



Delft University of Technology

Rethinking the Just Intelligence Theory of National Security Intelligence Collection and Analysis

The Principles of Discrimination, Necessity, Proportionality and Reciprocity

Miller, Seumas

DOI

[10.1080/02691728.2020.1855484](https://doi.org/10.1080/02691728.2020.1855484)

Publication date

2021

Document Version

Final published version

Published in

Social Epistemology

Citation (APA)

Miller, S. (2021). Rethinking the Just Intelligence Theory of National Security Intelligence Collection and Analysis: The Principles of Discrimination, Necessity, Proportionality and Reciprocity. *Social Epistemology*, 35(3), 211-231. <https://doi.org/10.1080/02691728.2020.1855484>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.



Rethinking the Just Intelligence Theory of National Security Intelligence Collection and Analysis: The Principles of Discrimination, Necessity, Proportionality and Reciprocity

Seumas Miller

To cite this article: Seumas Miller (2021) Rethinking the Just Intelligence Theory of National Security Intelligence Collection and Analysis: The Principles of Discrimination, Necessity, Proportionality and Reciprocity, *Social Epistemology*, 35:3, 211-231, DOI: [10.1080/02691728.2020.1855484](https://doi.org/10.1080/02691728.2020.1855484)

To link to this article: <https://doi.org/10.1080/02691728.2020.1855484>



© 2021 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.



Published online: 14 Feb 2021.



[Submit your article to this journal](#)



Article views: 1640



[View related articles](#)



[View Crossmark data](#)



Citing articles: 3 [View citing articles](#)

Rethinking the Just Intelligence Theory of National Security Intelligence Collection and Analysis: The Principles of Discrimination, Necessity, Proportionality and Reciprocity

Seumas Miller^{a,b,c}

^aAustralian Graduate School in Policing and Security Studies, Charles Sturt University, Canberra, Australia; ^b4TU Centre for Ethics and Technology, TU Delft, The Hague, Netherlands; ^cOxford Uehiro Centre for Practical Ethics, University of Oxford, Oxford, UK

ABSTRACT

In this article, it is argued that the constitutive principles of Just War Theory and the *jus ad bellum*/*jus in bello* duality do not transfer all that well to national security intelligence activity. Accordingly, while Just Intelligence Theory – comprising *jus ad intelligentiam* and *jus in intelligentia* – provides a useful starting point in the construction of a normative framework for national security intelligence activities, ultimately it is found to be wanting in a number of important respects. For instance, while war ought to be a last resort, intelligence collection and analysis ought to be a first resort. In addition, analyses are offered of the key principles of discrimination, necessity and proportionality, and it is shown in general terms how they apply, or ought to apply, to national security intelligence activity. Importantly, the principle of necessity has been given a novel analysis according to which it is in reality a set of different principles, depending on the institutional setting in which it is being used. Moreover, the analysis reveals that as typically used it consists (in part) of a means/end principle of rationality and one or other versions of a principle of harm minimisation. Finally, it is argued that there is a normative principle governing espionage, in particular, that is not a constitutive principle of Just War Theory; this is a principle of reciprocity.

KEYWORDS

National security intelligence; intelligence ethics; principle of necessity; principle of proportionality; principle of reciprocity; principle of discrimination

1. Introduction

In this article, it is argued that the constitutive principles of Just War Theory and the *jus ad bellum*/*jus in bello* duality do not transfer all that well to national security intelligence activity, in part because of the essentially epistemic character of intelligence activity. Accordingly, while so-called Just Intelligence Theory – comprising *jus ad intelligentiam* and *jus in intelligentia* – provides a useful starting point in the construction of a normative framework for national security intelligence activities, ultimately it is found to be wanting in a number of important respects. For instance, while war ought to be a last resort, intelligence collection and analysis ought to be a first resort. On the other hand, some constitutive principles of Just War Theory are, appropriately revised, applicable to national security intelligence activity, notwithstanding the essentially epistemic character of intelligence activity. Specifically, analyses are offered of the key principles of discrimination, necessity and proportionality? Importantly, the principle of necessity has been given a novel analysis according to which it is in reality a set of different principles, depending on the institutional setting in which it is being used. Moreover, the analysis reveals that as typically used it consists (in part) of a means/end principle of rationality and one or other versions of a principle of harm minimisation. In addition, it is

shown in general terms how the principles of discrimination, necessity and proportionality apply, or ought to apply, to national security intelligence activity. In doing so, a distinction is invoked between the macro level (e.g. the establishment of bulk databases by security agencies for national security purposes) and the micro level (e.g. the conduct of a specific operation utilising data from a bulk database). Finally, it is argued that there is a normative principle governing espionage, in particular, that is not a constitutive principle of Just War Theory; this is a principle of reciprocity. This principle is elaborated and it is argued that compliance with it is both possible and desirable.

2. National Security Intelligence

The definition of intelligence in criminal justice, national security and related contexts has proved problematic.¹ One reason for this pertains to what is being defined; there are multiple definienda. Intelligence can refer to an ability, e.g. an intelligence officer has the ability to collect information, or it can refer to an action that is an exercise of that ability, e.g. an intelligence officer has the ability to collect information, or it can refer to the content or product of that action, e.g. the information collected. Moreover, in the context of intelligence agencies, what is being defined can be the capacity of the agency itself (or its action), as opposed to the ability of one or more of its officers (or their actions); or it can be the institutional processes (e.g. the so-called intelligence cycle) and/or information storage systems (e.g. database of information), as opposed to the reasoning processes and/or knowledge in the 'heads' of officers. Further, definitions can be purely descriptive (defining what intelligence *is*) or they can be normative (defining what intelligence *ought to be*). Consider, for instance, this illustrative definition in the national security literature: "Intelligence is a corporate capability to forecast change in time to do something about it. The capability involves foresight and insight, and is intended to identify impending change which may be positive, representing opportunity, or negative, representing threat" (Breakspear 2013). This definition refers to an *ability*, as opposed to an action, of an *organization*, as opposed to an intelligence officer per se ("corporate capability"); and this definition has a normative feature ('capability involves foresight and insight') that might exclude poor intelligence as intelligence (since it did not manifest insight). Accordingly, this definition stands in contrast with, for instance, definitions that focus on intelligence as information or knowledge, i.e. on the content or product of intelligence activity.

Another reason that the definition of intelligence has proved problematic pertains to the notion of knowledge and, therefore, to those definitions that focus on the content or product of intelligence activity. Knowledge is by definition true. However, much intelligence is obviously false. Therefore, intelligence, it is argued, cannot be defined in terms of knowledge. However, others have made the counter-argument that knowledge is what is ultimately *aimed at*,² even if success is only ever partial and sometimes there is complete failure, e.g. as in the case of Saddam Hussein's non-existent WMDs.

Given the complexity of the issues, and of the associated debate, a complete and correct definition, and defense of it, cannot be offered here. We must limit our ambitions. We do so in part by narrowing our definitional focus to the content or product of the intelligence activity of individual intelligence officers, e.g. knowledge. Moreover, it will be assumed that the truth of the content or product of this activity is typically what is or, at least, ought to be aimed at (even if failure is often the reality). Accordingly, other things being equal, false information is poor intelligence. Note, however, that truth is not the only criterion of good intelligence since some information, although correct, is nevertheless useless, e.g. because it is not sufficiently timely. Moreover, some intelligence might not, strictly speaking, be assessable as true or false but rather only as, for instance, correct or incorrect, accurate or inaccurate, or (in relation to future events) have a high, medium or low degree of probability. Further, a number of additional features of intelligence are identified below.

Many items of what are referred to as intelligence, whether human intelligence or electronic intelligence, can be thought of as, (i) being expressed or expressible in a language and, therefore, communicable, (ii) being epistemic (or knowledge focused) – and, as such, even in their raw form

capable of being true or false, correct or incorrect, probable or improbable, evidence-based or not, etc. – (iii) stored somewhere, e.g. in an investigator's notebook, in a security organization's data bank, and (iv) elements or fragments of some larger network or structure in terms of which an given single such item can become intelligible. Intelligence, therefore, can be thought of as in large part consisting of statements, or material expressible as statements, stored in some information storage system.

Information and intelligence are closely related concepts (see, for instance, Miller and Gordon 2014; Kahn 2001). Both information and intelligence can be thought of as statements stored in some information system. Moreover, both information and intelligence, as these terms are used here, are either true or false (although, perhaps unverifiable), but neither is necessarily true. This is, of course, true of disinformation masquerading as bona fide intelligence and also of much 'raw' intelligence. But it is also true of intelligence that has been subjected to analysis and is well evidenced. Finally, neither information nor intelligence necessarily has a good and decisive justification. Accordingly, neither information nor intelligence is necessarily knowledge. On the other hand, a piece of information or of intelligence might be true and might have a good and decisive justification, i.e. *some* information and *some* intelligence is knowledge.

Notwithstanding that information and intelligence are closely related concepts, they are not the same thing; specifically, intelligence is information but information is not necessarily intelligence. In the context of criminal investigations (Miller and Gordon 2014, Chap. 2), intelligence is information that is utilized, actually or potentially: (i) to facilitate the outcome of specific criminal investigations (e.g. to identify and apprehend the Yorkshire Ripper); (ii) in the day-to-day tasking and deployment of an organization's sub-units (e.g. a local police station or local area command unit) in response to crimes of particular types in particular locations (e.g. violence at closing time outside certain late-night venues on weekends) (tactical intelligence); or (iii) in the long-term planning of the organization's deployment of resources and its strategic response to crime trends (e.g. home-grown terrorism) (strategic intelligence).

Notice that intelligence, whether it be criminal intelligence, market intelligence, or military intelligence, is defined relative to some institutional purpose or function and, relatedly, it is typically the result of the joint action of multiple intelligence personnel; specifically, the result of joint epistemic action (Miller 2015, 2018a). Police seek information for purposes such as investigation, arrest and the prosecution of criminals. That the information is sought in order to realize these kinds of purposes is what makes the information in question criminal intelligence. Likewise, members of military intelligence agencies seek information for the purposes of winning wars, deterring military aggressors, and so on, and that the information is sought for these purposes makes it military intelligence. Again, corporations seek information on other corporations for purposes of gaining commercial advantage in a context market-based competition, and that the information is sought for these purposes makes it market intelligence.

In short, in this article a teleological (purpose-based) or functional account of intelligence is assumed, i.e. intelligence is, by definition, information or data (typically expressible as a statement or, more likely, structured set of statements) that is acquired for various institutional purposes. Moreover, intelligence is institutionally relative in that it is relative to the purposes of some institution (Miller 2010, 2016a, Chap. 3). So military intelligence is a different category of intelligence from criminal intelligence because military institutions have a somewhat different institutional purpose than police organizations. That said, one and the same item of knowledge might be both an item of military intelligence and an item of criminal intelligence, e.g. the knowledge that Abū Bakr al-Baghdadi is engaged in war crimes. The fact that military intelligence and criminal intelligence are different categories of intelligence does not mean that a given piece of information might not belong to both categories.

What of so-called national security intelligence? National security intelligence is sometimes collected, stored, analysed and disseminated, as actionable intelligence, by military organisations, sometimes by police organisations, but paradigmatically by intelligence agencies the institutional purpose of which is internal and/or external national security, e.g. the CIA, NSA, GCHQ, MI5, MI6, Mossad, RAW,

ASIO, etc. Accordingly, what makes information or other data collected by these agencies national security intelligence is that these agencies collect, analyse and disseminate this information in the service of national security – national security being the primary institutional purpose of these agencies. This immediately raises the vexed question as to what national security is; after all, the content of the term ‘national security’ is notoriously ill-defined, indeterminate, shifting, open-ended and contestable. For instance, the US National Intelligence Strategy has as one of its purposes to promote American prosperity.³ Importantly, national security should not simply be understood as national interest (contrary to the view expressed in the US National Intelligence Strategy, 2017), since the latter notion is very permissive and could license all manner of individual and collective rights violations. For instance, it might be in the national interest of a nation-state to increase its territory by invading a neighbouring nation-state, e.g. Germany’s invasion Poland in the initial stages of World War 2 or a possible future Russian invasion of Ukraine, or to enslave a population, or otherwise engage in widespread, serious rights violations, to increase its own wealth, e.g. US’ enslavement of members of African tribes in the antebellum or the Chinese incarceration in the last few years of hundreds of thousands of Uighurs in oil- and resources-rich Xinjiang. However, let us assume that national security intelligence is intelligence pertaining to serious internal or external threats to the nation-state itself, or to one of its fundamental political, military or criminal justice institutions, and that these threats might emanate from state or non-state actors, e.g. terrorist groups.

Note that this account of national security intelligence would include military intelligence and some, but by no means all, criminal intelligence. Please also note that intelligence collection, analysis and dissemination is not an end-in-itself; rather the end-point of the intelligence process is actionable intelligence, i.e. intelligence to relevant decision-makers that is a means to non-epistemic, notably kinetic, action (or, as the case may be, intentional *in-action (omission)*). A final point to be noted is that many intelligence agencies focused on the epistemic activity of national security intelligence have traditionally also engaged in kinetic activity, e.g. sabotage (including by way of cyber-attacks, such as the Stuxnet virus), interference in elections, funding dissidents, cross border kidnappings, arming secessionists, assassinations (see Perry 2016; Miller 2016b). Moreover, some of their intelligence collection, analysis and dissemination is in the service of these latter covert operations. Accordingly, not all national security intelligence collected, analysed and disseminated by national security intelligence agencies is in the service of kinetic activities (often devised by policymakers) to be conducted by other institutions, e.g. by the military or the police and/or under the direction of the intelligence agencies’ political masters.

Armed with these working definitions of intelligence in general and national security intelligence in particular, let us now turn to the normative theory or model of intelligence of interest here, namely, what can be referred to as the Just Intelligence Theory (JIT) since it is based on the Just War Theory (JWT) of war.⁴ JWT is, of course, a normative theory, i.e. a theory consisting of a set of conditions under which waging war is (allegedly) *morally* justified. Likewise, JIT is a theory consisting of a set of conditions under which national security intelligence collection, analysis and dissemination is (allegedly) *morally* justified. Note that in both cases what is on offer is not a normative institutional theory (Miller 2010, 2016a, chap. 3), i.e. a theory of military forces or of intelligence agencies, but rather a theory of the constitutive activities of such institutions. This is important since, for instance, a normative theory of military institutions might hold that its principal institutional purpose was to deter any would-be aggressor whereas JWT is a theory of the activity of waging war, i.e. the conditions under which it is morally justified to wage war and the principles governing the conduct of war once underway. Let us begin by briefly outlining (serviceable versions of) JWT and JIT.

3. Just War Theory and Just Intelligence Theory

Roughly speaking, according to JWT, the armed forces of a collective⁵ political entity, A (e.g. a nation-state), are morally justified in waging war against the armed forces of another collective political entity B (e.g., a nation-state), if and only if:

- (1) A's purpose in waging war is collective self-defense against B's military aggression⁶;
- (2) A uses lethal force only to the end of bringing about the cessation of B's aggression;
- (3) A's armed forces are waging war under a legitimate political authority;
- (4) There is no alternative means of defense against B and so A wages war as a last resort;
- (5) A has a reasonable chance of winning the war against B;
- (6) It is probable that if A wages war against B the consequences, all things considered, will be better than if A does not;
- (7) A only deliberately uses lethal force against morally legitimate targets, i.e. against B's combatants (identifiable by virtue of their uniforms and the fact that they bear their arms openly) but not B's civilians ((i) principle of discrimination), an extent of violence that is necessary to the end of winning the constitutive battles of the war and, ultimately, the war itself ((ii) principle of military necessity), and the violence is not disproportionate ((iii) principle of proportionality).

Notice that condition (7) is essentially the so-called *jus in bello* of JWT according to which combatants on both the just and the unjust side are (at least on the traditional view⁷) moral equals. (Conditions (1)–(6) constitute the so-called *jus ad bellum*). However, I am assuming that a war that failed to comply with the *jus in bello* would not be a just war by the lights of JWT.

Let us now outline an account of JIT that parallels the above account of JWT. Note that JIT is intended by some of its advocates to apply not simply to national security intelligence but also to criminal intelligence that might not be national security intelligence, and perhaps also to counter-intelligence and covert action. For simplicity, this article is confined to its main focus: national security intelligence collection, analysis and dissemination.

Roughly speaking, according to JIT,⁸ the national security intelligence agencies of a collective political entity, A (a nation-state), are morally justified in collecting, analysing and disseminating intelligence in relation to a collective⁹ political actor, B (e.g. a nation-state), if and only if:

- (1*) A's purpose in undertaking these intelligence activities is to enable the removal, typically by A's security agencies (including A's military forces), of any national security threat posed by B;
- (2*) A undertakes these intelligence activities only for the ultimate purpose of enabling the removal of the threat to its national security posed by B;
- (3*) A's intelligence agencies are collecting, analysing and disseminating the intelligence in question under a legitimate political authority;
- (4*) There is no alternative means of enabling the removal of the national security threat posed by B and so A's intelligence activities are a last resort;
- (5*) A has a reasonable chance of enabling the removal of the national security threat posed by B;
- (6*) It is probable that if A undertakes these intelligence activities directed at B the consequences, all things considered, will be better than if A does not;
- (7*) A only deliberately collects, analyses and disseminates intelligence concerning morally legitimate targets, ((i) principle of discrimination*), an extent of intelligence activity that is necessary to the end of removing the security threat posed by B (principle of intelligence necessity*), and the harm caused is not disproportionate (principle of proportionality*).

Notice that condition (7*) is essentially the so-called *jus in intelligentia* of JIT. (Conditions (1*)–(6*) constitute the so-called *jus ad intelligentiam*). However, by analogy with JWT, it is assumed here that intelligence activities that failed to comply with the *jus in intelligentia* would not be just by the lights of JIT.

The analogy between JWT and JIT is problematic in a number of respects. First, the largely kinetic activity of waging war ('killing people and breaking things') is very different from the essentially epistemic activity of intelligence collection, analysis and dissemination.¹⁰ Importantly, epistemic

activity is in and of itself less harmful than killing and maiming – notwithstanding the fact that its consequences may well be just as harmful – and, therefore, the moral constraints on it more permissive. Moreover, since the end of epistemic activity is (roughly speaking) knowledge and, therefore as a matter of logic, one embarks on an epistemic project from a position of ignorance – and with a set of questions to be answered, e.g. Is there a threat? What is the nature of the threat? – the content of the end is in an important sense, and by definition, unknown. Accordingly, any prior moral assessment of the contemplated epistemic action is necessarily radically incomplete, since it depends in large part on the moral costs attaching to the realization of the epistemic end, i.e. of being in possession of the answers to the questions sought – something which is, to reiterate, by definition unknown.¹¹ Consider the following scenario. An intelligence unit is seeking to find out whether a terrorist group has the capability of producing a radiological ‘dirty’ bomb (a weapon of mass destruction or WMD) and, therefore, ‘turns’ a disaffected member of the terrorist group (Informant) and tasks him to get this information; an extremely dangerous job which puts the life of Informant at serious risk. Informant discovers that the terrorist group has no such capability and has no plans to develop it. Unfortunately, however, informant’s intelligence collection activities have raised suspicions and shortly thereafter he is murdered by the terrorist group. The upshot is that the terrorist group does not have the WMD capability – and is now known by the intelligence unit not to have it – and Informant is dead. Moreover, if the intelligence unit had not tasked Informant then the terrorist group would likewise not have a WMD capability – albeit the intelligence unit would not know this – but Informant would be alive. Therefore, as it turns out, there was no need for Informant to die since there was no risk of a WMD attack. However, the intelligence unit could not have *known* that the terrorist unit had no WMD capability and that, therefore, there was no need for Informant to put his life at risk and, indeed, to die, since it was precisely the end of Informant’s epistemic task to find out if the terrorist group had, or did not have, a WMD capability; that is, the intelligence unit’s prior moral assessment was not only radically incomplete, it was *necessarily* radically incomplete.

Second, the morally acceptable purposes of war (for the most part consisting of defence against military aggression) constitute a much narrower, determinate and less open ended (e.g. wars but not national security concerns in general have a start and finish date) set of national security concerns than those that might legitimately motivate intelligence activities (see Omand and Phythian (2018, 78)).

Third, unlike war, intelligence is not the last resort; rather it is typically the first resort (see Omand and Phythian 2018, 90; Bellaby 2014, 28). Indeed, it is logically prior to the decision whether to wage war. Whether to wage war, apply economic sanctions or merely use diplomacy are decisions predicated on intelligence. In short, the relationship between any putative JIM and JWT is not that of theories that mirror one another by virtue of the structurally similar activities (national security intelligence collection/analysis/dissemination and waging war, respectively) that they are theories of, but rather that of theories of dissimilar activities but activities that, nevertheless, stand in (roughly speaking) the relationship of knowledge to action; kinetic action presupposes epistemic action since the decision to perform a kinetic action (or not to do so) presupposes knowledge with respect to the why, how, what, when, where, who, etc., of the kinetic action in question (and its alternatives).

Fourth, intelligence activities often do not in themselves remove national security threats. Rather, being epistemic activities, they are the precursor to the removal of these threats by kinetic means, e.g. war, arrest. Accordingly, principles (1*), (2*), (4*) and (5*) are very different principles from (1), (2), (4) and (5), (respectively). The latter ones involve kinetic action that can directly remove the threat; the former are epistemic actions that are (usually) part of the precursor means provided to kinetic actors to remove the threat (supposing it turns out to exist).

Given the nature of epistemic action and its typically indirect relationship to significantly harmful outcomes, the moral principles governing it – and especially the application of these principles – is likely to be quite different.

Obviously, principles (1) and (1*) are very different by virtue of the different purposes they express and their direct versus indirect (respectively) relationship to the realization of those purposes. Moreover, JIT's analogue to (4), i.e. (4*), seems to be false. On the other hand, JWT and JIT do share some general principles, namely, (3) and (3*) (legitimate authority), (5) and (5*) (principle of reasonable prospects of success), and (6) and (6*) (principle of consequentialism) – albeit (5*) has an indirect relationship to the success of its ultimate purpose. Of these (5) and (5*) are instances of a general and rather abstract principle of rationality, and (6) and (6*) are instances of a general and rather abstract principle of morality. As such, these three sets of general principles do not do much by way of establishing a close relationship between JWT and JIT. For instance, they might also be regarded as principles to be applied to a morally committed market actor setting up his or her franchise business in a competitive market.

What of principles (2) and (2*), and of principles (7) and (7*) (respectively)? Principles (2) and (2*) are analogous in so far as one should not engage in harmful activity morally justified by one purpose (e.g. self-defence or national security, respectively) for a different purpose which does not morally justify the activity (e.g. territorial expansion or national interest, respectively). That said, as we have seen, national security is a much wider, less determinate and more open-ended notion than self-defence against a military aggressor. Accordingly, intelligence gathered for one national security purpose can often be very easily be argued to be useable for another national security purpose.

What of the analogy between (7) and (7*), i.e. the so-called *jus in bello* and *jus in intelligentia*, and, therefore, the principles of discrimination, necessity and proportionality? This is the subject of the following section. Here some general preliminary points are made.

We need to distinguish between national security intelligence activities directed at one's own citizens, e.g. home-grown terrorists or revolutionaries (internal threats to national security), and national security intelligence activities directed at foreign powers, including foreign spies (external threats to national security).

The targets of internal national security criminal intelligence activities such as terrorists are, presumably, criminals or suspected criminals. As such, they fall under what might be termed the law enforcement model under which law enforcement agencies such as the FBI conduct their activities, including their national security intelligence activities. However, the law enforcement model does not sit well with JWT nor, evidently, with JIT.

Unlike combatants fighting an unjust war but operating under *jus in bello* principles, criminals do not identify themselves as such and do not seek to differentiate themselves from 'ordinary civilians' by way of dress or other means; quite the reverse, they seek to conceal their real identity. Hence, the need for law enforcement agencies to utilize notions of probable cause and reasonable suspicion. Given their status as criminals, criminals are not regarded as moral equals by law enforcement officials (whether intelligence officers or others). In this respect, criminals are quite different from combatants waging an unjust war in accordance with JWT's principle of the moral equality of combatants. For instance, if apprehended, criminals are not treated as prisoners of war but are charged, tried and, if found guilty, sentenced and punished.

Foreign spies are also engaged in unlawful activity, as far as the nation-state being spied upon is concerned, and as such are also treated as criminals. On the other hand, there is a moral symmetry of sorts between foreign spies and combatants in so far as states spying on one another each regard their own foreign spies, not as criminals, but as intelligence personnel engaged in a morally legitimate national security activity – although (unlike combatants) doing so deceptively and using a cover, e.g. as diplomatic staff officially serving in embassies. However, this moral symmetry does not seem to hold between home-grown terrorists or revolutionaries engaged in an unjust insurrection, on the one hand, and (say) undercover operatives gathering intelligence on them with a view to prosecuting them for crimes against the legitimate (let us assume) state in question.

Bearing these general points in mind, let us now turn to a more detailed analysis of the principles of the *jus in intelligentia*; specifically, the principles of discrimination, necessity and proportionality (as they might apply to national security intelligence collection and analysis).

4. Principles of Discrimination, Necessity and Proportionality in National Security Intelligence

Consistent with relevant liberal democratic regulatory regimes, harmful internal national security intelligence activities (at least) are, or ought to be, conducted in accordance with the principles of necessity and proportionality. Moreover, these principles are typically applied in conjunction with a third principle, namely, the principle of discrimination. (See L.C. Green 1993. In the context of Just War Theory see Walzer 1997.) The principle of discrimination is involved in so far as it is generally assumed that innocent persons ought not to be harmed or have their rights violated. Accordingly, it would be one thing for police to intercept and access the metadata and content of the phone calls and emails of a known terrorist on an ongoing basis for intelligence purposes, and quite another for this to be done to a citizen known to be innocent of any crime, e.g. on the off-chance that some useful intelligence might be picked up (Miller and Gordon 2014). Surveillance of the terrorist would in this instance be a *morally justified infringement* of the right to privacy (Miller and Gordon 2014; Miller 2009), whereas surveillance of the innocent citizen would evidently be a *violation* of the right to privacy. Accordingly, the principle of discrimination ought to be applied to national security intelligence activities. However, its application in these activities is somewhat different from its application in kinetic military and policing contexts. Speaking generally, its application to national security intelligence activities is far more permissive.

Importantly, unlike the targets of, for instance, military combatants (Miller 2016a), the targets of intelligence activities can sometimes be innocent civilians, e.g. deliberately and deceptively gaining information about a terrorist from the terrorist's innocent relative might be morally justified whereas deliberately killing the terrorist's innocent relative would certainly not be. Moreover, intelligence activities ultimately aimed at identifying terrorists and thwarting acts of terrorism often now involve the application of machine learning techniques to bulk databases that consist in the main of the communication and other data of innocent civilians – indeed, frequently innocent fellow citizens; i.e. the data of innocent civilians are deliberately collected and accessed (or, at least, filtered and accessed). It can be argued that while the bulk data of these innocent persons is 'read' by a machine or, perhaps, 'seen' by human eyes in an anonymized form, it is for the most part not seen (in the appropriate privacy infringing sense). Of course, the particular data items that result from the application of the machine learning process are de-anonymized and, ultimately, seen by human eyes; however, the argument might continue, such data meets the standard of reasonable suspicion already applicable to intelligence gathering/investigation by law enforcement agencies, and does so by virtue of being the result of that very process. Whatever the merits of this argument as a justification for the application of machine learning techniques to bulk databases by way of mitigating the degree and extent of intrusion into the privacy of innocent citizens (see 2018 ; Miller 2018b); nevertheless, this intrusion into the privacy of innocent civilians is deliberately done, albeit as a means to an end. As such, it is not analogous to the principle of discrimination as it applies to the use of lethal force by combatants in war; combatants, to reiterate, are not permitted to *deliberately* kill innocent civilians, even as a means to some further legitimate end. The reason for this difference between the principle of discrimination as it applies to intelligence activities and as it applied to the use of lethal force reflects the much greater moral significance that attaches to deliberately overriding an innocent person's right to life than attaches to deliberately overriding their right to privacy. This difference in significance in turn reflects, indeed in large part is derived from, the much greater moral weight that attaches to life than to privacy. Hence, there is a (more or less) absolute legal prohibition on deliberately killing the innocent (even in wartime), but not on deliberately overriding their privacy (even in peacetime).

A final point regarding the principle of discrimination as it applies to intelligence activities pertains to differences between internal and external national security threats. Intelligence activities directed at external threats to national security, e.g. threats posed by foreign powers, are much less constrained than those directed at internal threats, e.g. home-grown terrorists; indeed, in war there are few, if any, constraints on intelligence activity. Arguably, this is not how it should be; after all, the

innocent citizens of enemy authoritarian states do have moral rights, including privacy rights (whatever their legal rights might be or, more likely, not be). However, it does seem that given the purpose of the intelligence activities in question is national security, and governments and their security agencies have special *partialist* duties in respect of their own citizens (Miller 2016a chap. 3), it is perhaps to be expected that the principle of discrimination and, for that matter, the principles of necessity and proportionality, might justifiably be applied in a more permissive manner externally than they are internally.

The principle of necessity that is of interest in this article is applied in circumstances in which harm (including, for ease of exposition, in the sense of infringement of a moral right¹²) is typically caused by members of security agencies, including members of intelligence agencies, to those targeted in their operations and caused by virtue of the inherently harmful method used, e.g. surveillance (infringement of the right to privacy), arrest (infringement of the right to freedom of movement), and use of lethal force. The principle of necessity is typically illustrated by recourse to a standard situation of personal self-defence in which the defender (Defender) has a choice between these two (effective) means (harmful methods) to preserve her life – killing or disarming her attacker (Attacker). It is generally held that she ought to disarm Attacker since it is *not necessary* for her to kill Attacker.¹³ By parity of reasoning, if an intelligence officer (Officer) has a choice between two (effective but harmful) means to collect information on a suspected terrorist (Suspect) – collecting metadata from Suspect's phone or intrusively surveilling Suspect by means of miniature cameras and listening devices placed in his home – Officer ought to simply rely on metadata since it is *not necessary* to engage in intrusive surveillance. However, from the mere fact that one of two available means is not necessary to realize some end it does not follow that it ought not to be chosen. After all, *ex hypothesi* neither of the two available means is a necessary means to achieve the end in question (unless the other means is not chosen) and it would be irrational (other things being equal) not to choose any of the available means to one's ends.

What is going on here? Clearly, the idea is that the less harmful means morally ought to be chosen and the harm in question is harm to the target, e.g. to Attacker or to Suspect. In our self-defence example, Defender ought to choose to disarm Attacker rather than kill him because disarming Attacker is the less harmful means to achieve the end of preserving Defender's life. Likewise, in our intelligence example, Officer ought to choose to collect Suspect's metadata but not his communicative content since this is the less harmful means – being less of an infringement of Suspect's privacy¹⁴ – to achieve the end of acquiring the desired information.

However, *qua means to the end of preserving Defender's life*, disarming Attacker is no better than killing Attacker. Indeed, disarming Attacker might be a worse choice *qua means* to that end since, for example, it might be less effective (the chances of failure are greater) or less efficient (the effort required is greater) than killing him. Again, *qua means* to the end of acquiring information, accessing Suspect's metadata is no better (we have assumed) than intrusively surveilling Suspect. Indeed, if anything, merely accessing Suspect's phone metadata is presumably a worse choice *qua means* to that end since, for example, it might be less effective (some relevant information might only be found in the content of his conversations). Nevertheless, it might continue to be insisted that the less harmful means morally ought to be chosen. Why so? Evidently, there is another end in play here. The end in question is the moral end to minimize harm to persons from which can be derived the moral principle to minimize harm to targets. Moreover, given the possibility of so-called collateral damage, there is also the derived principle of minimising harm to bystanders. Further, given the possibility of harm being done by targets or third parties to the user of harmful methods, we can derive a third principle from the general moral end of minimising harm to persons; the principle of minimising harm to the users of the harmful methods in question, such as Defender, police officers and intelligence officers (who will in each case henceforth be referred to as operators, i.e. users of a harmful method). We return to this threefold distinction between minimising harm to targets, bystanders and operators below.

As we have seen, in each of our scenarios there is a harmful means to an end. However, in each of our scenarios, there are two *conceptually independent* ends, e.g. in our personal self-defence scenario the end of preserving Defender's life and the end of minimising harm to Attacker. While conceptually independent, these two ends can come into conflict under some circumstances, e.g., "in a self-defence" scenario in which Defender must kill Attacker or be killed by Attacker. In the intelligence collection scenario, there is the end of collecting the desired intelligence and – as in the self-defence scenario – the end of minimizing harm to the target, i.e. in the intelligence collection scenario to those from whom intelligence is being collected (and the harm minimization in question is that of minimizing the degree or extent of infringement of the right to privacy). In each of the two scenarios, the two ends are conceptually independent. This is obvious from the fact that one could have as an end to defend oneself and yet not have as an end to minimize harm to others.¹⁵ This may well be so if, for instance, Defender decides to kill Attacker to preserve Defender's life, notwithstanding that Defender could easily have chosen the equally effective means of disarming Attacker. Likewise, the conceptual independence of the two ends in play in the intelligence collection scenario is obvious since one could have as an end to collect the desired intelligence and yet not have as an end to minimize harm to targets (by minimizing the degree and extent of the infringement of privacy rights). This may well be so if, for instance, Officer decided to intrusively surveil the target, notwithstanding that Officer could easily have chosen the equally effective means (let us now assume) of relying exclusively on the metadata.

So we need to distinguish between two conceptually independent ends in the application of the necessity principle in both self-defence scenarios and intelligence collection scenarios; the end definitive of the activity in question (e.g. preserving one's life, acquiring national security intelligence) and the end of minimizing harm to the target. However, in discussing applications of the principle of necessity we also need to stress the differences between the ends definitive of kinetic activity, such as interpersonal self-defence, and the ends definitive of intelligence activity. Clearly, the ends definitive of intelligence activities (e.g. knowledge of terrorists and their plans) and those definitive of kinetic activities (e.g. arrest of suspects, destruction of enemy military forces) are different and, accordingly, the end implicit in the application of the principle of necessity in an intelligence collection context will be different from that implicit in kinetic activity. Moreover, the (harmful) means to realize the ends definitive of intelligence activities and to realize the specific ends implicit in an application of the necessity principle in an intelligence collection context will also be different. So drawing analogies between military, law enforcement and intelligence activities in respect of the principle of necessity relies on moving to a high level of abstraction. Further, in intelligence activities, the notion of necessity in play is very often a permissive one. For instance, intelligence activities that utilize bulk data are often not necessary, strictly speaking; rather they are the most effective – and perhaps least resource intensive – means to a national security end (see [section 5](#) below).

The description of the principle of necessity that has been provided thus far omits a key feature of this principle as it operates in the contexts in question, namely, that its constitutive end – that of preserving Defender's life or, in the other scenario, of collecting desired intelligence – is a moral end. Accordingly, the principle of necessity is not after all *merely* a principle of rationality. Rather, it has, in the contexts in question, a moral loading by virtue of the moral quality of the ends in play; accordingly it is, after all, a moral principle, at least in these contexts. What morally justifies Defender's act of killing Culpable Attacker is in part the moral weight attaching to Defender's life, the preservation of which is Defender's end. It is also true that preserving Defender's life (the end) needs to be weighed against the loss of Attacker's life (the harmful means). Specifically, the moral principle of proportionality applies in relation to weighing morally significant ends and means. Moreover, an analogous argument operates in the case of intelligence collection that is the means to a legitimate national security end; Officer's acquisition of the national security intelligence (the end) needs to be weighed against the infringement of Suspect's right to privacy (the harmful means).

Now neither the negative moral weight of the harmful method nor the positive moral weight of the end of preserving Defender's life or of collecting national security intelligence is constitutive of the necessity principle *qua principle of necessity*. The necessity principle *qua principle of necessity* pertains to the necessity of a means (Defender killing Attacker or Officer infringing Suspect's right to privacy) to an end (preserving Defender's life or Officer acquiring national security intelligence, respectively). But evidently an action is a means to an end irrespective of the moral quality of either the action or the end that it serves.¹⁶ Moreover, an action that is a necessary means to some end is a *necessary* means irrespective of the moral quality of the action or its end. Nevertheless, the principle of necessity, or rather principles of necessity, as they apply in the inter-personal, military, law enforcement and intelligence contexts in question are moral principles by virtue of their implicit reference in each case both to a harmful means and a moral (indeed, morally worthy, let us assume) end. So each of these principles of necessity is a moral principle at the core of which is a means-end principle of rationality; and each of these principles of necessity *qua principle of necessity* is merely a principle of rationality.

The upshot of the discussion thus far is that the principles of necessity in the institutional contexts in question (military, law enforcement and intelligence) are moral principles each of which has at its core a means/end principle of rationality. The means/end principle of rationality states that (other things being equal) one ought to choose the means to one's ends and, if there is a necessary means, then one ought to choose it. However, each principle of necessity is different from the others by virtue of the different moral ends in play in these various institutional contexts (as well as, typically, the different harmful means used in the service of these ends). Thus, in military contexts, the principle of necessity is referred to as the principle of military necessity and implicitly refers to the moral end of winning the war, battle or other military engagement in question. On the other hand, in a law enforcement context, the principle of necessity might refer to the moral end of preserving the life of the defender, be that a citizen or the officer him/herself. By contrast with these essentially kinetic ends, the principle of necessity as it applies in intelligence activities has an epistemic (and morally significant) end; to acquire national security information.

Moreover, as we saw above, these principles of necessity are applied in circumstances in which harmful methods are being used, notably by members of security agencies (the operators) against their targets but also in interpersonal contexts by, for instance, Defender against their Attacker. Accordingly, as also mentioned, there is a harm minimization principle associated with each of these principles of necessity. The harm minimization principle is a principle of morality and it states that one ought to minimize harm to persons and, therefore, if there are two (or more) means to a given end then (other things being equal) one ought to choose the least harmful. Of course, to reiterate, the ends specified in the family of related principles of necessity applicable in different security contexts vary. As will become clear, this point is also relevant to our understanding of the harm minimization principle and, therefore, it is important to reflect on its scope in the interpersonal self-defense, kinetic military, kinetic law enforcement and national security intelligence contexts of its application. In each of these contexts we need to keep in mind the threefold distinction made above with respect to the harm minimisation principle, namely, harm to targets, harm to bystanders and harm to operators (i.e. to reiterate, Defender, police officer, intelligence officer or combatant using the harmful method in question). It is only the *harm to target* version of the harm minimization principle that has been identified above as implicated in the versions of the necessity principle thus far discussed. This is to be expected, given that the circumstances in question all involve the use of harmful methods by the operator (e.g. intelligence official) against the target.¹⁷

Let us consider further the various harm minimization principles implicated in different versions of the principle of necessity. The harm minimization principle in play in the application of the principle of necessity in internal national security intelligence operations is focused on minimizing harm to the criminal or suspected criminal (as well as to innocent third parties, i.e. (typically) fellow citizens of the intelligence officers), i.e. it is the harm to targets version of the principle (and also the harm to bystanders version). As such, it is not focused on minimizing harm to the intelligence officer,

him/herself, i.e. the one conducting the surveillance or collecting and analysing the data (the operator). Indeed, there is typically no need to go beyond the minimize harm to targets principle (and, perhaps, minimize harm to bystanders principle) and invoke the minimize harm to operators principle in such contexts unless, of course, the intelligence officer is an informant or undercover operative. Consider, by contrast, kinetic law enforcement, i.e. situations involving the use of lethal force (in particular) by police officers. In this latter case, police officers ought to minimise harm to themselves (as well as to offenders and to innocent members of the public). As such, the minimise harm to operators principle seems applicable and, indeed, is implicated in the application of the necessity principle in these contexts. The principle that police officers should not use lethal force against a suspect unless it is necessary implies, firstly, that they ought not do so if there are non-lethal methods available (minimise harm to the target) or if using lethal force would put the lives of bystanders at serious risk (minimise harm to bystanders) but, secondly, that they may do so if to do otherwise would put their own lives at serious risk (minimise harm to operators).¹⁸

Now consider kinetic military operations and, in particular, the harm minimization principle implicated in the principle of military necessity. Here there is also a dis-analogy with respect to harm minimization as it applies to intelligence activities. The harm minimization end implicated in the principle of military necessity applies when war is already underway and, important to our concerns here, might *not* have a focus on minimizing harm to enemy combatants in, for instance, an ambush or firefight. Indeed, quite the reverse; in such circumstances it might be focused on *maximizing* harm to enemy combatants, e.g. if the most effective military strategy is to degrade enemy forces. So the harm minimization principle implicated in the principle of military necessity is *not* the principle of minimizing harm to targets; rather it is the principle of minimizing harm to bystanders (as opposed to targets). In this respect, the harm minimization principle implicated in the principle of military necessity is also dis-analogous to the harm minimization principle implicated in the principle of necessity as it applies in law enforcement, since the police ought to minimize harm to criminals, and as it applies in internal national security intelligence activity, since intelligence officers ought to minimize harm to targets in their domestic intelligence activities. Related points can be made with respect to the operation of the principle of proportionality (about which more below). If it was necessary to slaughter most of a much larger enemy force in order to win a battle, this would not necessarily be regarded as a disproportionate measure, assuming innocent civilians' lives and the lives of one's own combatants – belonging to the much smaller armed force – were not put at significant risk. One reason for this difference is evidently that in internal (at least) national security operations, as in law enforcement more generally, there is a presumption of innocence and, therefore, a strong presumption in favor of refraining from harming criminals, i.e. 'criminals' are, legally speaking, only suspects prior to conviction.

What of the principle of proportionality? In the light of our above analysis of the necessity principle in terms of a means/end principle of rationality at its core, a moral end, and an implied principle of harm minimization in its contexts of use, what is to be morally weighed in applications of the proportionality principle? As already mentioned, the application of the principle of necessity in the above-described scenarios requires moral weight to be attached to the relevant moral ends in question (e.g. the life of Defender, the intelligence collected, winning the battle), relative to the harm caused by the harmful means, e.g. the death of Attacker, infringement of privacy, foreseen but unintended death of some innocent civilians. Moreover, obviously the ends in question need to be (actually or prospectively) realized if they are to be given moral weight; so the use of the harmful methods needs to be successful. As we have seen, the harm done as a result of the use of harmful methods is (at least potentially) harm to the target, to bystanders and/or to the operator (e.g. the user of the harmful method). However, as we have also seen, not all of these harms are present in some settings, e.g. there might be no harm done (or in prospect of being done) to intelligence officers engaged in surveillance; and not all of these harms, even if present, are necessarily given negative moral weight in the application of the proportionality principle in a given institutional setting, e.g. harm done to enemy combatants might not be. It follows from this that there are

different versions of the proportionality principle in play in different institutional settings (as we saw was the case with the necessity principle). However, speaking generally, the proportionality principle ascribes positive moral weight to the ends realized (the ends constitutive of self-defence, law enforcement, intelligence activity or military action) and negative moral weight to the harms caused by the use of harmful means. This is in part simply a reflection of a general principle of rationality. After all, other things being equal, it is surely irrational to choose to perform an action as a means to one's end in circumstances in which the benefit (or good) of the end is outweighed by the cost (including moral cost) of the means. To this extent, the application of the principle of necessity implies the application of a principle of proportionality with respect to means and end. However, the various proportionality principles give negative moral weight to some harms that are not part of the means to the end, e.g. harms to bystanders. Indeed, in theatres of war – settings in which the principle of military necessity applies – the accompanying principle of proportionality is largely focused on avoiding disproportionate harm to bystanders, i.e. to innocent civilians (Walzer 1997).

Note that the necessity principle and all relevant harm minimization principles might be complied with and yet the proportionality principle not be complied with. Thus, in order to prevent a pickpocket escaping, it might be necessary for a police officer to shoot the pickpocket in the leg and this might be the least harmful means available, e.g. the only alternative to allowing the pickpocket to escape would be for the officer to shoot the pickpocket dead. However, even this action of shooting the pickpocket in the leg would be disproportionate. So compliance with the necessity principle and all relevant harm minimization principles is not a sufficient condition for compliance with the proportionality principle. Conversely, it is possible to comply with the proportionality principle and yet fail to comply with relevant harm minimization principles. It might not be disproportionate for an intelligence officer to intrusively surveil a known terrorist, notwithstanding that an equally effective, less intrusive means was available. However, to do so would be a failure to comply with the principle of minimizing harm to targets and, given the latter is (as we saw above) implicated in the relevant necessity principle, it would also be a failure to comply with the necessity principle.

Please also note that the proportionality principle presupposes the application of the principle of discrimination. Thus, the *deliberate* bombing of these innocent civilians (as opposed to their foreseeable but unintended deaths) is ruled out by the principle of discrimination, irrespective of any proportionality considerations. As we saw above, by contrast with the bombing of innocent civilians, the principle of discrimination as it applies to intelligence collection is far more permissive. Does the principle of proportionality, nevertheless, presuppose the principle of discrimination in national security intelligence activities?

Harm in terms of privacy infringements is easy or, at least, easier to justify in the case of suspects – and certainly known offenders, e.g. known terrorists – than in the case of innocent citizens. Hence, the application of the principle of proportionality presupposes the principle of discrimination in play; it might be disproportionate to collect intelligence by means of an intrusive method from a person believed to be innocent of any serious crime but not disproportionate if the target were a known terrorist.

5. Jus Ad Intelligentiam, Jus in Intelligentia and the Macrolevel/Microlevel Distinction

Whereas there is a relatively clear cut distinction between the decision to wage war and decisions made in the actual conduct of war and, therefore, between the jus ad bellum and the jus in bello, matters are somewhat different when it comes to national security intelligence collection/analysis/dissemination and, therefore, the (alleged) distinction between the jus ad intelligentiam and the jus in intelligentia. National security intelligence activity (i.e. collection, analysis and dissemination) is a continuous, ongoing (indeed, cyclical – hence, the so-called intelligence cycle) activity in relation to threats and enemies that come and go; unlike war, it has no determinate end state, such as the

cessation of hostilities, that is being aimed at (perhaps understood in terms of winning the war).¹⁹ Moreover, as we saw above, national security intelligence activity does not mirror kinetic activity such as waging war but rather stands to it in the general relationship of knowledge to action, i.e. as its logical precursor. Accordingly, the existence and application of the alleged *jus ad intelligentiam*/*jus in intelligentia* dualism – as opposed to the existence and application of particular constitutive principles, e.g. the principles of necessity and proportionality – is, to say the least, open to question; there seems to be a lack of conceptual fit between the phenomenon (national security intelligence activity) and this dual theoretical framework.

That said, national security intelligence activity exists at both a micro and a macro level (as long as this distinction is understood as being a fairly loose one). The micro level is the level of specific operations. This is the level which has been the focus of most of the discussion in [section 4](#) above; a level at which the principles of discrimination, necessity and proportionality are manifestly applied. But national security intelligence activity also exists at a macro level, and this has implications for the application of the principles of necessity and of proportionality in particular. Consider in this connection national security intelligence bulk data collection.

At the micro level, the application of the principles of necessity, proportionality and discrimination is on specific intelligence operations directed at particular targets, e.g. collecting information concerning the associates of a suspected terrorist. Questions to be addressed include the following ones. Is intrusive surveillance necessary and proportionate? Would the less intrusive collection of metadata to determine callers/persons called be sufficient? What of the macro level?

Key ethical issues at the macro level pertain to the necessity and proportionality of the establishment and general uses of the bulk databases themselves.²⁰ In his influential UK report David Anderson (Anderson, David (Independent Reviewer of UK Terrorism Legislation) 2016) distinguishes between bulk interception, bulk acquisition, bulk equipment interference (e.g. hacking into computerised device and copying material), and bulk personal datasets (e.g. electoral roles, passport database, driving license database, national insurance numbers, passenger name records from flights (PNRs)). He also distinguishes between databases held by the security agencies and their accessing of databases held by other agencies, e.g. private sector firms. His concern was with the former. Regarding the necessity of establishing and utilizing these databases, Anderson said: 'The bulk powers play an important part in identifying, understanding and averting threats in Great Britain, Northern Ireland and further afield. Where alternative methods exist, they are often less effective, more dangerous, more resource-intensive, more intrusive or slower (Anderson, David (Independent Reviewer of UK Terrorism Legislation) 2016 chap. 5–8)'. Clearly given, for instance, the existence of alternative methods that are merely more resource-intensive, this is a relatively weak and, therefore, permissive notion of necessity.²¹

Anderson did not address the question of proportionality in his report. In order to do so, we would need to distinguish between the proportionality of establishing a particular database for national security purposes, as opposed to accessing and analysing (for national security purposes) an existing database created for a purpose other than national security. Moreover, the weight to be accorded to the right to privacy in any such application of the principle of proportionality is a complex matter, not the least because of the close (conceptual?) relationship between privacy and other fundamental rights, such as the right to individual autonomy in the context of the liberal democratic concern not to allow individual autonomy vis a vis the state to be compromised. Evidently, the application of the principle of discrimination at this macro level is problematic in so far as the databases in question necessarily contain the data of citizens innocent of any national security breach; indeed, most of the data in many of the databases in question pertain to innocent citizens. Nor is this problem necessarily entirely resolved, even if it is considerably mitigated, by virtue of, for instance (and as mentioned above), the anonymised form in which the personal data in these databases exists in the collection and filtering, etc. phases of the national security intelligence process.

There is also the question of the relationship of the micro level to the macro level from the perspective of the application of the principles of discrimination, necessity and proportionality. For instance, successful microlevel counterterrorism (CT) operations that rely on bulk data might be aggregated so as to justify the existence and accessing of bulk databases for national security purposes in terms of the principle of necessity (as per Anderson's report mentioned above). Again, taken in aggregate the nature and extent of the infringements of privacy of innocent citizens resulting from the accessing of databases of personal information might be held not to be disproportionate to the aggregated outcomes of successful CT operations that relied on the accessing in question. Note that compliance with the principles of necessity and proportionality at the macro level does not mutually entail compliance with these principles at the micro level, i.e. does not mutually entail compliance with these principles on each and every specific intelligence collection operation. This is in part because microlevel operations might be only ultimately justified in terms of their contribution to macrolevel outcomes. For instance, spreading the intelligence gathering net wide and over a long period of time might enable the joining of dots on a terrorist network and its activities, notwithstanding that the accessing of the personal data of a given person who was not a suspect, but merely thought (falsely, as it turns out, let us assume) to be a potential associate, might not – *considered on its own* – be justified by the principles of discrimination, necessity or proportionality. Conversely, a micro-level operation might be justified in its own terms without recourse to its contribution to macro-level outcomes.

The principle of proportionality needs to take into account not only the somewhat vague character of the end of national security (definitive, as we saw above, of the principle of necessity) and the obstacles faced by intelligence officers, e.g. high-level encryption, but also potential future harms arising from national security intelligence activities and, in particular, from the utilization of bulk data. To reiterate, privacy concerns in this area are somewhat mitigated by the fact that the bulk data collected and analysed is typically in an anonymised form (e.g. by means of machine learning techniques), and, therefore, arguably only the privacy rights of genuine suspects are infringed or, perhaps, seriously infringed (i.e. the individuals identified upon completion of the analysis). However, these harms, such as the weakening of individual autonomy vis a vis the state arising from extensive privacy infringements by intelligence agencies, and a diminution in public trust (a collective good [Miller 2010]) as a consequence of the secret nature of national security intelligence activities, may be incremental, difficult to quantify and collective in character. Please also note that considered at the macro level, the harms in question are potentially various in terms of our above-mentioned taxonomy of harms. For instance, since intelligence officers are themselves citizens, their intelligence activities might turn out to be (indirectly and incrementally) a form of collective *self-harm*, given their membership of the collective harmed.

Accordingly, it can be difficult to know exactly where to draw the line between proportionate and disproportionate intelligence activities when it comes to the utilization of bulk data for national security purposes. Consider in this connection the potential utilization of integrated biometric and non-biometric databases. One prominent concern about the inadequacy of privacy protections is the potential for 'function creep', where the use of information taken for a particular purpose is used for other purposes for which consent was not obtained. The underlying concern in relation to 'function creep' is the one adumbrated above; namely, the threat to individual autonomy posed by comprehensive, integrated biometric and non-biometric databases utilized by governments and their security agencies in the service of ill-defined notions of necessity and national security and, at least potentially, without appropriate regulatory constraints and democratic accountability.

6. Espionage and the Principle of Reciprocity

Thus far, it has been argued that the constitutive principles of JWT and the *jus ad bellum/jus in bello* duality do not transfer all that well to national security intelligence activity. Accordingly, while JIT is a useful starting point in the construction of a normative framework for national security intelligence

activities, ultimately it has been found to be wanting in a number of important respects.²² Analyses of the principles of necessity and proportionality, and of their relationship to one another and to the principle of discrimination in their application to national security intelligence activity, have also been provided. As argued above, versions of these principles are applicable to intelligence activity. However, these versions turned out to be very different from those applicable to kinetic military and law enforcement contexts. It is now time to argue that there is a normative principle governing external national security intelligence activities (call it espionage), in particular, that is not a constitutive principle of JWT and, therefore, not of JIT; this is a principle of reciprocity (see Miller and Walsh 2016; Walsh and Miller 2016). Accordingly, this argument, if successful, is an important additional problem for JIT. Note that the principle of reciprocity does not provide the basis for an ethical framework to replace JIT. Rather it is one of a set of principles (others being, as argued above, discrimination, necessity and proportionality) that ought to give direction to national security intelligence activity. The task of providing an ethical framework is a large one to be left for another day.

Verizon and PRISM²³ have raised legitimate privacy concerns, both for US citizens and for foreigners, for example in relation to metadata collection and analysis. Regarding metadata collection and analysis in the context of domestic law enforcement, including internal national security intelligence collection and analysis, the solution, at least in general terms, is evidently at hand; extend the existing principles of probable cause (or, in many non-US jurisdictions, reasonable suspicion), and the existing relevant accountability requirements, for example, the system of judicial warrants.

However, some of these privacy concerns pertain only to foreign citizens. Consider the FISA (Foreign Intelligence Surveillance Act) Amendments Act of 2008. It mandates the monitoring of, and data gathering from, foreigners who are outside the US by the NSA. Moreover, data collected but found not to be relevant to the foreign intelligence gathering purpose of, say, counterterrorism is not allowed to be retained. Importantly, however, there is no probable cause (or reasonable suspicion) requirement unless the person in question is a US citizen.

This is problematic insofar as privacy is regarded as a *human* right and, therefore, a right of all persons, US citizens or not. Moreover, these inconsistencies between the treatment of US citizens and foreigners are perhaps even more acute or, at least obvious, when it comes to the infringement of the rights to privacy and, for that matter, confidentiality of non-US citizens in liberal democratic states allied with the US, for example EU citizens.

Intelligence-gathering, surveillance and so on of citizens by domestic law enforcement agencies is reasonably well defined and regulated, for example in accordance with probable cause/reasonable suspicion principles and requirements for warrants; hence, the feasibility of simply extending the law enforcement model to metadata collection within domestic jurisdictions. However, this domestic law enforcement model is too restrictive, and not practicable, in relation to external national security intelligence gathering from, for example, hostile foreign states during peacetime, let alone wartime.

The privacy rights of the members of the citizenry during wartime are curtailed under emergency powers; and the privacy and confidentiality rights of enemy citizens are almost entirely suspended. Military intelligence-gathering during war-time has few privacy constraints and, given what is at stake in all-out wars, such as World War II, this may well be justified. However, these are extreme circumstances and the suspension of privacy rights is only until the cessation of hostilities. Accordingly, this military model of intelligence-gathering is too permissive in relation to secret national security intelligence gathering from, for example, fellow liberal democracies during peacetime.

The above-mentioned intelligence gathering activities, including cyber-espionage, of the NSA do not fit neatly into JWT or, for that matter, the law enforcement model. At any rate the question arises as to whether some different moral principle(s) needs to be invoked in relation to espionage, in particular. Evidently, a principle not to be found in JWT or JIT needs to be invoked and, in particular, a principle of reciprocity.²⁴ Here we should distinguish between a retrospective and a prospective principle of reciprocity.

The retrospective principle of reciprocity takes its inspiration from the ancient prescription, ‘an eye for an eye and a tooth for a tooth’ and, therefore, from *lex talionis*. On this version of the principle, if one is unjustifiably attacked, or otherwise unjustifiably harmed or wronged, then one is morally entitled to respond in kind, irrespective of whether it is necessary for the specific purpose of, say, self-defence. On the other hand, one is not entitled to do more harm to an attacker than the attacker did to oneself, whether one does so by mounting a single more harmful counter-attack or by mounting a series of counter-attacks which in aggregate are more harmful. However, the prescription ‘an eye for an eye and a tooth for a tooth’ is too permissive; it would license reciprocal attacks on others for any purpose whatsoever, just so long as this attack was not more harmful than the one it was in response to. Accordingly, we need to place a restriction on the principle; a restriction with respect to the purposes it is to serve. More specifically, a morally acceptable version of this retrospective principle would justify nation-state, A, engaging in espionage against nation-state (or non-state actor), B, in circumstances in which B had engaged, or was engaging, in unjustifiable espionage on A, but only if A’s espionage was in the service of A’s morally justifiable political purposes, namely, national security (as opposed to, for example, national interest).

The prospective principle of reciprocity is a tit-for-tat principle in the service of bringing about a morally desirable future state of affairs. The state of affairs in question is an equilibrium state among nation-states; more specifically, a morally justifiable equilibrium under the rule of international law. This is not tit-for-tat in the service of the very general purpose of doing whatever is in one’s national interest, legitimate or otherwise (in the manner of rational choice theories); nor is it tit-for-tat measures short of war in the service of the narrow purpose of averting a future large-scale lethal attack which would constitute war (as might be justified under some extension of Just War Theory). Of course, in this equilibrium state of affairs there would be no espionage activities among participating nation-states, e.g. liberal democracies, or, at least, they would be few and far between. So this principle does not justify harmful actions in the manner of its sister retrospective principle; rather it has as its purpose to eliminate, or at least greatly reduce, harmful actions and, in this case, espionage and, thereby, move relevant nation-states into some form of a social contract. However, the equilibrium which is its *raison d’être* is at best a long-term goal; it is unlikely to be achieved anytime soon.

On the one hand, the US and its allies cannot be expected to defend their legitimate national interests with their hands tied behind their backs. So their recourse to espionage seems justified and the retrospective principle of reciprocity provides a specific moral justification for this. On the other hand, understood as a prospective tit-for-tat procedure in the service of bringing about a social contract, the principle of reciprocity requires the moral renovation of espionage, including cyber-espionage, as it is currently conducted. Second, a couple of suggestions: (i) the clustering of nation-states and; (ii) a demarcation between government and security personnel on the one hand, and ordinary citizens on the other.

Under existing arrangements, the US, the UK, Canada, Australia and New Zealand – the so-called ‘Five Eyes’ – share information gathered from other states. These nation-states are, so to speak, allies in espionage, notably cyber-espionage; for example, they share intelligence. They are the members of my first cluster. There are, of course, other liberal democratic states outside the Five Eyes, such as various EU countries, which have ‘shared core liberal democratic values’ with one another and with the Five Eyes and, specifically, a commitment to privacy rights. This is a second cluster.

The members of these two clusters each have privacy respecting laws and associated accountability measures in their domestic settings. However, they ought to make good on their claims to respect privacy rights as human rights by developing privacy-respecting protocols governing their intelligence-gathering activities in relation to one another. Of course, determining the precise content of such protocols is no easy matter given, for example, that there are often competing national interests in play, even between liberal democracies with shared values and many common political interests. But there does not appear to be any in-principle reason why such protocols could not be developed; and the fact that this might be difficult is no objection to attempting to do so. Moreover, since adherence to the protocols in question would consist, in so far as it is practicable, in

ensuring compliance with some of the standard moral principles protecting privacy and confidentiality rights already in place in liberal democracies (but not authoritarian states), such as probable cause or reasonable suspicion and use of judicial warrants, these two clusters would essentially consist of an extension of the law enforcement model to espionage conducted within and between these countries. Clearly, such an extension is unlikely, if not impossible, in the case of authoritarian states since these lack any commitment to privacy (and other) individual rights even in their domestic settings.

Further, such a process of clustering of liberal democratic states would be in accordance with the prospective principle of reciprocity; each of these nation-states would need to agree to, and actually comply with, the privacy respecting protocols in question but each might be deterred from not doing so by the tit-for-tat procedure of the prospective principle. This is, of course, not to say that those who agree to comply with the protocols/laws will *never* breach them. But this is true of protocols and laws in general.

What of authoritarian states known to be supporting international terrorism and/or engaging in hostile covert political operations, including espionage and cyber-espionage, for example China and North Korea?

In respect of authoritarian states of this kind, the retrospective principle of reciprocity reigns. Accordingly, there are few, if any constraints on intelligence-gathering and analysis, including cyber-espionage, if it is done in the service of a legitimate political interest such as national security.²⁵ Nevertheless, it is important to demarcate within such an authoritarian state between the government and its security agencies, on the one hand, and private citizens, on the other.²⁶ Notwithstanding the applicability of the retrospective reciprocity principle, the need to respect the privacy rights of private citizens in authoritarian states remains; perhaps all the more so given these rights (and, for that matter, human rights in general) are routinely violated by their own governments.

So a stringent principle of discrimination ought to govern espionage, including cyber-espionage, directed at authoritarian states. At the very least, the citizens of these states ought to be able to differentiate between morally justified infringements of the privacy and confidentiality rights of members of their government and its security agencies, on the one hand, and violations of their own privacy and confidentiality rights, on the other, and be justified in believing that whereas the former might be routine the latter are few and far between.

7. Conclusion

In this article, it has been argued that the constitutive principles of Just War Theory and, therefore, the *jus ad bellum*/*jus in bello* duality do not transfer all that well to national security intelligence activity. Accordingly, while Just Intelligence Theory – comprising *jus ad intelligentiam* and *jus in intelligentia* – provides a useful starting point for the construction of a normative framework for national security intelligence activities, ultimately it has been found to be wanting in a number of important respects. For instance, while war ought to be a last resort, intelligence collection and analysis ought to be a first resort. In addition, analyses of the key principles of discrimination, necessity and proportionality have been provided, and it has been shown in general terms how they apply, or ought to apply, to national security intelligence activity. Importantly, the principle of necessity has been given a novel analysis according to which it is in reality a set of different principles, depending on the institutional setting in which it is being used. Moreover, the analysis reveals that as typically used it consists (in part) of a means/end principle of rationality and one or other versions of a principle of harm minimisation. Finally, it has been argued that there is a normative principle governing espionage, in particular, that is not a constitutive principle of Just War Theory, namely, a principle of reciprocity.

Notes

1. For a useful recent theoretical contribution to this neglected topic see Pili (2019). For definitions offered in the national security literature, e.g. Ronn and Hoffding (2013), Stout and Warner (2018) and Marrin (2018).
2. See Ronn and Hoffding (2013) for this argument.
3. *National Security Strategy of the United States of America*, 2017, p. 4 'Second, we will promote American prosperity. We will rejuvenate the American economy for the benefit of American workers and companies.' <https://www.hsd.org/?view&did=806478>. See Miller (2010), Baldwin (1997) and Wallace (2018).
4. The literature on Just War Theory is vast. However, the most salient contemporary account is that of Michael Walzer (1997).
5. For an analysis of collective entities see Miller (2010).
6. War might also be justified in defense of A's allies or other defence against B's massive ongoing rights violations (notably genocide) of a non-ally of A. However, here I set these complications aside.
7. See Walzer (1997). For criticisms and reductionist accounts see Rodin (2002), McMahan (2009), Fabre (2013) and Frowe (2014). For an alternative 'institutionalist' perspective, see Miller (2016a).
8. For accounts of the so-called Just Intelligence Theory in which Just War Theory principles are applied to intelligence see Gendron (2005), Quinlan (2007), Bellaby (2014) and David Omand and Phythian (2018, chap. 3). For criticism of the Just Intelligence Theory, see Diderichsen and Ronn (2019).
9. There is question as to whether JIT is to be applied at the macro or micro levels (see section 5). Theoretically, if it parallels JWT it can only apply at the macro level. However, in practice, its advocates shift between the macro and the micro levels.
10. See Phythian in Omand and Phythian (2018, 84–86) for this kind of criticism of *ius ad intelligentiam*.
11. This point has implications for Bellaby's ladder of escalation (Bellaby 2014, 30–31) which seeks to calibrate justify levels of harm from harmful epistemic means relative to the justification available by the lights of the principles of JWT or, at least, JIT. Whereas one can determine the level of harm arising from the epistemic means, e.g. the degree of the privacy infringement, the problem is determining whether there are any benefits from the knowledge gained, as we can see in the case of torture. Accordingly, there is an epistemic problem in the application of principles of JWT/JIT in the kinds of examples in question.
12. In other contexts, it is important to distinguish harms from rights violations since, arguably, there can be rights violations without harm, e.g. a violation of a right to privacy which is never disclosed to the person whose right to privacy is violated.
13. There is a voluminous literature on the moral justification for killing in personal self-defense. See, for instance, Leverick (2009). Regarding the necessity principle, see, for instance, Lazar (2012). Regarding the proportionality principle see, for instance, Hurka (2005) and Uniacke (2011, 253–272). See also Miller (2016a).
14. There might also be resource considerations; intrusive surveillance is more resource intensive. But this is another matter.
15. And vice-versa, since one could have as an end to minimise harm to other but not have as an end to defend one's own life.
16. There might be some exceptions to this, e.g. the moral end of being treated respectfully.
17. Naturally, there might be scenarios in which the defender could minimize harm to him/herself by, for instance, using his arm to shield his head from a baseball bat wielding attacker, thereby, incurring a broken arm rather than a broken skull. However, such scenarios are not at issue here.
18. Or would put the lives of bystanders at risk (minimise harm to bystanders). The applicability of two or more of the harm minimisation principles given rises to important moral questions, if the principles come into conflict, e.g. should a police officer put his own life at risk to save an innocent bystander's life?
19. See Mark Phythian in Omand and Phythian (2018, 85) for this kind of point and David Omand (91–2) for a response to it.
20. David Omand recommends the distinction between laws and their application as being serviceable in the attempt to understand how *ius ad intelligentiam* and *ius in intelligentia* might relate to national security intelligence activities. See Omand and Phythian (2018, 99).
21. Assuming, of course, that the principle of necessity is what Anderson had in mind. But if he did not have the principle of necessity in mind what principle did he have in mind? See Macnish (2018, chap. 5) for an account of the ethical issues in this area.
22. This has also been argued in effect by Mark Phythian, if not David Omand, in their *Principled Spying* (2018), albeit the argumentative detail is somewhat different.
23. Verizon – the collection by the NSA of the metadata from the calls made within the US, and between the US and any foreign country, of millions of customers of Verizon and other telecommunication providers; PRISM – the agreements between NSA and various US-based internet companies (Google, Facebook, Skype etc.) to enable NSA to monitor the on-line communications of non-US citizens-based overseas.
24. Reciprocity-based principles are related to, but distinct from, consent-based principles. In relation to the latter applied to espionage, see Piaff and Tiel (2004).

25. There are important questions here concerning what counts as a legitimate purpose, particularly in the context of the blurring of the distinction between a political interest and an economic interest, for example China's cyber-theft operations. For reasons of space, I cannot pursue these here.
26. As an anonymous referee pointed out, the distinction between members of government and private citizens in authoritarian states may be hard to draw in practice, e.g. Is Huawei a private company or a part of the Chinese security apparatus? Here we need to distinguish between citizens and organisations (see Miller 2010 for analysis of organisations). Huawei is a company, not a private citizen. Nevertheless, managers/employees at Huawei and other citizens in authoritarian states can be required, indeed coerced, to engage in spying activities on behalf of their governments. When this happens they cease to be *merely* private citizens; they are now functioning as spies.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This article has been written as part of a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 670172).

Notes on contributor

Seumas Miller B1 is the author or co-author of 20 books, including *Moral Foundations of Social Institutions* (CUP, 2010), *Investigative Ethics (with I Gordon)* (Wiley-Blackwell, 2014), *Shooting to Kill: The Ethics of Police and Military Use of Lethal Force* (OUP, 2016) and *Institutional Corruption* (CUP, 2017), and over 200 academic articles.

References

- Anderson, D. (Independent Reviewer of UK Terrorism Legislation). 2016. "Report of the Bulk Powers Review". <https://terrorismlegislationreviewer.independent.gov.uk/wp-content/uploads/2016/08/Bulk-Powers-Review-final-report.pdf>
- Baldwin, D. A. 1997. "Concept of Security." *Revue of International Studies* 23 (1): 5–26.
- Bellaby, R. 2014. *Ethics of Intelligence*. Abingdon: Routledge.
- Breakspear, A. 2013. "A New Definition of Intelligence." *Intelligence and National Security* 28.
- Diderichsen, A., and K. V. Ronn. 2019. "Intelligence by Consent: On the Inadequacy of Just War Theory as a Framework for Intelligence Ethics." *Intelligence and National Security* 32 (4): 479–493.
- Fabre, C. 2013. *Cosmopolitan War*. Oxford: Oxford University Press.
- Frowe, H. 2014. *Defensive Killing*. Oxford: Oxford University Press.
- Gendron, A. 2005. "Just War, Just Intelligence: An Ethical Framework for Foreign Espionage." *International Journal of Intelligence and Counter-Intelligence* 18 (3): 398–434.
- Green, L. C. 1993. *The Contemporary Law of Armed Conflict*. Manchester, Canada: Manchester University Press.
- Hurka, T. 2005. "Proportionality in the Morality of War." *Philosophy and Public Affairs* 33 (1).
- Kahn, D. 2001. "An Historical Theory of Intelligence." *Intelligence and National Security* 16: 79–92.
- Lazar, S. 2012. "Necessity in Self-Defence and War." *Philosophy and Public Affairs* 40: 3–44.
- Leverick, F. 2009. *Killing in Self-Defence*. Oxford: Oxford University Press.
- Macnish, K. 2018. *The Ethics of Surveillance*. London: Routledge.
- Marrin, S. 2018. "Evaluating Intelligence Theories: The State of Play." *Intelligence and National Security* 33.
- McMahan, J. 2009. *Killing in War*. Oxford: Oxford University Press.
- Miller, S. 2009. *Terrorism and Counter-Terrorism: Ethics and Liberal Democracy*. Oxford: Blackwell.
- Miller, S. 2010. *The Moral Foundations of Social Institutions: A Study in Applied Philosophy*. New York: Cambridge University Press.
- Miller, S. 2015. "Joint Epistemic Action and Collective Moral Responsibility." *Social Epistemology* 29 (3): 280–302.
- Miller, S., and P. Walsh. 2016. "NSA, Snowden and the Ethics and Accountability of Intelligence Gathering." In *Ethics and the Future of Spying: Technology, Intelligence Collection and National Security*, edited by J. Galliot and W. Reed, 193–204. New York: Routledge.
- Miller, S. 2016a. *Shooting to Kill: The Ethics of Police and Military Use of Lethal Force*. New York: Oxford University Press.
- Miller, S. 2016b. "Cyber-attacks and 'Dirty Hands': Cyberwar, Cyber-Crimes or Covert Political Action?" In *Binary Bullets: The Ethics of Cyberwarfare*, edited by F. Allhoff, A. Henschke, and B. J. Strawser, 228–250. Oxford: Oxford University Press.

- Miller, S. 2018a. "Joint Epistemic Action: Some Applications." *Journal of Applied Philosophy* 35 (2): 300–318.
- Miller, S. 2018b. "Machine Learning, Ethics and Law." *Australian Journal of Information Systems* 22: 1–13.
- Miller, S., and I. Gordon. 2014. *Investigative Ethics: Ethics for Detectives and Criminal Investigator*. Hoboken, NJ: Wiley-Blackwell.
- Omand, D., and M. Phythian. 2018. *Principled Spying: The Ethics of Secret Intelligence*. Oxford: Oxford University Press.
- Perry, D. 2016. *Partly Cloudy: Ethics in War, Covert Action and Interrogation*. 2 ed. London: Rowman & Littlefield.
- Piaff, T., and J. R. Tiel. 2004. "Ethics of Espionage." *Journal of Military Ethics* 3 (1): 1–15.
- Pili, G. 2019. "Intelligence and Social Epistemology: Towards a Social Epistemological Theory of Intelligence." *Social Epistemology* 33: 6.
- Quinlan, M. 2007. "Just Intelligence: Prolegomena to an Ethical Theory." *Intelligence and National Security* 22 (1): 1–13.
- Rodin, D. 2002. *War and Self-Defense*. Oxford: Oxford University Press.
- Ronn, K. V., and S. Hoffding. 2013. "The Epistemic Status of Intelligence." *Intelligence and National Security* 28.
- Sorell, T. 2018. "Bulk Collection, Intrusion and Domination." In *Philosophy and Public Policy*, edited by A. Cohen, 39–60. Lanham MA: Rowman and Littlefield.
- Stout, and M. Warner. 2018. "Intelligence Is as Intelligence Does." *Intelligence and National Security* 33: 4.
- Uniacke, S. 2011. "Proportionality and Self-Defence." *Law and Philosophy* 30 (3): 253–272.
- Wallace, W. C. 2018. "National Security." In *The SAGE Encyclopedia of Surveillance, Security, and Privacy*, edited by B. A. Arrigo, 647. Thousand Oaks, CA: Sage.
- Walsh, P., and S. Miller. 2016. "Rethinking 'Five-eyes' Security Intelligence Collection Policies and Practices Post 9/11/Post-Snowden." *Intelligence and National Security* 31 (3): 345–368.
- Walzer, M. 1997. *Just and Unjust Wars*. 2 ed. New York: Basic Books.