



AARHUS UNIVERSITY



Coversheet

This is the accepted manuscript (post-print version) of the article.

Contentwise, the post-print version is identical to the final published version, but there may be differences in typography and layout.

How to cite this publication

Please cite the final published version:

Rønn, K. V., & Lippert-Rasmussen, K. (2020). Out of Proportion? On Surveillance and the Proportionality Requirement. *Ethical Theory and Moral Practice*, 23(1), 181-199.

<https://doi.org/10.1007/s10677-019-10057-z>

Publication metadata

Title: Out of Proportion? On Surveillance and the Proportionality Requirement
Author(s): Kira Vrist Rønn, Kasper Lippert-Rasmussen
Journal: *Ethical Theory and Moral Practice*, 23(1), 181-199
DOI/Link: <https://doi.org/10.1007/s10677-019-10057-z>
Document version: Accepted manuscript (post-print)

General Rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Out of proportion?

Making the case for *wide* and *narrow* proportionality in the context of surveillance

Introduction

In public critiques of surveillance, it is often claimed that specific deployed measures are *disproportionate*. Following Snowden's revelations concerning the extensive mass surveillance programs of the NSA and subsequently of GCHQ, the disproportionality label has been frequently applied in debates surrounding mass surveillance. This is true, for example, in public discussions of NSA's so-called PRISM program and GCHQ's TEMPORA program. These programs respectively allow storage of huge amounts of data from cables and seek to store telephone and internet communication data from random users for potential later use (Stahl 2016).

Here, disproportionality seems to be associated with the idea that specific means of surveillance range from unjust, unnecessary, indiscriminate, to illegitimate. The blurred meaning of the concept of 'disproportionate surveillance' in everyday language is unsurprising. What is more surprising is the apparent blurriness of the concept of disproportionate and proportionate surveillance in the context of moral philosophy (Macnish 2015). Nearly all attempts to provide an ethical framework for surveillance, both from a practical hands-on perspective, as exemplified in the ethical code of the British Home Office (Home Office 2010), and from the work of moral scholars such as Gary Marx (1998), John Kleinig (2009), and Kevin Macnish (2014, 2015), emphasize the importance of proportionality of surveillance measures. In surveillance, proportionality can, roughly speaking, refer to the equation that requires the benefits of surveillance to outweigh, morally speaking, the costs of surveillance. Kleinig, Marx, and Macnish all specify that the costs of surveillance – often phrased as 'harms' – can be understood as violations of the right to privacy, yet the elements of the proportionality balance rarely receive much further specification (with Macnish 2015 being an exception). As a result, both policy documents and scholarly work on the ethics of surveillance generally agree that proportionality constitutes at a minimum a necessary condition for ethical surveillance, yet further clarification of the proportionality principle is needed (Macnish 2015).

In this article, I scrutinize how the concept of proportionality can be understood in surveillance. The paper is structured as follows: In the *first section*, I present the main elements of the proportionality principle in surveillance in order to enable a framing of my argument within the existing literature on the topic. In the *second section*, I specify two aspects present in the existing scholarly work on proportionate surveillance: First, I show that not only harms but also *harmless wrongs* should count on the cost side of the proportionality principle equation in the context of surveillance. Second, I argue that the analogy from just warfare and self-defense, entailing an aggressor and a defender, will not always apply in just surveillance since, in some types of surveillance (e.g. mass surveillance), *the surveilled* is not an aggressor. The analogy from just warfare and self-defense thus requires adjustment. In *section three*, I return to my main argument: that proportionality should be understood differently depending on *the liability* of the surveilled and *the intentions* of the surveiller. I show that McMahan's JWT categorization of *wide* and *narrow proportionality* generally applies to surveillance. This entails that the liability of the surveilled and the intentions of the surveiller play crucial roles in determining whether an act of surveillance is proportionate. Previously, both Nathan (2017) and Macnish (2015) have regarded the liability of the surveilled as a central element in determining proportionality in surveillance. The intention of the surveiller has, however, not been specifically addressed and combined with the question concerning liability. I argue that the intentions of the surveiller also matter for the understanding of proportionate surveillance. Accordingly, I discuss how the notion of wide proportionality applies to cases in which the surveilled is not liable to suffer from the intended and unintended intrusion entailed in the surveillance measure, as exemplified by two types of mass surveillance. I argue that a stricter understanding of wide proportionality is needed in cases, characterized by *intentional* intrusion upon non-liable agents, and that the proportionality principle of such cases of surveillance is more demanding than merely balancing in favor of *the lesser evil*, as prescribed by McMahan's notion of wide proportionality.

Elements of the proportionality principle

In this section, I present the existing distinctions pertaining to the scholarly literature on proportionate surveillance. Proportionate surveillance¹ has previously been addressed by the scholars Kevin Macnish (2015) and Christopher Nathan (2017) (with specific attention to the liability to harms inflicted as a result of covert police activities), both of whom draw upon the proportionality condition

¹ In this paper I understand 'surveillance' as "... any systematic and routine attention to personal details, whether specific or aggregate, for a defined purpose" (Lyon 2014:2).

from other philosophical domains, the Just War Theory (JWT) and the theory of permissible individual self-defense respectively.

The fact that distinctions from the contexts of just warfare and just self-defense are transferred to surveillance does not imply an equation between warfare and surveillance. On the contrary, the ‘just surveillance context’ is considered one that could benefit from distinctions and discussions from other domains similar to an analogy or a principle with ‘common roots’, as emphasized by Nathan (2017, p. 7). This clarification is important since the transfer of principles from just wars to other fields of research often creates misguided resistance to these arguments since opponents, for example, read an equation of the application of JWT as an endorsement of, for example, the *militarization* of government activities, such as surveillance programs, which are not normally considered part of warfare (Stoddart 2014; [OMITTED]).

Both JWT and just self-defense distinguish between an identified *threatener* and an identified *defender*, and when defining proportionate surveillance, both scholars pay specific attention to the conditions of the permissible infliction of defensive harms as a result of counteracting an unjust threat (Macnish 2015; Nathan 2017). Accordingly, in surveillance, the role of the threatener is occupied by *the surveilled*, and the role of the defender is occupied by the *surveiller*.

According to Macnish, the proportionality principle in surveillance includes the following three elements: 1) *the benefits* (or the avoided harms) connected to employing the specific surveillance activity, 2) *the harms* (or costs) of a specific surveillance activity, and 3) *the appropriate balancing* of the harms and the benefits of a specific surveillance activity.

First, the benefits or goods of a specific surveillance activity could be phrased as the acquisition of information that would (amongst other means) contribute to preventing or disrupting a specific unjust threat.² Accordingly, this side of the equation will reflect the severity of the threat that is potentially prevented/disrupted as a result of actions based on the specific surveillance activity. The proportionality balance is also often phrased in terms of *harms* on both sides of the equation since the benefits of surveillance include the foreseeable *avoided harms* connected to acquiring information. Thus, if a specific threat concerns a terrorist attack foreseeably resulting in the deaths of hundreds of individuals, avoidance of such harms would count heavily as potential benefits of a specific

² Macnish furthermore appeals to the distinctions provided by Thomas Hurka including a specification of which types of goods to include in this part of the proportionality balance (Macnish 2014, pp. 534-535).

surveillance activity addressed at the threatener. In contrast, the harms entailed in a case concerning a pair of parents' false claims of residency in a specific school district in order to gain access to a popular school would count significantly lower (an example provided by Macnish 2015). As a result, 'the gravity and extent of the perceived crime or offence' counts significantly on this side of the equation (Macnish 2015, p. 530).

The second aspect of the proportionality principle concerns the harms associated with applying the specific surveillance activity. According to Macnish, all types of 'reasonably foreseeable harms arising from surveillance' should be included in the equation, and this side of the equation should not include the same kind of leveling as in the case of the goods of surveillance (Macnish 2015, p. 537, p. 543).

The third element of the proportional principle of surveillance is the comparison of the goods and the harms of the specific surveillance activity. Susan Uniacke, for example, specifies (in the context of just self-defense) that proportionality is both *a relational* and *a normative* concept: It is relational in the sense that 'it involves a ratio or comparison of scale between x and y' in which the two variables should be suitable or adequate, and it is normative in the sense that proportionality requires *an appropriate balance* between the two elements (Uniacke 2011, p. 255). Disproportionate self-defense thus entails that 'x is either excessive or deficient on the relevant side' (Uniacke 2011). A tricky aspect of the proportionality principle is thus assessing the appropriateness of a specific balance between x and y. Macnish lists some versions of how the balance could be understood in the context of surveillance, for instance, in terms of the goods *outweighing* the harms; the production of a *great deal more* goods than harms; or the middle path, in which 'the benefits should *equal* the harms' (Macnish 2015 – my emphasis).

Critiques of the proportionality principle state that such a threshold or balance is impossible to strike (D'Amico 2015 (in the context of proportionate punishment), O'Donovan 2003 (in the context of just war). Macnish, like Hurka, acknowledges the unrealistic aim of 'strict proportionality', especially in the so-called 'twilight zone' (Macnish 2015, p. 539). Macnish, however, argues that despite the difficulties related to determining the balance in a strict sense, proportionality should still play a crucial role when judging the moral permissibility of surveillance, assuming disproportionality in the blurred cases and leaving it to the specific case officer to prove proportionality (Macnish 2015, p. 539).

Harms and harmless wrongs of surveillance

The above presentation of the existing distinctions of the proportionality principle in the context of surveillance will constitute the starting point for my elaboration of proportionate surveillance. Hence, before applying to surveillance McMahan's notion of wide and narrow proportionality, I briefly specify two elements of the previous use of the principles from JWT in a surveillance context, which is not entailed in the existing literature.

First, *harm* is generally applied as the central term when explaining why an action in warfare should be morally restricted (McMahan 2009a). When applying the principles from JWT and self-defense to surveillance, the harms inflicted upon the surveilled as a result of a specific surveillance activity are similarly emphasized as relevant to the moral permissibility of that specific activity (Bellaby 2012, 2014; Macnish 2014, 2015; Nathan 2017). Furthermore, the harms connected to surveillance activities are often specified with reference to the violation of the individual right to privacy. The degree of harm thus varies depending on the specific surveillance activity, and Bellaby and Macnish argue as a result that some surveillance measures obviously inflict more harm on the surveilled than others. The violation of the right to privacy is most often referred to as the central harm, whereas deception and manipulation are applied when stating, for example, why the use of false flag or bugging is harmful to the affected individual. Thus, as stated by Marx (1998, p. 173), 'privacy' is 'a vague catch-all phrase that includes a variety of concerns, such as respect for the personhood, dignity, and autonomy of the individual'. I will not go into more detail as to how various types of surveillance might harm affected individuals. I will, however, argue that the transfer of the term 'harm' as the key concern of surveillance activities (e.g. by Macnish and Bellaby) is inaccurate in some cases. If harm is generally understood as a decrease in the welfare of the affected individuals, then some types of surveillance do not necessarily harm the affected individuals, simply because they are never aware of the surveillance activity, and their welfare is, all else being equal, unaffected by the surveillance activity. Harms are much more obvious in warfare and self-defense, where most harms are physical. Indeed, surveillance activity will not inflict the same kinds of harm upon the affected individuals as will warfare, and harm might not be the best term in this context. Other substantial moral norms besides the 'do-no-harm-principle' (Bellaby 2012, p. 96) are at stake when it comes to surveillance. Mass surveillance constitutes just one example in which the affected individual of a surveillance activity will not necessarily be *harmed* as a result of the surveillance activity. This is the case because the surveilled individual is unaware of the surveillance activity, and her welfare is consequently not

decreased due to the surveillance activity. Arguably this individual is unaffected. Her life would have been indistinguishable from what it actually is had she not been surveilled in this manner.

The claim that, for example, mass surveillance does not inflict harms does not, however, entail that such activities are morally unproblematic. My claim is that some surveillance activities, for instance mass surveillance, entail *wrongs* but not necessarily harms. As a result, *harmless wrongs* may be a more precise term than *wrongs* when discussing the negative effects of various surveillance activities such as mass surveillance. Consider, for example, a case in which person A takes nude photos of person B and shares them without B knowing and without the episode affecting B's subsequent life. In this case, B is not harmed by the episode because B's welfare is unaffected. B has nevertheless been wronged by A's actions. The distinction between harms and harmless wrongs has not previously been addressed in discussions of ethical surveillance, yet it is important when addressing the proportionality of surveillance activities. The cost side of the proportionality equation should thus be specified in the context of surveillance: It is necessary to address not just harms but also harmless wrongs when considering the question of proportionate surveillance.

Additionally, I argue that mass surveillance and other types of surveillance entailing only harmless wrongs can still be harmful, just not in the direct way in which we normally think about harms. Imagine, for example, a situation in which people know that government authorities apply mass surveillance. Each individual does not, however, know whether she is a target of the government mass surveillance activity. The harm entailed in the mass surveillance activity will, then, best be understood in terms of *the fear of being watched* and will affect the ways in which individuals behave in their everyday lives.³ This type of harm can be present independent of whether a person is in fact surveilled and in this sense becomes a *wrongless harm* since the action will not constitute the rights-violation of a specific individual but will nevertheless be a morally relevant aspect of surveillance.

Bearing this specification between harms and harmless wrongs in mind, a second specification is relevant when applying the analogy from JWT and self-defense to surveillance. Previously, when applying principles and distinctions from JWT and self-defense, the main focus has been on a simple notion of the 'defender-aggressor' allegory (Macnish 2014, 2015; Nathan 2017). Both Macnish and Nathan focus on types of surveillance in which a threatener is identified and in which surveillance activity is initiated to counteract or defend against that specific threat. Some surveillance activities

³ This description of the wrongless harms of surveillance resembles the so-called "chilling effect" described by e.g. Macnish 2015.

do not, however, fit this original notion of who is in fact playing the respective roles of threatener and defender. Imagine again the case mass surveillance in which personal details about the internet behavior of random individuals are collected by government authorities in order to potentially apply them in future situations. In this case, there is no aggressor, and one could even argue that the roles of threatener and defender are switched since the surveilled in no way constitutes an aggressor or an unjust threat. Thus, in this case, the surveiller seems to become the aggressor while the surveilled becomes the defender, and the situation is blurrier than when considering cases in which the surveiller constitutes the defender and the surveilled is the aggressor.

Emphasizing liability and intentions in surveillance

As I argued above, some of the distinctions from JWT and self-defense require clarification and nuancing when applied to surveillance. In the following sections, I consider the roles played by the liability and intentions of the surveilled individuals when specifying the concept of proportionality in surveillance.

Being liable to surveillance

Besides the three elements (costs, benefits, and an appropriate balance) presented above, many scholars dealing with the ethics of surveillance would agree that the *liability* of the surveilled should be included when addressing proportionality as a moral principle. Although Macnish addresses the importance of distinguishing between liable and non-liable agents of surveillance and focuses his attention on surveillance of liable agents, he does not provide a specific account of what it takes to become liable to surveillance (Macnish 2015, p. 542).

Other scholars working in self-defense and warfare have elaborated upon the specific notion of being liable to suffer from defensive harms (e.g. McMahan 2005, 2009; Uniacke 2011; Nathan 2017). Jeff McMahan, for example, regards proportionality as a constraint on an action that causes harm. The harms are thus ‘instrumental to the achievement of some valuable goal against which the harm can be weighted’ (McMahan 2009a, p. 19). Furthermore, he specifies that ‘The goal sought must usually be a solution to a problem’. Additionally, it is not enough that a person poses a threat in order to be liable – it should be an unjust threat (McMahan 2009a, p. 157). Thus, being liable to a defensive harm in war requires, according to McMahan, that the inflicted harms solve the specific problem and that the threat avoided by inflicting harm upon a threatener is unjust. Being liable to suffer from a

defensive harm then entails that ‘a person *is implicated* in the existence of a problem in such a way that harming him in a certain way on the course of solving the problem would not wrong him’ (McMahan 2009a, p. 19 – my emphasis).

Nathan (2017) has transferred some of the distinctions from self-defense into discussions concerning the permissibility of covert police work. According to Nathan (2017, p. 5 – my emphasis), ‘proportionality in self-defense involves not only the harm threatened and the harm used to avert the threat, but also the attacker’s *culpability* for the threat’. Thus, in Nathan’s ‘liability model’, emphasis is not just placed on ‘implication’ in a problem, as stated by McMahan, but also on culpability for a specific threat. Furthermore, Nathan argues that culpability means that an individual is ‘blameworthy for a threat that is unjustified’ (Nathan 2017, p. 4). The fact that the target of covert police actions is ‘culpable for a possible harm’ will ‘cancel the moral complaints’ that the violation of her rights would otherwise entail (Nathan 2017, pp. 3-5). Nathan thus favors an epistemic approach to culpability in which emphasis is placed on whether a police officer has a reasonable belief that somebody is posing an unjust threat (Nathan 2017, p. 8). As a result, according to Nathan, actions that cause government authorities to believe that one is responsible for creating an unjust threat are important for a moral principle of proportionality of surveillance activities. Nathan further specifies that if a person A has acted in way that makes others believe that A is posing an unjust threat even though A is in fact not a threat (e.g. exemplified by a case in which someone dresses up as a bank robber and walks out of a bank with a bag full of her own money as part of an art project), then A is not wronged if the police reasonably but falsely believe she poses an unjust threat (Nathan 2017).

The question is whether Nathan’s liability model is applicable to surveillance. Imagine a case in which individuals who make frequent visits to certain websites with content reflecting radical extremist perspectives come under government surveillance as a part of a counterterrorism strategy. Suppose a researcher B frequently visits such websites in the process of conducting research into radical networks. B subsequently undergoes government surveillance as a result of the counterterrorism strategy. According to the liability model, B would be liable to the surveillance since B is causing government authorities to believe that B poses an unjust threat even though this is in fact false. Accordingly, B will not be wronged by the surveillance since she is liable to the surveillance. Even though B might know that she runs a high risk of being surveilled on the basis of her visits to the webpages, the conclusion that she is liable to subsequent surveillance is discussable. My claim

would be that B is wronged even though she might cause government authorities to believe that she poses an unjust threat. McMahan's understanding of liability as involvement in a specific unjust threat thus constitutes an alternative to Nathan's liability model, in which culpability and reasonable belief plays the central role. I will not go into more detail here concerning the distinctions of liability, but I will return to the question of moral justification later. First though, I will address another necessary specification when applying the liability model to surveillance.

The distinct nature of the threat in surveillance

The nature of the threat in surveillance is often somewhat different than in either warfare and self-defense, in which threats are often manifested and already occurring (e.g. pointing guns or falling bodies) (Frowe 2011). In surveillance, a threat might often better be understood in terms of *risk factors*, and a threat seems (at least in some cases) to best be understood as *a proxy for probability* that an unjust event will occur. Consider again the case of counterterrorism and the application of surveillance based on knowledge of the internet behavior of specific individuals who visit extremist websites. The internet behavior itself does not constitute a threat. Instead, the frequent visits to extremist websites, together with other risk factors, indicate a specific probability of a future unjust threat that has not yet materialized.

Nathan (2017, p. 9) himself recognizes a distinction between ascribing liability in the contexts of permissible self-defense and covert police work as a difference in 'immediacy'. He highlights this difference in order to emphasize that the police often have much more time available to consider and analyze the situation before acting, which is not typical in cases of self-defense, in which a person acts more intuitively in the specific moment of an occurring threat.⁴ This point does not, however, capture the distinct difference in the nature of the threat in cases of self-defense and cases of surveillance. Means of surveillance are thus sometimes applied to *enhance the epistemic situation* of the authorities in order to enable their identification of potential unjust threats (e.g. in the case of some types of mass surveillance). Some surveillance activities thus constitute means of identifying (future) unjust threats, creating an epistemic challenge for justifying specific surveillance activity since knowledge about the threat is first needed to determine liability and moral permissibility [OMITTED]. In other cases, the authorities' epistemic situation is better, and surveillance is initiated

⁴ The professional skills of the police investigators (seeking the truth) will minimize the risk of reasonable but false beliefs according to Nathan (2017).

on the basis of evidence that somebody is in fact (or most likely) posing an unjust threat. It is thus useful to distinguish between *offensive (preventive)* and *defensive surveillance activities* in order to denote the specific nature of the unjust threat entailed in each case. The literature on surveillance ethics has thus far primarily addressed cases of (defensive) surveillance in which government authorities have identified agents of manifested, unjust threats (Macnish 2014, 2015; Nathan 2017). The circularity entailed in the epistemic requirement of assessing liability is, however, not addressed in the previous literature. Thus, Nathan (2017, p. 11) argues that cases with no crime or ‘only possible future threats’ should be considered ‘analogous to harming the innocent bystander in exchange for a significant benefit’. I will return to this claim concerning the moral justification of harming innocents later. For now, I will continue on the assumption that identification of liability to surveillance is possible in cases in which the epistemic requirements for identifying unjust threats have been fulfilled.

The question of liability in the case of surveillance is in general somewhat underdeveloped in surveillance. Macnish mainly addresses liability as a question of the ability to distinguish between combatants and non-combatants, but he does not account for how legitimate targets of surveillance should be identified (Macnish 2015, pp. 542-543). As shown above, Nathan provides some convincing arguments in the context of covert policing. Yet, the question concerning the role and understanding of liability in surveillance has only been tentatively addressed in the literature on surveillance ethics in general. In the rest of the paper, I mainly address cases of surveillance (primarily mass surveillance) in which the affected individuals are not liable to the surveillance activities since the surveilled individuals are neither blameworthy nor implicated in the existence of the problem that is potentially solved by the activity.

Wide and narrow proportionality

Before elaborating further upon the role of liability in surveillance, I consider the existing discussions concerning proportionate surveillance by introducing McMahan’s distinction between *narrow* and *wide proportionality*. The distinction between narrow and wide proportionality captures precisely the difference between the respective harms addressed at individuals who are liable and at those who are not liable to defensive harms in the case of warfare.

Narrow proportionality is articulated as a relationship between a specific threatener and a specific defender and is therefore phrased as an agent-based condition and furthermore as based on a liability justification – since the crucial element here is whether the threatener is liable to suffer harms

potentially inflicted on him by the defender, meaning that the threatener is not wronged by the inflicted harm (McMahan 2009a, p. 20). The nature and degree of the harm inflicted upon the threatener (together with his degree of involvement in the threat) defines *an upper limit* for the harm that may be inflicted upon the threatener by the defender. Inflicting more harm is then disproportionate ‘since he is only potentially liable to lesser harm’ (Uniacke 2011, p. 270). As a result, narrow proportionality does not suggest a specific degree of harm for which the threatener is liable but suggests only an upper limit. Inflicting less harm on the threatener than the upper limit dictates would not be a case of disproportionality (McMahan 2009a).

Wide proportionality is articulated as inflicted harms upon non-liable agents. Inflicting such harms is considered aggressive, in contrast to the defensive harms inflicted in the case of narrow proportionality, in which the agent is liable to suffer from defensive harm due to his involvement in the problem (McMahan 2009a, p. 21; Uniacke 2011). Determining wide proportionality depends upon the specific context and entails the assumption that harming non-liable individuals creates a better situation than the projected alternative, which is prevented by inflicting harm on those individuals. Wide proportionality is thus judged in accordance with *lesser evil justifications*. As a result, wide proportionality rests on the assumption that if one must choose between 1) a scenario in which some smaller harm is inflicted upon non-liable individuals in order to avoid a greater degree of harm for the non-liable individuals and 2) a scenario in which the minor harm for the non-liable individuals is not inflicted and the greater harm occurs, then the first scenario should be preferred (Uniacke 2011, p. 266). The lesser evil justification is thus not simply a consequentialist assessment of costs and benefits associated with a specific action but instead involves a preference for cases entailing foreseeable harming of non-liable individuals in order to avoid greater harm for the non-liable.

McMahan provides even more detail to his distinction between the two types of proportionality. Besides liability, the *intentions* of the defender when inflicting harm are also relevant. McMahan divides the various cases of intention and liability into the following four categories of proportionality requirements: 1) intentionally harming the liable, 2) unintentionally but foreseeably harming the liable, 3) intentionally harming the non-liable and 4) unintentionally but foreseeably harming the non-liable (McMahan 2009a, pp. 19-20)

The first type of action is the one most often addressed in philosophical discussions on self-defense since this type includes cases with a specific defender who intentionally inflicts harm on a specific threatener who is (potentially) liable to suffer from the specific harm (McMahan 2009a). The fourth case is, on the other hand, the type of case most often referred to in discussions on disproportionality in warfare in connection to ‘collateral damage’, when bystanders are foreseeably but unintentionally harmed as a result of military action (McMahan 2009a). Narrow proportionality would, according to McMahan, apply to the first two types of cases entailing liable agents, whereas wide proportionality would (intuitively) apply to the third and the fourth cases, in which the harm is inflicted upon non-liable individuals (McMahan 2009a). The respective roles of liability and harm in the two types of proportionality are illustrated in the table below:

Table 1: Requirements of wide and narrow proportionality

	Intentional harm	Unintentional but foreseeable harm
Liable individuals	<i>Narrow proportionality</i>	
Non-liable individuals	<i>Wide proportionality</i>	

In discussions concerning proportionate surveillance, both Nathan and Macnish primarily consider cases and distinctions pertaining to the first category, in which the harms are intentionally inflicted upon a liable individual (Macnish 2015, p. 542; Nathan 2017). In the following, however, I go into greater depth concerning the proportionality principle of surveillance when the surveillance affects the lives of non-liable agents. I begin, however, with a specification of the four categories of proportionality requirements in surveillance.

Proportionality requirements and cases of surveillance

1) Intentional intrusion upon liable individuals

An illustration of the first type of case could be the example provided by Nathan concerning a man who we reasonably believe is taking part in an online child pornography distribution ring and who we wish to target with covert police actions, i.e. by surveilling his online communication (Nathan 2016, p. 4). In this case, the intrusion entailed in the surveillance measure are intentionally inflicted

upon an individual who we believe is liable to suffer from defensive harms due to his active involvement in online child pornography distribution, which could be considered an unjust threat to those children who are harmed by his (and the group's) actions.

If McMahan's distinctions between narrow and wide proportionality apply to surveillance, then the surveillance activities presented in this first case would adhere to narrow proportionality. More explicitly, the surveillance measures applied against a man involved in online child pornography would be disproportionate if the level of intrusion imposed upon him as a result of the applied surveillance measure exceeds the upper limit for the permissible level of intrusion based on the severity of the threat and the degree of his involvement in the case.

When applying a concept of narrow proportionality in such cases of means of surveillance, the intrusion of surveillance measures is deemed defensive, and the surveilled is not wronged as a result of the intrusion caused by the surveillance measure since he is liable to suffer from the intrusion if the level of intrusion does not exceed the threshold pertaining to the specific case. Inflicting more harm than the upper threshold would be a case of disproportionate surveillance since the surveilled would only be liable to suffer from less intrusion.

Naturally, it is difficult to identify a precise upper limit between the maximum level of intrusion caused by a surveillance activity (and other potential costs of the surveillance activity) and the benefits obtained by the specific activity. The challenges of narrow proportionality, especially the fact that only an upper limit for the degree of intrusion is relevant, mirrors existing logic within police, intelligence, and security authorities that dictate application of the *least intrusive means* in the course of achieving an important end (Kleinig 2009). I claim that narrow proportionality seems plausible as a proportionality principle in cases of surveillance of liable individuals. This is also the dominant view in the literature on surveillance ethics and reflects the views of both Macnish (2015) and Nathan (2017).

2) *Unintentional but foreseeable intrusion upon liable individuals*

An example of the second case would be a case in which another individual who is involved in the online child pornography case is unintentionally wronged by the covert police work addressed at the primary target. The person is involved in the distribution, and the police is aware of his affiliation, but he is not the primary target of the police action. This person would then be liable to suffer from

defensive intrusion, and the intrusion would be inflicted upon him unintentionally but foreseeably. The narrow proportionality principle likewise seems plausible in the second type of cases for the same reasons presented in the example above.

3) Intentional intrusion upon non-liable individuals

An example of the third case is a bit trickier to identify. When are surveillance measures initiated intentionally against individuals who are not potentially liable? Suppose a surveillance program is initiated to acquire information concerning relatives and friends of an identified threatener. The surveilled individuals would then be law abiding citizens, who are simply related to a liable person. Suppose furthermore that a government authority initiates targeted surveillance of the close relatives and friends, for example, by collecting personal details about them from telephone and internet providers in order to learn more about their behavior and their relations with the threatener and ultimately to know more about the threatener. In this case, the surveillance would be directed at non-liable agents, who are intentionally intruded upon. The surveillance of the relatives is, so to speak, used as a means of acquiring information that might be useful for mitigating a specific threat but that is not directed at a liable individual. The third case should apparently be assessed in accordance with the notion of wide proportionality and a lesser evil justification. Below, I question further whether wide proportionality would be the correct way of assessing the proportionality of cases in which intrusion is intentionally applied to non-liable individuals and not as a result of collateral intrusion.

4) Unintentional but foreseeable intrusion upon non-liable individuals

An example of the fourth type of proportionality requirement would, for example, be a case in which the relatives of a primary target of covert policing are wronged by the police action. It could be a case in which personal friendships or relationships are developed between an undercover police officer and relatives and friends of a liable threatener and in which the undercover police officer obtains substantial information concerning the lives of the relatives. These pieces of information are, however, irrelevant, and it is thus foreseeable but unintended that the officer obtains these pieces of information in order to play the 'undercover' role and not reveal the operation. In such cases, the main target of the intruding actions will presumably be liable, yet the relatives and others who are affected will be non-liable to the intrusion, which is unintentionally but foreseeably inflicted upon them.

This case would according to table 1 be assessed in accordance with wide proportionality, in which the pertinent factor in terms of proportionality would be that the intrusion upon the non-labile would cause the avoidance of a greater intrusion upon non-labile individuals. The intrusion into the lives of the relatives in the specific case of surveillance constitutes so-called collateral intrusion since the main purpose of the surveillance activity is to acquire information concerning the liable individual. As a result, the more or less direct application of McMahan's understanding of wide proportionality seems intuitively reasonable in these types of cases.

Differentiating between foreseeable and intentional intrusion upon non-labile individuals

I now return to the third type of case and emphasize the relevance of distinguishing between foreseeing and intending the intrusion entailed in a specific surveillance activity affecting non-labile individuals. Like Macnish, Nathan distinguishes between harms imposed upon the liable individuals and harms inflicted as a result of so-called 'collateral intrusion' of covert police work (Nathan 2017). Nathan, for example, specifies:

Even where there are no grounds for suspicion of involvement in wrongdoing, or even positive evidence of innocence, it is still possible that deception and manipulation are useful and justified, but this justification will need to cross the high bar of the kind that is analogous to a *justification of harming an innocent bystander* for the sake of some very significant benefits (Nathan 2017, pp. 13-14 – my emphasis).

Nathan thus opens up for the moral permissibility of harming bystanders as collateral intrusion, though neither Nathan nor Macnish distinguish between *the unintended but foreseeable* intrusion upon non-labile individuals and *the intentional* intrusion upon non-labile individuals in the course of a surveillance activity. Their sole focus is on the unintentional but foreseeable harming of innocent bystanders in the course of some activity serving a significant good.⁵ This is, however, an important distinction to draw since some means of surveillance are characterized by intruding intentionally upon non-labile individuals (e.g. some types of mass surveillance). In order to illustrate my claim that there is a distinction between intending and barely foreseeing intrusion upon the non-labile individuals, consider the following two examples:

Example 1: Casting a wide net I

⁵ Macnish argues that when a surveillance activity intrudes innocent by-standers the relevant moral assessment would be to adhere to the Doctrine of Double Effect (2015), yet he does not specify further how this should be done.

Imagine a surveillance program that is initiated in order to acquire information about the online behavior of an identified threatener, who we have reasonable belief to be engaged in planning an act of terrorism. These pieces of information will (among other means) help guide the authority when initiating counteractions against the planned aggression. In order to obtain the relevant pieces of information, the authority also collects information containing personal details of 10,000 individuals who are not involved in the criminal act. The personal lives of the 10,000 individuals are thus intruded upon as a consequence of collecting information on the threatener, but the collected pieces of information on the non-labile individuals are of no interest to the authority and are regarded as irrelevant after identifying the relevant pieces of information on the liable individuals. It is simply the case that surveiller cannot distinguish in advance between liable and non-labile individuals when collecting the pieces of information.

Example 2: Casting a wide net II

Imagine then a different surveillance program, which is initiated in order to acquire information on risk factors leading to radicalization and subsequently risks of conducting an act of terrorism. The program is designed to collect electronic information on the online behavior of 10,000 individuals, including communication via social media platforms, without any prior knowledge on whether the affected individuals are in fact considered radicalized. In other words, such a surveillance program would systematically collect personal details with the aim of gaining an impression about some of the factors potentially leading to radicalized behavior. In this case, intrusion is intentionally inflicted upon non-labile individuals in order to potentially prevent some future crimes.

The overall level of intrusion upon the non-labile individuals would apparently be equal in these two cases. Yet all else being equal, the second case is more difficult to morally justify than the first case. The main reason is that example 2 entails the intentional intrusion upon the personal lives of the non-labile individuals whereas example 1 entails the intentional intrusion upon known threateners in which intrusion upon the non-labile constitutes an unintentional but foreseeable side effect of the surveillance program.

The next question is whether wide proportionality would be an appropriate means of assessing cases of surveillance in which non-labile agents are wronged. McMahan (2009a, p. 21) specifies that narrow proportionality usually concerns harms that are inflicted intentionally, whereas ‘issues of wide

proportionality usually concern harms inflicted foreseeable but unintentional'. Thus, as already stated, as far as wide proportionality is concerned, interventions are proportionate 'if the expected good would outweigh the unintended harms to civilians' (McMahan 2009a, p. 20). Yet if wide proportionality would apply to cases of intentional surveillance of non-labile individuals, then the specific surveillance should enable the avoidance of greater intrusions upon the non-labile than if they were not surveilled. As argued above, this lesser evil justification seems intuitively appealing in cases in which the intrusion upon the non-labile is unintended but foreseeable. The relevant question is, however, whether the same proportion of benefits to harms is proportionate if the intrusions are intentionally inflicted upon innocent individuals? Intuitively, my answer would be no since the intended intrusion upon the non-labile is a more demanding claim than the foreseeable but unintended intrusion.

The difference between these two cases of casting a wide net could be described as the same kind of difference appearing in the paradigmatic case concerning the tactical bomber versus the terrorist bomber from discussions on 'The Doctrine of Double Effect' in the case of just warfare, where it is applied to explain the moral permissibility of inflicting harms (often death or injury) upon non-labile individuals in order to achieve a good end (McMahan 2009a, pp. 173, 232; 2009b, p. 346). The terrorist bomber will detonate a bomb to intentionally inflict harm upon non-labile individuals, whereas the strategic bomber aims to bomb a munitions factory, which produces weapons for the enemy, and thereby end the war. However, in so doing, the tactical bomber will harm both liable and non-labile individuals. If the doctrine is broadened to include not just harms but also harmless wrongs, it would thus intuitively apply to cases of surveillance. It seems that the first case should be assessed more harshly than the latter in terms of proportionality (according to JWT, the terrorist case would be morally impermissible, whereas the latter case would be morally permissible if the side effects were unavoidable and proportionate). The same logic would apparently hold for the two examples of 'casting a wide net'. The intention of the authorizing body when initiating the surveillance measure is thus relevant for the permissibility and understanding of proportionality in the specific context of surveillance. Intentional intrusion upon the non-labile could then be deemed impermissible with reference to the distinction between the tactical and the non-tactical bomber. However, the degree of harm in warfare and self-defense relative to the much less severe degree of intrusion of surveillance might render a different result more plausible.

McMahan appeals to a commonsense notion that a stricter principle of proportionality should count in such cases concerning warfare simply because the harms are inflicted intentionally. He does not, however, elaborate upon this apparent logic (McMahan 2009a, p. 22). Following JWT, cases of intentional harms of non-labile individuals would not meet to the criteria of discrimination of targets since only so-called combatants are liable targets. Yet most scholars working with JWT, McMahan included, are not absolutists in the sense that they would never permit intentional harms of non-labile individuals.⁶ Thus, for example, Michael Walzer would allow such actions in cases of ‘supreme emergency’ with reference to the claim that ‘it is a matter of how much worse the effect of obeying it [the requirements of discrimination] would be than the effect of violating it’ (McMahan 2009a, p. 22). Transferring this viewpoint to the domain of surveillance, it is indeed doubtful whether this type of preventive surveillance would meet the conditions of a case characterized by ‘supreme emergency’ since such surveillance activities are most often characterized by obtaining pieces of information that are ‘nice to know’ (potentially useful) rather than are urgent or emergent. Would this then mean that preventive surveillance measures are *per se* disproportionate, indiscriminate, and morally impermissible? This would be an appealing consequence of applying the concepts and distinctions from the Just War Theory and from the philosophical literature on self-defense to the case of surveillance, yet let us consider a different conclusion.

A stricter notion of wide proportionality?

A good place to start when elaborating upon the initial concept of wide proportionality could be to restrict the use of such surveillance measures by introducing a concept of *wide proportionality+* to cases in which the intrusion is intentionally inflicted upon non-labile individuals. Such a principle of proportionality could then be phrased as follows: Suppose that, in the course of eliminating a particular threat, it would be proportionate to unintentionally (but foreseeably) impose a particular level of intrusion upon non-labile individuals (wide proportionality). In this case, the threat to be eliminated would have to be of a much more serious nature for the intentional imposition of intrusion upon non-labile individuals to be proportionate (wide proportionality+).

The introduction of a stricter wide proportionality principle in cases in which non-labile individuals are intentionally intruded upon will require further specification in terms of, for example, what a

⁶ This is indeed a highly contested topic within the just war theory. One of the reasons why this is controversial is that allowing the intentionally harming of non-labile individuals brings JWT disturbingly close to terrorism (MCMahan 2009a).

‘much more severe threat means’. A more severe threat could, for example, be defined by the degree of the potential harms expected to be imposed if a specific and manifested threat is not avoided. Following this understanding of a threat, the principle would result in the fact that some of the existing mass surveillance programs, which are characterized by wide collections of personal electronic pieces of information in the name of general crime and terrorism prevention, would be disproportionate. This is because such means do not lead to the avoidance of a specific threat in the majority of cases, and the threat is of such a general and potential nature that it cannot be characterized as being ‘serious in nature’ since it has, for instance, not yet manifested itself (e.g. there is a general fear that a terrorist attack will occur but no knowledge of a specific threat). Such programs mainly serve to provide potentially useful information if a specific threat appears and the pieces of information are collected ‘just in case they are needed’. In such cases (e.g. in the case of the so-called PRISM and TEMPORA programs presented in the introduction), the surveillance activities would be disproportionate, and the intentional surveillance of the non-liable individuals would be morally impermissible.

Above, I have argued that a stricter notion of wide proportionality is necessary to guide cases in which non-liable individuals are intentionally intruded upon by means of surveillance. It is also necessary to consider the expected beneficiaries of the intended intrusion in order to further specify a wide proportionality+ principle. A relevant question to ask is thus: who benefits from the intentional intrusion upon the non-liable, for instance in the case of a specific mass surveillance program?

Are all non-liable individuals expected to be beneficiaries of the mass surveillance program, and if so, would this then affect the degree of morally permissible intrusion inflicted by the program? In contrast, would the wide proportionality principle look different if other individuals besides the affected non-liable individuals are the expected beneficiaries? It is beyond the scope of this paper to discuss this further, and I will leave this discussion to future work on the ethics of surveillance.

Concluding remarks

The existing literature on ethical surveillance mainly address cases of surveillance entailing intentional intrusion upon liable individuals (Macnish 2014, 2015; Nathan 2017). As a result, surveillance of non-liable individuals and the distinction between *foreseeing* and *intending* intrusion upon non-liable individuals in the course of surveilling have not previously received specific attention. This focus is, however, increasingly important, given that a precautionary principle is

generally considered an important instrument with regard to *anticipation* of criminal acts before they are committed as well as *preventive surveillance* of non-labile actors (e.g. some cases of mass surveillance) (McCulloch & Wilson 2015; Innes & Sheptycki 2004).

In this article, I have addressed the principle of proportionality as it is applied in the context of ethical surveillance. I have mainly focused on specifying which conditions should be present in order to be liable to suffer from the intentional intrusion entailed in specific surveillance measures. I have transferred to surveillance McMahan's distinction between narrow and wide proportionality, and I have argued that his notion of narrow proportionality applies more or less directly to cases of surveillance, including surveillance of liable individuals. I have additionally argued that McMahan's wide proportionality principle applies to cases in which intrusion is inflicted unintentionally but foreseeably upon non-labile individuals – understood as cases of collateral intrusion. However, I have argued that surveillance, which is characterized by intentional intrusion upon those who are non-labile, should be assessed with a stricter understanding of wide proportionality.

I have provided some more nuances to this claim by arguing that the harms potentially avoided in such cases should count much more than in a balance between lesser evils entailed in McMahan's wide proportionality principle from the context of warfare. My claim is that the threat should be more serious for the surveillance activity to be proportionate. Thus, the fact that the intrusion associated with such surveillance measures are inflicted intentionally upon non-labile individuals renders necessary a *wide proportionality+* principle, which can account for cases in which non-labile individuals are (more trivially) affected by surveillance in the course of avoiding very serious and specific threats.

References

Bellaby RW (2012) What's the Harm? The Ethics of Intelligence Collection. *Intelligence and National Security* 27/1: 93-117.

D'Amico DJ (2015) Knowledge Problems and Proportionality. *Criminal Justice Ethics* 34/2: 131-155.

[OMITTED]

Frowe H (2011) *The Ethics of War and Peace*. London/New York: Routledge.

Home Office (2010) Covert Human Intelligence Sources Code of Practice. Pursuant to section 71 of the Regulation of Investigatory Powers Act 2000. The Stationery Office. https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/514059/code-practice-human-intel.pdf.

Hurka T (2005) Proportionality in the Morality of War. *Philosophy and Public Affairs* 33/1: 34-66.

Innes M, Sheptycki JWE (2004) From Detection to Disruption: Intelligence and the Changing Logic of Police Crime Control in the United Kingdom. *International Criminal Justice Review* 14: 1–24.

Kleinig J (2009) The Ethical Perils of Knowledge Acquisition. *Criminal Justice Ethics* 28/2: 201-222.

Lyon D (2014) Surveillance, Snowden, and Big Data: Capacities, consequences, critique. *Big Data & Society*. doi: <https://doi.org/10.1177/2053951714541861>

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 Practise* 18: 529-548.

Marx GT (1998) Ethics for the New Surveillance. *The Information Society* 14/3: 171-185.

McCulloch J, Wilson D (2016) Pre-crime. Pre-emption, precaution and the future, Oxon/New York: Routledge.

McMahan J (2009a) *Killing in Wars*. Oxford: Oxford University Press.

McMahan J (2009b) Intention, Permissibility, Terrorism, and War. *Philosophical Perspectives*, 23: 345-372.

McMahan J (2005) Self-Defense and Culpability. *Law and Philosophy* 24: 751-774.

Nathan C (2017) Liability to Deception and Manipulation: The Ethics of Undercover Policing. *Journal of Applied Philosophy* 34/3: 370-388.

O'Donovan, O (2003) *The just war revisited*. Cambridge: Cambridge University Press.

Stahl T (2016) Indiscriminate mass surveillance and the public sphere. *Ethics of Information Technology* 18: 33-39.

Stoddart E (2014) Challenging 'Just Surveillance Theory': A Response to Kevin Macnish's 'Just Surveillance? Towards a Normative Theory of Surveillance'. *Surveillance & Society* 12/1: 158-163.

Uniacke S (2011) Proportionality and self-defense. *Law and Philosophy* 30/3: 253-272.