance, which fills out the schema by holding that surveillance is proportional when the good (P_1) that surveillance brings about is not exceeded by (R) the bad (P_2). Specifically, I shall say of **teleologically proportional surveillance** that:

an act φ (of surveillance; I shall mostly omit this specification) is proportional in the sense relevant to moral permissibility *iff* the good that φ 'ing brings about is equal to or greater than the bad that φ 'ing brings about.

I focus initially on this account both because it seems to me highly plausible and because setting it out first helpfully lays the groundwork for discussion of potential variations.

It is worth noting that the account's being teleological does not mean that it is or presupposes consequentialism. Teleology is (roughly) the domain of reasons for action grounded in resulting states of moral value; consequentialism (roughly) the ambitious theory that such reasons exhaust the space of moral reasons. It is thus perfectly possible to hold that the teleological account captures the condition of proportionality, while maintaining that there are other conditions of permissible surveillance, and that these are deontological in nature. Nonetheless, I anticipate that some friends of deontology will be critical of a teleological account. To these I offer two inducements to read on. First, teleology plays a part in mainstream deontological ethics, and presumably should do so in some form in a deontological account of proportional surveillance. Second, as mentioned above I will shortly explore three potential deontological modifications. Although I argue against these modifications, they illustrate ways in which one might "deontologize" the account. Thus, I hope that even the sceptical reader might take elements of the analysis to come as a starting point.

Disclaimers and enticement complete, let us fill out the teleological account.

3.1. The goodness and badness of proportional surveillance

The teleological account of proportional surveillance takes as its two properties the good and the bad brought about by an act of surveillance, and so requires that we consider the values at stake. There are two complicating factors for this issue. The first is that we want to say something about

values specific and concrete enough that our reflections can inform actual decision-making; a completely general axiology will not serve this purpose. The second that there are many different forms of surveillance, and what we say must either take into account these differences or be general enough to apply across contexts. An account that fails to do the first is uninformative; one that fails to do the latter is inaccurate. As is evident, these considerations are at odds with each other.

Consider for illustration of the challenge two analyses of the badness of surveillance. Kevin Macnish contends that surveillance can be bad (“cause harm”) in at least 12 ways: “1. Privacy violations. 2. Chilling effect. 3. Social sorting – stereotyping, stigmatization, discrimination. 4. Paternalism (harm to autonomy). 5. Social fatalism. 6. Behavioural uniformity. 7. Imbalance of distribution of costs. 8. Diminution of trust. 9. Vulnerability. 10. Fear of control. 11. Human error and abuse of power. 12. Fear of being “found out” when hiding legitimate information.” (Macnish, 2015, p. 543)

Daniel Solove’s concern is the badness of losing privacy. However, if we plausibly accept that surveillance decreases privacy, then we can adopt his analyses as an at least partial account of the badness of surveillance. Solove suggests that the mere awareness or experience of surveillance can cause harm:

“These harms need not be physical or emotional; they can occur by chilling socially beneficial behavior (for example, free speech and association) or by leading to power imbalances that adversely affect social structure (for example, excessive executive power).” (Solove, 2007, p. 758)

Beyond these, Solove suggests, problems can arise in a number of ways:

Aggregation involves the combination of various pieces of data about a person. *Identification* is linking information to particular individuals. *Insecurity* involves carelessness in protecting stored information from leaks and improper access. *Secondary use* is the use of information collected for one purpose for a different purpose without the data subject’s consent. *Exclusion* concerns the failure to allow the data subject to know about the data that others have about her and participate in its handling and use. [...] *Breach of confidentiality* is breaking a promise to keep a person’s information confidential. *Disclosure* involves the revelation of truthful information about a person that impacts the way others

judge her character. *Exposure* involves revealing another's nudity, grief, or bodily functions. *Increased accessibility* is amplifying the accessibility of information. *Blackmail* is the threat to disclose personal information. *Appropriation* involves the use of the data subject's identity to serve the aims and interests of another. *Distortion* consists of the dissemination of false or misleading information about individuals. [...] *Intrusion* concerns invasive acts that disturb one's tranquility or solitude. *Decisional interference* involves the government's incursion into the data subject's decisions regarding her private affairs."⁹ (Solove, 2006, pp. 490-491)

Both Macnish's and Solove's lists give insightful overviews of many of the ways in which surveillance may be cause for concern but considered as accounts of the badness of surveillance they are problematically *ad hoc*. The connection between the items on the lists is unclear, which leaves uncertain both the extent to which the lists are complete, and whether the suggested items genuinely merit inclusion.¹⁰ What we require is a more general picture of the good and bad of surveillance grounded in intrinsic value.

At a general level we can distinguish broadly between the good and bad of the act of surveillance, the good and bad of the information surveillance provides, and the good and bad of the responses to surveillance.

Consider first how carrying out the act of surveillance affects the agent. An act of surveillance requires time, effort and often material resources on behalf of the agent. Expending these in carrying out surveillance is a bad that attaches to the act.¹¹ Conversely, the surveillance might e.g.

⁹ Solove intends for this list to illustrate his point that there is no distinctive bad of losing privacy, but rather a cluster of problems that emerge in the context of privacy-losses. There are two ways of reading this claim. The first is that losing privacy is not intrinsically bad, but can be instrumentally bad in a number of ways. I believe this is Solove's intended reading, and it seems to me a very plausible claim. The second is the stronger claim that there is nothing that unites the various ways in which losing privacy can be instrumentally bad. This is too strong. The various ways in which losing privacy might be bad can be instrumentally bad for ultimately the same reason, e.g. because they lead to decreases of human wellbeing, as indeed I think is the case.

¹⁰ Macnish argues that the harm of surveillance is generally tied to the degree to which it is intrusive. As a very general observation this sounds plausible, but if the concept of intrusion is to have the kind of explanatory power that would allow it to support an account of the badness of surveillance, then we need a more sophisticated and rigorous understanding of it than the common-sense conception. Unfortunately, despite exploring several examples and dimensions of intrusion, Macnish does not ultimately clarify what this sense of intrusion might be.

¹¹ Does counting this use of resources as a bad of surveillance sneak opportunity costs into the account of proportionality? Not necessarily. If we assume that resources have value independent of the opportunities they provide then this is the value at stake here. If we assume the opposite then we can discount them at this stage, although the

provide training, the opportunity to test equipment or techniques, or gratification to the agent. This is a good that attaches to the act. Something similar applies to other persons, most prominently the target of surveillance. The act can be bad simply by virtue of imposing physical or psychological harm on the target, e.g. because the constant humming from a drone hovering above causes stress. On the other hand, we can also imagine targets who benefit from the surveillance, such as the exhibitionist who enjoys the exposure and attention. Surveillance also decreases privacy and some believe that this is in itself bad - Macnish's concern with privacy violations suggests that he thinks so, as does Solove's concern with intrusion – although this view faces powerful objections. (e.g. Doyle, 2009self-reference deleted; Ryberg, 2007; Thomson, 1975)

As for the further effects of surveillance, the most obvious result of (successful) surveillance is increased information, which in turn allows the agent to make better decisions. Whether this ultimately constitutes a good depends on several factors, including crucially whether the agent *actually* employs the information to improve her decisionmaking, and for *what* the agent employs the information. A police officer who employs improved information to better target criminal offenders may do (more) good because of the information; a corporation that employs improved information about voters to manipulate their information stream so as to deceive them into voting against their own interests may do (more) bad because of the information; an agent who ignores the information in favour of her preconceived opinion will do neither (although she may of course do good or bad for reasons unrelated to the information).¹² However, we must also consider the possibility of increased misinformation as a result of unsuccessful surveillance. That is, the agent may come to believe something that is in fact false as a result of surveillance, and to employ this misinformation in her decisionmaking, which will (all else equal) make her decisions worse.

Finally, it is widely recognized that experienced surveillance can have behavioural effects, that is, people respond to surveillance by acting differently than they otherwise would have. As this effect relies on the experience of surveillance, it does not pertain, of course, to covert surveillance, and depends crucially on how persons *perceive* the surveillance. Whether such responses are good or

fact that they intuitively ought to count as costs strengthens the case for the charge, to which we shall return, that narrow proportionality is insufficient because it ignores opportunity costs.

¹² Macnish's concerns with social sorting, distribution of costs, and human errors or abuse of powers belong in this category, as do Solove's with insecurity, disclosure, exposure, and blackmail.

bad depend on the nature of the response. Deterring potential criminals from offending presumably counts as a good, as might e.g. increased exercise from employing an app that allows users to share data about their training regimens, while deterring legal political activism or strengthening stereotypes through fear of being documented deviating from the norm might be bad.¹³

As is apparent, the goods and bad I discuss above are highly contingent. This is inevitable given that surveillance works differently in different contexts, but it leaves us with the question of how to decide whether something is (generally or actually) instrumentally (dis-)valuable, and hence merits inclusion, or morally irrelevant. Ultimately, this matter turns on whether it affects something of intrinsic value. This is the theoretical grounding that at once explains why we include certain factors as generally instrumentally (dis-)valuable and serves as the condition for whether they are actually good or bad on any given occasion. That is, the good and bad of surveillance, generally speaking, are those factors of the act, information and response that generally are instrumentally valuable or disvaluable since they promote or diminish something of intrinsic value.¹⁴

What one takes the good and bad of surveillance to be will therefore depend on what views one holds about axiology. If we assume welfarism, it follows directly that the goods and bads I have discussed above can at most be instrumentally (dis-)valuable, since none of them involve wellbeing directly. However, it is also clear that many of them can and often will be instrumentally (dis-)valuable, since they can cause lives to go better or worse. This explanation should appeal to non-welfarists too – after all, few deny that well-being is one central component of a complete axiology, and as such the conclusions I have drawn about well-being continue to apply. A value-pluralist axiology might require somewhat more complex views on which states of affairs are instrumentally valuable but will not fundamentally alter the account I have given.

Before we move on to the sufficiency relation, it is worth briefly broadening the perspective on the properties of teleologically proportional surveillance. Specifically, why compare *good* and *bad*? One

¹³ Macnish's concerns with chilling effects, social fatalism, behavioural uniformity, fear of control and loss of trust belong in this category.

¹⁴ Note that in all cases, the good and bad is measured at the margin: an act of surveillance that results in an offender being successfully prosecuted, and a resulting increase in general deterrence, contributes the value of this increase, not, for example, value equal to the average deterrence created by successfully prosecuted cases, or even of successfully prosecuted cases resulting from surveillance. This follows directly from the way the baseline of proportionality is defined, but it is important enough, I think, to merit underlining here.

might focus instead more abstractly on reasons, or more concretely on specific values. As for the first, we might be tempted to say more generally that an act φ is permissible with respect to proportionality *iff* the reasons that count in favour of φ 'ing outweigh and/or cancel the reasons that count against φ 'ing. Formulated at this level, however, proportionality collapses into a general principle of permissibility. The second approach is more promising, and frequently employed. McMahan sums up the conventional view in just war theory when he writes that “[p]roportionality is a constraint on action that causes *harm*”. (McMahan, 2009, p. 19, my emphasis) In the context of surveillance, one might for example take as the relevant properties the degree of intrusiveness and the target of surveillance’s culpability for wrongdoing.¹⁵ I have deliberately left the axiology of my account somewhat open, so as to allow readers with differing views to slot in their favourite version. Thus, it is compatible with my account to adjust certain values according to the liability of the target, as Macnish prefers, e.g. in the shape of desert-weights for the moral value of well-being.¹⁶ I do want to resist the suggestion that proportionality could pertain to a set of properties that dramatically narrows the values at stake. If, for example, we restrict considerations on one side to intrusiveness, this would entail that two cases that vary dramatically with respect to other morally relevant values, such as harm, could be identical with respect to proportionality. This seems to me clearly mistaken. I will discuss a more modest type of restriction as a potential modification below but let us first take a closer look at the sufficiency relation.

3.2 The sufficiency relation and the baseline issue

Now that we have explored the two properties, it is time to do the same for the relation between them. The sufficiency relation of the teleological account of proportionality I have sketched is itself simple value weighing: the good that the act of surveillance brings about must be *equal to or greater than* the bad. This sufficiency relation is strongly intuitive – this is the threshold where the act does no bad on balance, and where there is therefore at least no overall telic reason against it - but not

¹⁵ This approach could be compatible with the account I have set out if one adopted a suitably unorthodox axiology. Such an axiology strikes me as wildly implausible, however. Hence the more charitable interpretation is the approach that takes these values as pertaining to proportionality specifically.

¹⁶ On teleological accounts of desert-adjusted weights for well-being see (Feldman, 1995a, 1995b) One might also consider as starting points David Rodin’s account of 14 factors that affect proportionality in the context of just war, and Susanne Uniacke’s detailed discussion of the effect of liability to harm on proportionality. (Rodin, 2011, pp. 80-81; Uniacke, 2011) For a critical overview of some of the challenges desert-adjusting well-being encounters, including centrally the problem that future decreases in well-being can ground desert-based reasons to do wrong, so as to decrease the moral disvalue of the future decreases, see (Ryberg, 2011)

of course the only relation conceivable. We will explore some more complex possible sufficiency relations shortly. Before we do so, however, the notion of bringing about deserves to be developed in slightly more detail, particularly since elements of this notion will become important once we turn to the differences between narrow and wide proportionality.

What are we to understand by the act bringing about certain effects? To put it differently, how do we determine which good or bad effects exist *because* of an act of surveillance? These problems are not particular to the context of the proportionality of surveillance, of course. There is a large and complex literature on the concept of causality and the attribution of effects to causes, particularly as it pertains to benefitting and causing harm. Although it would take us too far afield to explore the problem in depth here, I will briefly sketch three possible answers and the reasons why we ought to prefer the third.

In order to attribute effects to a particular act of surveillance, we must compare the situation in which the agent carries out surveillance with a baseline. One initial suggestion for the baseline might be the situation prior to surveillance. That would be a mistake. Consider:

Stunting. Capital city is rife with public unrest. Police install comprehensive CCTV in order to deter rioters and looters. Subsequently, crime decreases moderately. However, unbeknownst to authorities, the mood was just beginning to improve when CCTV was introduced, the surveillance caused widespread resentment, and the drop in crime would therefore have been both steeper and deeper, had surveillance not been introduced.

Stunting suggests that the baseline is counterfactual. We are interested in *what would have happened* but for ϕ 'ing.

A second suggestion is to adopt as the baseline the scenario where we, so to speak, "remove the agent". This too would be a mistake, because the mere presence or absence of the agent can have effects, which intuitively we do not want to attribute to the act. In *Stunting* we can readily imagine that the disappearance of the authorities would have a series of effects, such as the removal of a general deterrent effect, but that situation is not plausibly the proper comparison.

A third and more acceptable suggestion is that the baseline is the situation in which the agent acts so as to bring it about that there is no surveillance. In many situations this will involve what we often label “not acting”, i.e. not initiating a form of surveillance. Strictly speaking, what we tend to label “inaction” is a course of action (say, “twiddling one’s thumbs while staring into empty space”), or a set of actions, defined by the way in which they have no relevant causal effects. However, it is easier here to stick to the common way in which we refer to such actions as inaction. Perhaps the most common baseline for proportionality, then is something along the lines of the scenario in which the agent through inaction has no relevant effect on the outcome. However, we might also want to apply the proportionality condition to those situations in which the agent can bring it about that there is no surveillance by interrupting a form of surveillance that would occur but for the agent’s intercession.¹⁷ It will often seem odd, for instance, to say that a form of surveillance is impermissible because disproportional for an agent who initiates it, but not for the agent who takes over to supervise the surveillance after it has begun, and refrains from stopping it. Hence, although I shall mostly be concerned with the more standard examples of the agent bringing about surveillance by initiating it, we shall stick to the more inclusive definition that the counterfactual baseline is the situation in which the agent acts so as to bring it about that there is no surveillance.¹⁸

Note that this baseline’s restriction of comparison to only the action pair [φ ’ing/ $\neg\varphi$ ’ing] means that teleological proportionality does *not* take into account opportunity costs. That is, the proportionality considerations as set out here do not include the cost of failing to pursue an alternative course of action that would result in less bad, more good or both. In the literature on just war theory this is sometimes used to distinguish the proportionality condition from a necessity condition, or narrow proportionality from wide proportionality. As previously mentioned, I will return to this point in the last section of this article to argue in favour of a wide principle of proportionality.

¹⁷ There are further difficulties with this distinction, familiar from the debate on the differences, moral and conceptual, between doings and allowings (Woollard, 2016 provides an excellent overview). These are difficulties, however, that might be said to affect the concept of permissible action more so than the conception of proportionality I here set out.

¹⁸ Some might be inclined to accept the idea of a counterfactual baseline along the lines above, but to object that it makes a difference to permissibility whether the act of surveillance is a “doing” or an “allowing”. This is the central claim of the doing-allowing doctrine (or acts-omissions doctrine). In the context of surveillance, it could function either as a deontological modification of the teleological account of proportionality, or as an independent constraint on permissible surveillance (alongside e.g. necessity, appropriate intentions, justified cause, etc.). The doctrine faces a long series of well-known objections, which seem to me very persuasive. (see e.g. McMahan, 1993; Norcross, 2003; Rachels 1975; although cf. Woollard, 2015) I will not here assume that the doctrine is mistaken, but I shall leave the task of developing an account of proportional surveillance that applies the doctrine to more sympathetic thinkers.

Consider one last complication for the sufficiency relation. The conception of teleological proportionality I have set out above is objective and factualist. It assesses proportionality independently of the beliefs and intentions of the agent, and on the basis of what actually does happen or would have happened, rather than the probabilities of various outcomes occurring. But it is possible to take the opposite view on either or both of these issues. That is, proportional surveillance may be held to pertain to what the agent believes, or could reasonably believe, to be the good and bad that φ 'ing will cause, or to the probability-adjusted good and bad that φ 'ing will cause, rather than to the good and bad that φ 'ing will *actually* cause.

Should we be subjectivists or probabilists about proportionality? One reason that may seem to count in favour of subjectivism and/or probabilism is that surveillance in practice is typically risky. As Macnish points out, a police investigation, for example, which subjects a suspect to surveillance, will rarely have justified absolute confidence that the surveillance will provide evidence that the suspect is engaged in criminal activity. (Macnish, 2014) The effects of surveillance, which determine proportionality, will differ depending on whether or not it turns out to be the case that the suspect is engaged in criminal activity, as well as on whether the surveillance uncovers useful information, whether that information allows authorities to proceed with a beneficial course of action, such as prosecution or coercing the suspect to turn informer, etc. Similar considerations apply to other types of surveillance. Given this riskiness, should we not judge proportionality according to the probable outcomes the agent faces, or her reasonable perception of such probable outcomes? After all, we otherwise risk condemning the actions of an agent as disproportionate, simply because the slimmest of probabilities ends up going against her.

A standard response is that the apparent appeal of probabilism and/or subjectivism can be explained by virtue of their relevance in decision-making and the moral assessment of decisions rather than by their relevance to a condition of permissibility. Consider as a contextualised illustration:

Suspicious. Sneaky is acting very strangely, such that a reasonable and well-informed observer would judge the probability that she is engaged in serious mischief very high. If she was in fact engaged in serious mischief, then a moderately intrusive form of surveillance would reveal it and be proportional.

However, she is not. Her suspicious behaviour is merely the result of a series of coincidences and her slightly neurotic character.

Is the moderate surveillance of Sneaky proportional or not? I want to say that, although a reasonable and well-informed observer will judge that it is proportional, it really is not. The fact remains that Sneaky is not *actually* up to any mischief, and that moderate surveillance will therefore be disproportionate. The reason we may be tempted to say otherwise is the sympathy we feel, and ought to feel, for the reasonable and well-informed observer who engages in surveillance of Sneaky regardless. Surely it would be mistaken to chastise that agent? Yes, but we need not say that surveillance was proportional to explain why this is the case. We can say simply that her action was innocent, in the sense that it was reasonable for the agent to think it justified, and that therefore her wrongful surveillance merits no blame.

As mentioned, this is a familiar response in the wider debate on wrongdoing and blameworthiness, associated prominently with Tim Scanlon's work. (Scanlon, 2008) I cannot hope here to persuade those familiar with the debate and committed to subjectivism or probabilism – the issue seems to me one about which informed thinkers can reasonably disagree – the ambition has been more modestly to render my objectivist and factualist approach plausible by showing how one can respond to perhaps the most obvious challenge. Ultimately, however, I do not think the remaining arguments of this article hinge on taking a particular approach, so I invite readers who are committed to subjectivism or probabilism to assume that we are dealing with their preferred version of teleological proportionality of surveillance.

On that ecumenical note, let us turn from the teleological account to some potential modifications that may appeal to friends of deontology.

4. Deontological variations

The teleological account of proportional surveillance I have set out above holds that an act of surveillance is morally permissible with respect to proportionality *iff* the act brings about more good than bad through its direct effects, the results produced and the responses, as measured against a counterfactual baseline of doing nothing. As I initially emphasised, the teleological account is compatible with both consequentialist and deontological overall accounts of permissible surveillance. There may be some, however, who prefer not merely to combine a teleological account

of proportionality with deontological supplementary conditions, but to construe proportionality itself in ways that are incompatible with classical consequentialism.¹⁹

The most detailed and sophisticated non-teleological treatment of proportional surveillance in the philosophical literature is Kevin Macnish's "An Eye for an Eye". (Macnish, 2015) Hence, it is largely to his treatment of the topic that I will look in this discussion of modifications of the teleological account of proportional surveillance. As I will devote most of my energies to demonstrating the problems I believe such modifications face, it is only fair to note here that I believe these problems pertain to deontological accounts in general more so than to the specific version that Macnish presents.²⁰ In this section I will review three ideas: i) that P_1 should restrict the range of goods, ii) that the sufficiency-relation of surveillance should require some degree of excess good over bad, and iii) that the sufficiency-relation ought to be able to condemn not merely too much surveillance but also too little.

4.1 Restricting the good of proportional surveillance

The teleological account allows a certain flexibility in that it does not require a specific account of the good (or bad); any "non-crazy conception of moral value", as Richard Arneson has put it, will do. (Arneson, 2009) It does, however, count all goods and bads equally. The first suggestion for a deontological modification I wish to discuss is the idea that we should restrict the range of goods relevant to proportionality. Macnish, following Thomas Hurka, distinguishes between "peripheral benefits", which do not provide justification for an action at all, "contributory benefits", which cannot (ever) by themselves make an action proportionate, but which nonetheless affect actions that are otherwise justified, and "sufficient benefits", which can justify an action with respect to proportionality. (Hurka, 2005; Macnish, 2015, pp. 534-536) Note that the idea here is therefore not to challenge a particular account of the good, but rather to say that some of the good effects brought about by surveillance are peripheral, and therefore do not make a difference to whether an act of surveillance is proportional regardless of their being good. If true this would have dramatic

¹⁹ I say classical consequentialism because almost any theory can be interpreted as a form of consequentialism, or "consequentialised", if we adopt a suitably complex axiology. (Brown, 2011)

²⁰ Indeed, Macnish deserves credit both for introducing the idea of applying concepts developed in just war theory to the context of surveillance, and for, in doing so, substantially advancing the debate on the ethics of surveillance in terms of theoretical sophistication. Though his work provides a convenient foil on which I focus my criticism, I do not mean thereby to detract from that achievement.

effects, radically altering the permissibility of many instances of surveillance as compared with the teleological account I have set out.

There are, however, several reasons why I am sceptical about the notion that some benefits are merely peripheral. The first is that the distinction is suspiciously *ad hoc*. Macnish admits that he can find no principled way to distinguish between the three types of benefit but suggests that it is obvious in at least some cases to which of the three categories a given benefit belongs. My worry with this suggestion is whether the intuitions Macnish employs to sort different benefits into the three categories are sufficiently clear and robust to provide the required support for the distinction in the first place. Perhaps it does sound odd to say that allegedly peripheral benefits of economic development or the professional fulfilment of career soldiers (to take two of Macnish's examples) can count as justifications for waging war, but is that because they give *no* reason to do so, or because that reason is vastly inadequate in the realistic scenarios we tend to imagine? Macnish needs to claim the former, but it is not clear on what grounds he can rule out the latter, since the two seem to provide roughly equally good explanations of our intuitive hesitation.

A second argument against these goods being peripheral is that they appear capable of functioning as tie-breakers. Consider:

Choice. Country A can justifiably wage war for humanitarian reasons against country B or country C, but it is impossible for A to do both. Waging war against B will leave A to languish in its current state of immiseration, whereas waging war against C will open trade routes, dramatically improving A's economic development to the benefit of its people, without making anyone else worse off.

Is it odd to say that the economic benefits give A at least *a reason* to prefer waging war against C rather than B? It does not seem to me to be so, and yet, this is incompatible with describing the benefit as peripheral.

A third reason to be sceptical follows from the fact that the same tri-partite distinction does not apply, according to Macnish, to bad effects. (Macnish, 2015, p. 537) This discrepancy between good and bad effects introduces very serious problems when combined with the counterfactual baseline I have argued for above, where it is not clear that we can meaningfully distinguish at least some effects in this way. Consider:

Flip-flop: Agent A ponders the permissibility of an act of surveillance ϕ . If she ϕ 's, it will bring it about that person B lives. Leaving aside this outcome ϕ 'ing is just barely disproportional. If B's survival is a good, then it is peripheral.

Is ϕ 'ing proportional? It might appear that we should deny this – after all, bringing it about that B lives rather than dies appears to be a good, in which case it is a peripheral good, and so does not change the fact that ϕ is disproportional. However, the bringing it about that B dies seems equally clearly to be a bad, and therefore non-peripheral. This leads to a strange form of instability. If the agent ϕ s then the outcome for B as measured against $\neg\phi$ 'ing grounds no reasons for action, but if the agent $\neg\phi$ s then the outcome for B as measured against ϕ 'ing grounds reasons for action. Given that the outcome itself is only determined against comparison with the other case, how can it generate reasons against one action but not for the other? A mere shift in evaluative focus should not, it seems to me, be capable of affecting reasons for action in this way.²¹ A supporting reason is that this could in some cases implausibly leave agents vulnerable to unavoidable wrongdoing. Suppose that the balance of reasons is such that $\neg\phi$ 'ing is just barely permissible disregarding the outcome for B. If B's death counts as a bad, it seems reasonable to suppose that taking the outcome for B into consideration could swing the balance of reasons so as to make $\neg\phi$ 'ing impermissible. In this case $\neg\phi$ 'ing is impermissible as an alternative to ϕ 'ing, *because* B's death counts as a bad. But if the agent now considers ϕ 'ing instead, she will face the fact that ϕ 'ing is impermissible by virtue of being disproportional, *because* B's survival does not count as a good. This implication strikes me as absurd.²²

Consider finally Macnish's argument against including all good effects:

“The alternative, consequentialist position would be to count *all* the benefits of surveillance in the balance. This would, for example, include the enjoyment which the surveillant received from the surveillance. It would then be feasible, given

²¹ My argument here is structurally somewhat similar to, and to a large extent inspired by, the predictable regret objection against relativized pure-time discounting. (Broome, 2005; 2012, p. 151; Mintz-Woo, 2019, p. 412)

²² Note that the distinction at stake here is different from and unrelated to the doing-allowing distinction, and that the objections at stake are not at odds with the doing-allowing doctrine. This follows directly from the fact that my account means to include the possibility of ϕ 'ing being an act that allows an already initiated form of surveillance to continue. But it is also worth noting that the idea at stake here is that certain goods brought about by an act of surveillance should not affect the proportionality of that act, whereas the doing-allowing doctrine holds (roughly) that the badness resulting from an allowing grounds weaker reasons than the badness resulting from a doing.

enough pleasure on the part of the surveillant, that his voyeuristic pleasure could outweigh the harms visited upon the surveilled subject by the surveillance. This would render the surveillance proportionate and, in those terms at least, permissible. Yet both this consideration and this conclusion are, to me at least, ethically unacceptable.” (Macnish, 2015, p. 536)

This is a variant on a standard objection to utilitarianism, and as such a first, minor quibble is that it is inaccurate to label it an objection to consequentialism – it is not the consequentialist principle of holding an action right *iff* it has consequences at least as good as any alternative action that appears to cause problems here, but the welfarist inclusion of intuitively unappealing voyeuristic pleasure as a good. The more serious challenge follows immediately upon the heels of this observation, for as has been recognized at least since G.E. Moore’s discussion of the issue in “Principia Ethica”, the objection is easily answered by adopting a more restrictive axiology, where ethically dubious forms of pleasure, such as voyeuristic satisfaction, have no value (or even negative value).²³ (G. E. Moore, 2000) The Moorean response illustrates that the problem Macnish identifies cannot uniquely recommend the solution of restricting the range of goods relevant to proportionality, but even interpreted as an argument against welfarism it is not clear that Macnish’s objection is persuasive. Consider:

Estranged. Grey is a failure as a parent, who has alienated her children, and leads a lonely, sad life. She generally respects the wishes of her children to be left alone, but once a month she goes to the park and surreptitiously watches her children and grandchildren playing for an hour or two on the park playground. Doing so is the one highlight in an otherwise bleak existence and gives her enormous joy.

Is Grey’s surveillance morally permissible? Perhaps not, but Macnish needs to claim something far stronger – not merely that her doing so is wrong, but that the value she derives from observing her

²³ Macnish makes a related point, that Hurka wants to avoid the situation where a small benefit to a sufficiently great number of people can justify imposing a great cost on a few. (Macnish, 2015, p. 234) However, holding some benefits to be peripheral will do little to avoid this situation, if, as appears to be the case, a benefit’s being peripheral depends on the *nature* of the benefit, since the situation emerges due to aggregation of benefits of different *size*. Furthermore, this is the familiar “one million headaches”-objection to consequentialism, and is subject to the standard responses to that objection. (e.g. Norcross, 1997)

family makes *no* moral difference to the situation. This seems to me at the very least intuitively doubtful.²⁴

4.2 Surplus-good sufficiency

My teleological account of proportionality takes the sufficiency relation to be the equal to or greater of good over bad. One can vary either component of this relation, by setting the threshold for proportionality at another spot, or by more radically changing the direction of comparison. I will review both suggestions, beginning in this section with the idea that we should set the threshold so that the good must be not merely equal to, and not merely greater (by any amount), but *sufficiently* greater by some amount, to justify surveillance with respect to proportionality. Call this the surplus-good version of the sufficiency-relation for proportionality, or SG-proportionality.

In his argument for this claim, Macnish relies on the existence of epistemic and ontological vagueness along the boundary between proportionate and disproportionate actions. (Macnish, 2015, pp. 533-534) This vagueness means that proportionality cannot employ the threshold of “equal to”, since, in Macnish’s words: “[vagueness] means that any call for strict equality between benefits and harms is unrealistic. This would require the precise measurements that I have argued are available for mathematicians but unavailable for philosophers.” The question, Macnish contends, is whether we ought to treat cases falling within the zone of vagueness as proportionate or disproportionate. His answer is that we ought to favour a presumption of disproportionality since “there are very real harms which are apparent and which the ethical framework is intended to guard against”. (Macnish, 2015, p. 539)

The argument for SG-proportionality is:

- 1) There is a zone of irreducible vagueness around the boundary between acts of surveillance where the good of the act equals or exceeds the bad and acts of surveillance where the bad exceeds the good. (vagueness-claim)

²⁴ Does the fact that the voyeurism occurs in public invalidate *Estranged*? It seems to me strange to say, since it is widely held that we do have claims to privacy in public. (cf. Ryberg, 2007)

- 2) If it is uncertain whether the good of an act of surveillance equals or exceeds the bad, then we should hold the act to be disproportional.²⁵ (presumption against bad)
- C) So we should only hold those acts to be proportional, where the good of the act exceeds the bad sufficiently for the act to be out of the zone of vagueness.

The idea at stake here bears a certain resemblance to precautionary principles in environmental and medical ethics but is also different in important respects. Moral precautionary principles generally claim (roughly) that we have strong or decisive reason to not perform acts, when some possible outcomes of the act are very bad and uncertain (i.e. we cannot assign probabilities to them), and the good of the act is comparatively small.²⁶ Three prominent differences between the SG-proportionality condition and precautionary principles merit mentioning: i) uncertainty in the former attaches to the balance of good over bad for cases that are relatively closely balanced, and uncertainty in the former to particular, very strong bads; ii) precautionary principles only apply to cases with very strong bads and comparatively small goods, whereas SG-proportionality only applies to cases that are relatively closely balanced; and iii) precautionary principles are grounded in concern for the possibility of realizing uncertain but very bad outcomes, SG-proportionality in concern for doing overall more bad than good.²⁷

We should not assume at the outset, that these differences need constitute an overall disadvantage for SG-proportionality. After all, precautionary principles face a slew of challenges, which the

²⁵ I take it to be necessary but uncontroversial that the acts of surveillance become uncertain by being within the zone of vagueness. If we could assign probabilities to the act bringing about good equal to or greater than bad, then the presumption against bad does not apply, and a probabilistic variant of the presumption is implausibly strong, but then such probabilities are conventionally unassignable if there is irreducible vagueness.

²⁶ As is now widely recognized, the literature is best understood as containing a family of precautionary principles, rather than *a* precautionary principle. Apart from its moral interpretation, which seems to me most pertinent here, we should also distinguish between decision-making, epistemic and pragmatic variants, which differ from the version I set out above in that they take the principle to be a decision-rule, a belief-forming principle, or a rule-of-thumb respectively. See (Gardiner, 2006; Hartzell-Nichols, 2013; Sandin & Peterson, 2019) It does not seem to me, however, that my arguments in the following hinge on which precautionary principle is at stake. As such, we can set the differences aside.

²⁷ Could we formulate a version of the sufficiency-relation that, unlike SG-proportionality, is a proper member of the family of precautionary principles? It seems we could. However, it would restrict proportionality constraints to the narrow set of cases that feature uncertain and very strong bads. We presumably want proportionality to apply much more broadly. Friends of precaution are likely better off by assuming that it is a separate and supplementary constraint on permissible action.

proponent of different principles can hope to avoid.²⁸ It does mean, however, that SG-proportionality cannot draw credibility from membership of the family of precautionary principles. It must stand or fall on its own.

As it turns out, there are, I think, at least three challenges for SG-proportionality. First, it is not clear why there need be irreducible vagueness. Macnish mentions as possibilities both epistemic and ontological vagueness. Ontological vagueness appears to be the more serious of the two threats, and yet, after initially mentioning it, Macnish does not explain why any such vagueness applies to the goods and bads of surveillance, nor does he argue that it will generate a zone of vagueness. The latter is particularly important. SG-proportionality is a universal condition. This means that there must be a zone of vagueness in all cases of surveillance. For there to be a zone of vagueness in all cases, it needs to be true that all goods and bads are ontologically vague, not merely some of them. If there were discrete goods and bads, after all, then there would presumably be cases where the discrete goods and bad would be sufficiently unbalanced to shift the zone of vagueness above or beyond the threshold of equal to or greater than. If, on the other hand, all goods and bads are vague, it is not clear how we can do proportionality assessments at all. How do we compare goods and bads, all of which are ontologically vague? Those who wish to defend SG-proportionality on the basis of ontological vagueness must develop an account that avoids these issues.

Epistemic vagueness is perhaps more plausible. After all, Macnish is right that measuring goods in practice cannot be done with any degree of precision. We can note in passing, that such vagueness is only a problem for subjectivist accounts of proportionality – on the factualist account I defended above an act remains proportional or disproportional for all that agents might encounter epistemic vagueness in their attempts to assess it – however, even for subjectivists it is not clear that the difficulties Macnish refers to are enough to support genuine vagueness. After all, even if we cannot achieve precision it might well be that we could assign probabilities to different values of each good and bad, in which case it would not be uncertain but probable whether the good is equal to or greater than the bad.

²⁸ Very broadly, the challenge for precautionary principles is to formulate a version of the principle that is at once strong enough to be non-trivial, and weak enough to not have (very) implausible implications by condemning vast swathes of intuitively permissible human behaviour. See (Gardiner, 2006; Harris & Holm, 2002; Peterson, 2006)

Setting aside the difficulties of the vagueness-claim, the second challenge is that the argument for the presumption against bads strikes me as unpersuasive. Why give priority to avoiding a balance of bads over goods? The bads that surveillance brings about are, of course, very real, but then so are the goods. It does not seem to be a sensible rule in general to always avoid a potential balance of bads over goods – that would be a very strong form of risk aversion – and it is not clear why we should think that it is so in this case.²⁹

Third, SG-proportionality applies to all situations that are fairly evenly balanced and where there is for that reason vagueness around the exact balance, but this can require us to hold as disproportional cases that certainly bring about no bads. Consider:

Harmless. Subject consents to wearing a monitoring device that may provide moderately useful health data. The probability of its providing this data and the exact size of the resulting health benefits are uncertain. The device, however, is carefully designed and implemented so as to bring about no moral bads.

I have deliberately left the case somewhat abstract, because it is my experience that intuitions about what constitutes a bad from surveillance differ widely. I hope the reader can fill out such details, if necessary, so as to imagine a case where we can know in advance that surveillance will not bring about any bad. Such cases may exist only in theory. This should not deter us. The problem arises for SG-proportionality as a theoretical account even if it is unlikely or even impossible for it to occur in practice. Now consider what SG-proportionality must say about *Harmless*: given that the good harmless brings about is uncertain, it seems clear that the case falls within the zone of vagueness. That is, insufficient good will be brought about by wearing the device to meet the threshold of a certain surplus of good over bad – it might for example bring about *neither* good nor bad at all. As such, the surveillance is disproportional according to SG-proportionality. This is strongly counter-intuitive.

4.2 Upper and lower boundaries for proportionality

The sufficiency-relation of the teleological account of proportional surveillance I have set out is bounded in one particular dimension. It says that an act of surveillance must bring about *at least*

²⁹ This is another juncture in the argument, at which the doing-allowing doctrine might be thought to be relevant. I think that it is not, whatever one thinks of the doctrine's merits, for reasons parallel to those I mentioned in the context of restricting the goods above (see footnote 22).

this much good to be proportionate, where “this much” is determined by the amount of bad the act of surveillance brings about. We can say that this proportionality imposes an upper boundary on the badness of surveillance, and in practice thereby on the amount of surveillance permissible.³⁰

We can imagine, however, proportionality that instead or also imposes a lower boundary on the badness of surveillance. One could hold, for example, that proportional surveillance must be at least *this* extensive, or *this* bad. Indeed, as we observed in passing at the beginning of this article, some criminal justice theorists, prominently positive retributivists, believe there are deontic reasons why a punishment must be neither too severe *nor* too lenient. Macnish suggests that this can be accounted for by saying that punishment must be proportional in the sense that the harm it inflicts is neither greater nor less than but exactly equal to what the offender deserves, and that something similar may apply to proportionality in surveillance. (Macnish, 2015, p. 532)

I have three concerns with this suggestion. The first is that it is not clear how the notion of a lower boundary should translate to the context of surveillance. In part this is because it is unclear what a lower bound ought to apply to – is surveillance required to be sufficiently extensive, or sufficiently bad? The two may correlate but need not do so. More importantly it is also unclear how or where to set the threshold – what is supposed to determine the “this” in “at least *this* extensive/bad”? Note that the answer cannot be pragmatic, in the shape for instance of the very reasonable consideration that we can have reasons to impose at least a certain amount of surveillance since this surveillance provides goods. Such considerations ground only *pro tanto* reasons, not the hard limit of a proportionality condition.³¹

Second, it seems to me that there are theoretically stronger conceptions of positive retributivism than one which imposes a double-bounded proportionality condition. Positive retributivists can hold that we have strong *pro tanto* reasons grounded in desert to impose punishment of a certain severity, and a proportionality constraint (also, perhaps, grounded in desert) to impose no more than this. This, it seems to me, is a more defensible position because it does not leave retributivists

³⁰ Assuming, plausibly I think, diminishing marginal goods from surveillance, and that at least some marginal bads are constant or even increasing.

³¹ Unless, that is, one is a consequentialist, in which case the balance of telic reasons determine a uniquely permissible action set, in a way that is structurally similar to imposing an upper and lower boundary. But this is a rather different story, which includes opportunity costs (i.e. evaluates a wider set of potential actions), applies to permissibility *tout court* rather than just proportionality, and is in any case not what Macnish and deontologists more broadly might want.

with the embarrassment of having to defend the claims that we *inevitably* do wrong when we punish more leniently than the offender deserves, and that there are *equally* strong reasons not to over- and underpunish.³² (cf. Duus-Otterström, 2013)

Third, and I think most tellingly, a lower bound on proportional surveillance would have very strange implications, crucially that for any relevant lower bound agents routinely do wrong by subjecting some ordinary persons to insufficient surveillance. Consider the following argument, which we might call **the lower-bound dilemma**:

1. Ordinary persons regularly face very different levels of surveillance but are roughly similar in the respects that might determine the level of proportionate surveillance.
 2. If proportionality imposes both an upper and a lower bound on surveillance, then either a) some surveillance of ordinary persons is wrong because they are subject to not enough or not bad enough surveillance, or b) the bound is so low that every ordinary person is subject to either exactly proportional or excessive surveillance.
 3. Surveillance of ordinary persons is not regularly wrong because some are subject to not enough or not bad enough surveillance.
- C) Therefore, there either is no lower bound on proportional surveillance, or the bound is so low that every ordinary persons is subject to either exactly proportional or excessive surveillance.

The proponent of a lower-bound might choose the later of the conclusion's two horns, but this is genuinely a horn of a dilemma, since at least some ordinary persons are subject to no or almost no surveillance. If the lower bound is effectively zero, then it is trivial.

5. Narrow or wide proportionality?

So far I have presented a teleological account of proportional surveillance, and surveyed some potential deontological modifications of that account. I have argued against the modifications, so

³² Notoriously, Immanuel Kant defended roughly this view, when he claimed that even the members of a society that had decided voluntarily to disband were under a stringent duty to first impose full punishment on convicted offenders. (Kant, 1996) Contemporary retributivists often explicitly distance themselves from Kant on this point.

as it stands, the teleological account still seems to me the best account of proportional surveillance. There is, however, a last complication that we need to consider pertaining to the difference between what is sometimes labelled narrow and wide proportionality. As I earlier sketched the distinction, the difference pertains to the range of actions considered as alternatives to the action at stake, and as a result thereof whether proportionality weighing includes opportunity costs. I have set out the teleological account of surveillance as a narrow condition of proportionality, which compares the action at stake only with a baseline of the agent performing no causally relevant action, but I want to argue now that we should prefer a wide conception of teleologically proportional surveillance.

The basis of the argument against narrow proportionality is that meeting the proportionality condition is not in itself sufficient to justify surveillance. This is generally accepted for proportionality conditions in moral theory, but an example can illustrate how it applies to the context of surveillance. Consider:

Serial killer. Police are attempting to capture a murderer on a killing spree in the city. They petition the court for permission to covertly install video- and audio-recording devices in all of the city's public spaces, in order to catch the killer before she strikes again. The court rejects the petition on the grounds that investigation has already narrowed down the list of people who could be the killer to a handful of suspects, and that police can identify the killer simply by shadowing these persons.

Suppose, as I think *could* be the case, that the comprehensive surveillance requested by the police is proportional, because the importance of preventing further murders justifies even such dramatic surveillance measures. Even so, the court's rejection seems to be clearly right because an alternative means of achieving this objective is available, and that alternative brings about (we can suppose) much smaller bads. Although proportional, the requested surveillance is *unnecessarily costly*, and this is enough to make it morally unjustified.³³

³³ This point may be so blindingly obvious that the example strikes the reader as unrealistic – surely the police would pursue the superior course of action rather than waste time petitioning the court for permission to conduct more widespread but no more effective surveillance? In response I readily accept that this is the case. The example is meant only as illustration, and I do not mean to disparage the judgment of real-life police officers. However, in a philosophical literature full of aliens, teleportation devices and bizarre trolley-accidents, I think begging the reader's indulgence in suspending disbelief on this point is reasonably modest.

A conventional explanation of our intuitions in cases like the above is that permissible action must satisfy not merely proportionality, but also a condition of necessity or last resort. Consider the conception of necessity in self-defense developed in an influential article by Seth Lazar:

“Defensive harm H is necessary to avert unjustified threat T if and only if a reasonable agent with access to the evidence available to Defender would judge that there is no less harmful alternative, such that the marginal risk of morally weighted harm in H compared with that in the alternative is not justified by a countervailing marginal reduction in risked harm to the prospective victims of T.”
(Lazar, 2012, p. 14)

Lazar’s conception is carefully phrased and detailed, but the gist of it remains the core of necessity conditions: that in bringing about some bad (i.e. causing harm) in the pursuit of some good (i.e. avoiding future harm), one must bring about as little of this bad as possible.³⁴ This idea is simple and intuitive enough to be uncontroversial – how could it be permissible to do more bad, if one can achieve the same good while doing less?

However, as Lazar recognizes, once we include a necessity condition, we may have made narrow proportionality superfluous, since it appears that “[a] defensive harm’s satisfying *Necessity* is a sufficient but not a necessary condition for it’s being proportionate.” (Lazar, 2012, p. 21) How so? Necessity might initially seem to require only that the agent review all actions that would bring about some specific good G, e.g. averting a threat, and chooses one of the actions from this set which brings about the least bad. However, by insisting on reviewing only actions that bring about a specific good, this conception of necessity does not allow any trade-offs between goods and bads. In the context of self-defence, this would mean that for the good of averting a lethal attack, i.e. escaping unharmed, then killing the attacker is necessary even if one could have stunned the attacker, so long as stunning entails suffering any harm at all, e.g. a mild blow to the shoulder. As Lazar observes, this is implausible.

³⁴ The astute reader may have noticed that Lazar’s conception is both subjectivist and probabilist. Translated to its factualist version, the conception would be: “Defensive harm H is necessary to avert unjustified threat T if and only if there is no less harmful alternative, such that the marginal morally weighted harm in H compared with that in the alternative is not justified by a countervailing marginal reduction in harm to the prospective victims of T.” As before, I attempt in the following to remain neutral between factualism and the alternatives.

The proper range of actions to review therefore are not only actions that bring about the most good, but all actions that bring about the good at stake to some extent, and necessity must prefer the action from among this set that achieves the best balance of good over bad.³⁵ However, that set includes, at the limit, those actions that make no difference to the good, e.g. suffering the attack without resisting. Including those actions ensures that a necessary action is proportional, because, as we have previously seen, proportionality just requires that performing the action achieves a better balance of good over bad than inaction.³⁶

A suitable condition of necessity that in this way incorporates proportionality is a wide conception of proportionality. Revising the (narrow) account I set out at the beginning of this article, we can say of **wide teleologically proportional surveillance** that:

An act ϕ is widely proportional in the sense relevant to moral permissibility *iff* the balance of goods over bads that ϕ 'ing brings about is better than the balance of the same goods over bads of any alternative acts.

This, it seems to me, is the account we ultimately ought to prefer.

5.1 Is there still a role for narrow proportionality?

Lazar recognizes the possibility and advantages of a wide conception of proportionality in later writing:

³⁵ Is this an overly consequentialist reading of necessity? No. Necessity remains a different condition because it restricts its consideration to the bringing about of a particular good, rather than any effect on other goods. That is, it considers opportunity costs, but restricts the range of opportunity costs taken into consideration. Whether this is plausible is a different matter. Kieran Oberman argues, persuasively to my mind, that just war theorists have conventionally ruled out far too much, and that as a result waging war will often be impermissible simply because alternative uses of the resources would do far more good, e.g. by alleviating global poverty. (Oberman)

³⁶ As Macnish notes, Thomas Hurka makes a similar argument, albeit in the opposite direction. (Macnish, 2015, pp. 532-533) Hurka writes: "[The proportionality] condition must not be read literally, but must assess the alternatives to war in the same way as it does war: for the relevant good they may produce, their probability of producing it, and any costs that will result if the alternatives are tried and fail, such as making an eventual war more bloody. But then the last resort condition is in effect a comparative version of the initial, simple proportionality condition. For war and each of its alternatives it does a proportionality calculation, identifying the relevant goods and evils it will produce compared to a baseline of doing nothing, or continuing to act as one would have had there been no just cause. This yields the net good or bad effects of each, and it then says war is permitted only if its net outcome is better than those of all alternatives." (Hurka, 2005, pp. 37-38) McMahan too, although he considers them separate conditions, recognizes that proportionality and necessity are connected: "[T]he requirements of minimal force and narrow proportionality do begin to blur together when considerations of probability and risk arise, as they usually do in conditions of war." (McMahan, 2009, p. 23)

“The only difference between the proportionality and necessity constraints is that the former involves comparing one’s action with a very specific counterfactual scenario—in which we don’t act to avert the threat—while the latter involves comparing it with all your available options that have some prospect of averting or mitigating the threat. In my view, we should simply expand this so that the necessity constraint compares all your available options bar none. Then proportionality would essentially involve comparing each option with the alternative of doing nothing, while necessity would involve comparing all options (including doing nothing) in terms of their respective balances of goods and bads. On this approach, necessity would subsume proportionality.” (Lazar, 2017)

However, in his earlier work, Lazar presents an argument for the relevance of narrow proportionality which is worth briefly considering:

“It remains a necessary condition on permissible self-defense that the harm inflicted be proportionate to the threat averted. When the harm inflicted is disproportionate, it cannot be justified in the ethics of self-defense even if it satisfies *Necessity* (though it can be justified as a lesser evil). We can therefore see the proportionality calculation as a first stage in classifying the options available to Defender, setting out those that are admissible in the *Necessity* calculation, as potentially justified acts of self-defense.” (Lazar, 2012, p. 25)

How might an action fail to satisfy proportionality, while still satisfying necessity? Because, on Lazar’s definitions, proportionality compares the harm that various options inflict not with each other, but with an option-independent baseline of no harm. Consider Lazar’s example:

“**Pressure-1:** Attacker is going to unjustifiably break Defender’s wrist. Defender’s foot is on a pressure pad; if he keeps it there a lethal dart will be fired into Attacker. But if Defender removes his foot, that will trigger a second dart, which will hospitalize Attacker for a year.” (Lazar, 2012, p. 22)

Harm if you do; harm if you don’t. The harm caused by leaving your foot on the plate, Lazar argues, is unnecessary, because you have the option of lifting the foot and causing less harm, while equally effectively averting the same threat. However, neither option is proportionate, Lazar claims,

because even the milder harm, of spending a year in a coma, is vastly greater than the threat, of a broken wrist.³⁷

This seems to me mistaken, and the problem, I believe, concerns the way in which Lazar attributes harm to agents and actions. On the most plausible account of harm, the harm an action causes is evaluated counterfactually so that the amount of harm is itself a function of the available alternatives. (Holtug, 2002; Søbirk Petersen, 2014) Lazar's notion of the proportionality principle comparing options with doing no harm then collapses into necessity, because what "doing no harm" means is itself determined by reference to the agent's option set.

To see why this is more plausible, consider a slight variation on the pressure example:

Sticks and stones: Attacker is going to unjustifiably break Defender's wrist. If Defender stands still, then a falling rock will harmlessly bump into Defender, while a falling log hits attacker, killing her. If Defender ducks, the rock will hit Attacker, which will hospitalize her for a year, but also knock her out of the path of the log.

As in Lazar's example, suppose that ducking and standing still ("doing nothing") are the only two options available to Defender. Does ducking harm Attacker? Off-hand this sounds odd. Certainly, she is *better* off if Defender ducks than if Defender stands still. She will be worse off than she was immediately prior to Falling Rocks, but there is no scenario in which she gets to stay that way. Furthermore, in very similar situations, we would want to say that ducking prevents her suffering harm and thus *benefits* Attacker. Supposing, for example, that the two are friends, that there is no attack, and that ducking is an attempt to save Attacker's life, we would want to say that ducking is both permissible and praiseworthy.

However, the situation is exactly analogous in Lazar's *Pressure-1* example. As such, we must say that lifting your foot from the plate, which brings it about that Attacker suffers a year in hospital, does not harm Attacker; only keeping the foot on the plate causes any harm, and that harm is precisely equal to the difference in well-being between the two outcomes. If this seems counterintuitive, it is likely because we are misled by the way both outcomes in *Pressure-1* seem to originate with Defender, as they are determined by the position of her foot, in a way that it might superficially

³⁷ If this strikes the reader as not entirely intuitive, then simply scale the harms involved, e.g. a bruised wrist, instead of a broken one.

seem they do not in Ledge. This would be a mistake, based on an overly simplistic understanding of causality. Given that Defender has control over which of the two outcomes to bring about in both scenarios, she is equally causally responsible for the outcome that occurs. We may also be misled by implicitly supposing that Defender has created the situation in which Attacker will be either killed or hospitalized, e.g. stepping onto the pressure-plate in the first place. If this is the case then it is certainly morally relevant, and it means that Defender harmed Attacker by creating the very situation in which she now faces the choice of lifting her foot or standing still, but it does not support Lazar's idea of an action-independent baseline, and should be left out of the equation when evaluating the harm caused in *Pressure-1*.³⁸

Conclusion

Over the course of this article, I have explored the concept of proportionality, set out a teleological account of it, and critically reviewed some possible deontological modifications. In the final part of the article, I have argued that narrow proportionality should be replaced with wide proportionality, in practice a suitably conceived necessity-condition.

Narrow proportionality may retain some use in practice, as a form of (more) easily applicable initial test. Narrow Proportionality may often be a relatively simple matter to assess, compared to wide proportionality, and practices of surveillance that fail to satisfy narrow proportionality can be rejected at the outset, which will leave fewer candidates for the complete, more demanding test.

Some, particularly friends of deontology, may want to modify the teleological account in places. I have discussed three potential modifications that feature prominently in the literature, and argued that we should not adopt them, but this does not, of course, exhaust the field of potential modifications. Until efforts to develop and substantiate such modifications succeed, the teleological account can hopefully serve as a useful element in a theory of permissible surveillance.

It can likely serve as no more than that. Both consequentialists and deontologists will consider even the wide proportionality condition insufficient for moral permissibility. Consequentialists will want

³⁸ The doing-allowing distinction may also affect intuitions here. I assume that it should not since it is not a difference upon which Lazar's argument rests. He would presumably want to treat as equal for the purposes of assessing proportionality scenarios in which we reverse the outcomes attached to the actions, so that e.g. lifting your foot leads to death and resting your foot to hospitalization. As such, we should attempt to set aside any intuitions about the difference between doing and allowing, e.g. by assuming that the difference has been incorporated in the "moral weights" we assign to the harm(s).

to expand the set of actions evaluated to include actions that bring about entirely different goods. Deontologists will want to add supplementary conditions, such as just cause and morally benign intentions. Developing a full account of permissible surveillance is a complex task. I hope only to here have advanced some way towards that goal.

References

- Arneson, R. J. (2009). What do we owe to distant needy strangers? In J. A. Schaler (Ed.), *Peter Singer Under Fire: The Moral Iconoclast Faces His Critics* (pp. 267-293). Chicago and LaSalle: Open Court.
- Ashworth, A. (2005). *Sentencing and Criminal Justice* (4th ed.). Cambridge: Cambridge University Press.
- Austin, L. M. (2003). Privacy and the question of technology. *Law and Philosophy*, 22(2), 119-166.
- Bellaby, R. (2012). What's the Harm? The Ethics of Intelligence Collection. *Intelligence and National Security*, 27(1), 93-117. doi:10.1080/02684527.2012.621600
- Boonin, D. (2008). *The Problem of Punishment*. New York: Cambridge University Press.
- Broome, J. (2005). Should We Value Population?*. 13(4), 399-413. doi:10.1111/j.1467-9760.2005.00230.x
- Broome, J. (2012). *Climate Matters - Ethics in a Warming World*. New York: W. W. Norton & Co.
- Brown, C. (2011). Consequentialise This. *Ethics*, 121(4), 749-771.
- Doyle, T. (2009). Privacy and perfect voyeurism. *Ethics and Information Technology*, 11(3), 181-189. doi:10.1007/s10676-009-9195-9
- Duff, A. (2001). *Punishment, Communication, and Community*. Oxford: Oxford University Press.
- Duus-Otterström, G. (2013). Why Retributivists Should Endorse Leniency in Punishment. *Law and Philosophy*, 32(4), 459-483. doi:10.1007/s10982-012-9147-0
- Feldman, F. (1995a). Adjusting Utility for Justice: A Consequentialist Reply to the Objection from Justice. *Philosophy and Phenomenological Research*, 55(3), 567-585.
- Feldman, F. (1995b). Desert: Reconsideration of Some Received Wisdom. *Mind*, 104(413), 63-77.
- Gardiner, S. M. (2006). A core precautionary principle. *Journal of Political Philosophy*, 14(1), 33-60.
- Hadjimatheou, K. (2016). Surveillance Technologies, Wrongful Criminalisation, and the Presumption of Innocence. *Philosophy & Technology*, 1-16. doi:10.1007/s13347-016-0218-2
- Harris, J., & Holm, S. (2002). Extending human lifespan and the precautionary paradox. *Journal of Medicine and Philosophy*, 27(3), 355 - 368.
- Hartzell-Nichols, L. (2013). From 'the' Precautionary Principle to Precautionary Principles. *Ethics, Policy & Environment*, 16(3), 308-320. doi:10.1080/21550085.2013.844569
- Holtug, N. (2002). The Harm Principle. *Ethical Theory and Moral Practice*, 5(4), 357-389. doi:10.1023/a:1021328520077
- Hurka, T. (2005). Proportionality in the Morality of War. *Philosophy & Public Affairs*, 33(1), 34-66. doi:10.1111/j.1088-4963.2005.00024.x
- Kant, I. (1996). *The Metaphysics of Morals* (M. Gregor, Trans.). Cambridge: Cambridge University Press.
- Lazar, S. (2012). Necessity in Self-Defense and War. *Philosophy and Public Affairs*, 40(1), 3-44.

- Lazar, S. (2017). War. In E. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*.
- Macnish, K. Surveillance Ethics. In *The Internet Encyclopedia of Philosophy*.
- Macnish, K. (2014). Just Surveillance? Towards a Normative Theory of Surveillance. *Surveillance & Society*, 12(1), 142-153.
- Macnish, K. (2015). An Eye for an Eye: Proportionality and Surveillance. *Ethical Theory and Moral Practice*, 18(3), 529-548.
- Macnish, K. (2016). Government Surveillance and Why Defining Privacy Matters in a Post-Snowden World. *Journal of Applied Philosophy*, 33(3).
- Macnish, K. (2017). *The Ethics of Surveillance: An Introduction*: Routledge.
- McCloskey, H. J. (1980). Privacy and the Right to Privacy. *Philosophy*, 55(211), 17 - 38.
- McMahan, J. (1993). Killing, Letting Die, and Withdrawing Aid. *Ethics*, 103(2), 250-279.
- McMahan, J. (2009). *Killing in War*. Oxford: Oxford University Press.
- Mintz-Woo, K. (2019). Principled Utility Discounting Under Risk. In *Moral Philosophy and Politics* (Vol. 6, pp. 89).
- Moore, G. E. (2000). *Principia Ethica*. Cambridge: Cambridge University Press.
- Moore, M. S. (1997). *Placing blame: a general theory of the criminal law*: Oxford University Press, Incorporated.
- Nissenbaum, H. (1998). Protecting Privacy in an Information Age: The Problem of Privacy in Public. *Law and Philosophy*, 17(5), 559-596. doi:10.1023/A:1006184504201
- Norcross, A. (1997). Comparing Harms: Headaches and Human Lives. *Philosophy & Public Affairs*, 26(2), 135-167. doi:doi:10.1111/j.1088-4963.1997.tb00079.x
- Norcross, A. (2003). Killing and letting die. In R. G. Frey & C. H. Wellman (Eds.), *A Companion to Applied Ethics* (pp. 451-463).
- Oberman, K. War and poverty. *Philosophical Studies*, 1-21.
- Parent, W. A. (1983). Privacy, morality, and the law. *Philosophy and Public Affairs*, 12(4), 269-288.
- Peterson, M. (2006). The Precautionary Principle Is Incoherent. *Risk Analysis*, 26(3), 595-601. doi:10.1111/j.1539-6924.2006.00781.x
- Powers, M. (1996). A cognitive access definition of privacy. *Law and Philosophy*, 15(4), 369 - 386.
- Rachels, J. (1975). Active and Passive Euthanasia. *New England Journal of Medicine*, 292(2), 78-80. doi:10.1056/nejm197501092920206
- Rachels, J. (1975). Why privacy is important. *Philosophy & Public Affairs*, 4(4), 323-333.
- Rodin, D. (2011). Justifying Harm. *Ethics*, 122(1), 74-110.
- Rubel, A. (2007). Claims to Privacy and the Distributed Value View. *San Diego Law Review*, 44(4), 921-956.
- Rubel, A. (2011). The Particularized Judgment Account of Privacy. *Res Publica*, 17(3), 275-290.
- Ryberg, J. (2007). Privacy rights, crime prevention, CCTV, and the life of Mrs. Armac. *Res Publica*, 13(2), 127-143.
- Ryberg, J. (2011). Punishment and Desert-adjusted Utilitarianism. In M. Tonry (Ed.), *Retributivism Has a Past - Has it a Future?* (pp. 86-100). Oxford: Oxford University Press.
- Sandin, P., & Peterson, M. (2019). Is the Precautionary Principle a Midlevel Principle? *Journal of Ethics, Policy and Environment*, 22(1), 34-48.
- Scanlon, T. (1975). Thomson on Privacy. *Philosophy & Public Affairs*, 4(4), 315-322. doi:10.2307/2265076
- Scanlon, T. (2008). *Moral Dimensions: Permissibility, Meaning, and Blame*. Cambridge: Belknap Press.

- Solove, D. J. (2002). Conceptualizing Privacy. *California Law Review*, 90(4), 1088-1155.
- Solove, D. J. (2006). A Taxonomy of Privacy. *University of Pennsylvania Law Review*, 154(3), 477-560.
- Solove, D. J. (2007). 'I've Got Nothing to Hide' and Other Misunderstandings of Privacy. *San Diego Law Review*, 44, 745-772.
- Søbirk Petersen, T. (2014). Being Worse Off: But in Comparison with What? On the Baseline Problem of Harm and the Harm Principle. *Res Publica*, 20(2), 199-214.
- Tadros, V. (2011). *The Ends of Harm: The Moral Foundations of Criminal Law*. Oxford: Oxford University Press.
- Tavani, H. T. (1999). KDD, data mining, and the challenge for normative privacy. *Ethics and Information Technology*, 1(4), 265-273.
- Tavani, H. T. (2007). Philosophical theories of privacy: Implications for an adequate online privacy policy. *Metaphilosophy*, 38(1), 1-22.
- Tavani, H. T., & Moor, J. H. (2001). Privacy protection, control of information, and privacy-enhancing technologies. *SIGCAS Comput. Soc.*, 31(1), 6-11. doi:10.1145/572277.572278
- Thomson, J. J. (1975). The Right to Privacy. *Philosophy & Public Affairs*, 4(4), 295-314. doi:10.2307/2265075
- Uniacke, S. (2011). Proportionality and Self-Defense. *Law and Philosophy*, 30(3), 253-272.
- Von Hirsch, A. (2000). The Ethics of Public Television Surveillance. In A. Von Hirsch, D. Garland, & A. Wakefield (Eds.), *Ethical and Social Perspectives on Situational Crime Prevention* (pp. 59-76). Oxford: Hart.
- von Hirsch, A. (2009). Proportionate Sentences: A Desert Perspective. In A. Von Hirsch, A. Ashworth, & J. V. Roberts (Eds.), *Principled Sentencing* (3rd ed., pp. 115-125). Oxford: Hart Publishing.
- Von Hirsch, A., & Ashworth, A. (2005). *Proportionate Sentencing*. Oxford: Oxford University Press.
- Woollard, F. (2015). *Doing and Allowing Harm*: Oxford University Press.
- Woollard, F. (2016). Doing vs. Allowing Harm. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*.