Criminal Law and Philosophy
https://doi.org/10.1007/s11572-023-09704-5

ORIGINAL PAPER

Examining the Ethics of Spying: A Practitioner's View

David Omand[1]

Accepted: 7 October 2023

Abstract
This paper examines from the point of view of an intelligence practicioner the utility of the philosophical method that Professor Cecile Fabre has applied to intelligence ethics. Her emphasis on the duty that lies on governments to be sufficiently well informed about those who pose a real risk of serious violations of fundamental human rights is seen as a valuable addition to discourse on the ethics of intelligence activity. The just war tradition is put forward as an alternative framing of key ethical issues that can be translated into a practical code for intelligence officers that can be adapted to changing levels of threat in a way that is difficult to derive from a timeless philosophical analysis.

Keywords Intelligence · Espionage · Ethics · Moral philosophy

I write as a former practitioner who served at senior levels in the British security, intelligence and defence community but is now in an academic second career[1] and has wrestled with the question of whether 'principled spying' is an oxymoron.[2] And, given how that is the intent of successive British governments (at least in peacetime), I now assess whether that principled approach places our national security at risk from potential adversaries who have no such scruples, creating an asymmetry in their ability to steal our secrets as against ours to steal theirs. No eyebrows would be raised within the UK intelligence and security agencies at being required to demonstrate they have ethical standards by the Parliamentary Intelligence and Security

[1] As a visiting Professor in the War Studies Department of King?'s College London since 2005.

[2] David Omand and Mark Phythian, *Principled Spying: the Ethics of Secret Intelligence*, 2018, Oxford: Oxford University Press.

✉ David Omand
DAVIDOMAND@MAC.COM

[1] War Studies Department, King's College London, 47 Corinne Road, London N19 5EZ, UK

Published online: 20 October 2023

Committee (ISC) overseeing their activities. But what we have not had, until Professor Fabre's work, is a rigorous philosophical examination of ethical duty. Her book, *The Ethics of Espionage and Counter-Intelligence* (OUP, 2022), is much needed but is likely to be a hard read for those who have had no philosophical training, such as most of those who hold or have held senior positions in British intelligence. But as a former practitioner I found in her book many insights into the ethical positions taken up (or not) over the years by people like me. I am glad that Professor Fabre devoted the significant effort needed to write it.

A casual examination of intelligence practice (not least as it is described in intelligence fiction) reveals many issues that can generate ethical issues. There are many moral hazards associated with the practice of running covert human agents, both domestically and overseas. Issues will arise over how such agents are identified and recruited and what inducements may be offered and once in place what risks they (and their families) may run if exposed, including after they are resettled (such as demonstrated by the attempt in Salisbury by the Russian GRU to murder former MI6 agent Col Sergei Skripal and his daughter with the banned nerve agent Novichok). Where agents engage necessarily in criminal behaviour, or collude with such behaviour as part of their cover, it is an ethical as well as a legal decision as to what should be the limits imposed on them.[3] In recent years the growth of digital capabilities has also led to ethical concerns over the acquisition and use of personal information on individuals for intelligence purposes, including surveillance based on facial recognition systems, and the compatibility of such activity with the provisions of the Human Rights Act 1998 and the European Convention on Human Rights. Finally, there are now thriving international networks for intelligence sharing and operational cooperation, including with nations whose attitude to the rule of law and whose coercive interrogation practices are far from those of the UK.

Given all these potential areas of moral hazard and risk, training in the application of ethical principles – such as the principles of necessity and proportionality – is compulsory these days in the UK, especially since legislation has translated key principles into black letter law.[4] Each intelligence agency has, for example, ethics counsellors to provide training and (when necessary, confidential) guidance to staff. The wider intelligence community itself has since 1987 had a Staff Counsellor appointed by the Prime Minister[5] available to be consulted by any member of the agencies regarding matters of conscience about the work of their service, or a personal grievance or other problem which has not been resolved internally. The underlying thought, integral to modern British intelligence, is that operations should be assessed not just in terms of '*can* we do it', with a reasonable chance of delivering desired results, but also '*should* we do it' in terms of there being no ethically less risky course of action to achieve the

[3] The history of intelligence in the Norther Ireland campaign reveals many such instances, https://www.ppsni.gov.uk/news/statement-director-public-prosecutions-northern-ireland-relation-decisions-prosecution-arising.

[4] Investigatory Powers Act 2016, https://www.legislation.gov.uk/ukpga/2016/25/contents/enacted.

[5] See PM David Cameron's statement on 21 April 2016, https://hansard.parliament.uk/commons/2016-04-21/debates/16042142000020/StaffCounsellorForTheSecurityAndIntelligenceServices.

objective and that potential ethical risks are weighed in proportion to the harms that the objective of the operation is designed to manage.

In my case and that of most practitioners, especially those with a military background, the ethical framework within which such issues are evaluated tends to be one informally derived from longstanding 'just war' theory and practice.[6] That tradition recognises states have a duty to defend their citizens and uphold justice, and that protecting the innocent and defending moral values sometimes requires willingness to use violence. But it recognises that taking human life or deliberately seriously harming individuals is in itself morally wrong.

That leads to a loosely defined *jus in intelligentia*[7] adapting tenets such as right intention, right authority, proportionality, discrimination (in terms of the ability to assess and manage collateral harm) and last resort (meaning using other less ethically risky methods where these exist to achieve the same end).

A very much more rigorous philosophical analysis is applied by Professor Fabre. Of especial interest to me was her examination of the duty that lies on governments to be sufficiently well informed about those who pose a real risk of serious violations of fundamental human rights. In practice that means being prepared to uncover their secret plans and hidden activity through the practice of espionage; in Professor Fabre's words, 'that what matters is that one should seek to acquire information which, we have reason to believe, the other side does not want us to have'. The other side here must certainly include the dictators, autocrats and despots, terrorists and cyber gangs and others whose actions threaten serious harm (that would involve rights violations). This aligns with my own approach, which starts with a general definition of the purpose of human intelligence that is to improve the quality of decision making by reducing the ignorance of the decision-taker of what she faces. This is on the reasonable assumption that the more she knows about the threats she may face the more likely it is that her decisions will be effective in protecting her and her interests. The practice of secret intelligence has to achieve that purpose in respect of information that those others who mean harm do not want her to have (and may go to very violent ends to prevent her acquiring) so has to be stolen through cunning, deception and invasions of privacy. From that observation flow all the ethical issues around the practice of secret intelligence leading to the trio of key conditions for its use Professor Fabre identifies of effectiveness, necessity and proportionality.

Of particular interest in this book is the logical progression from establishing the conditions under which spying on others using ethically problematic methods may be morally justified to considering the conditions that may impose a moral duty on a state thus to spy. The latter argument recognises the duty of care to protect the rights of others, specifically those for whom a government has a direct responsibility, both members of the public threatened with serious rights violations or state servants who may be placed in greater danger through ignorance of the serious risks they face (such as members of the armed forces sent into action against a foe equipped with missiles, torpedoes or other deadly weapons of unknown characteristics). I find this compel-

[6] As for example can be found in Michael Waltzer, *Just and Unjust Wars*, 1977, New York: Basic Books.

[7] A term introduced, along with *jus ad intelligentiam*, by Sir Michael Quinlan, 'Just Intelligence: Prolegomena to an Ethical Theory', 2007, Intelligence and National Security 22 (1): 1–13.

ling. A comparable argument has been deployed in defence circles in part justification of the use of armed drones (under human control) since they remove the potentially lethal risk to the pilots of manned aircraft that would otherwise have to be sent to carry out the mission in the presence of air defences.

The book illustrates the strengths of the philosophical tradition in separating out sound arguments from the self-serving or simply specious reasoning that we all at times use to defend our actions in support of even our most principled objectives. There were few of her carefully argued conclusions that I would in principle cavil at when considering intelligence activity against foreign states or other overseas actors (the exception is in the treatment of surveillance, which I will describe later).

The philosophical method used in the book argues from a starting axiom using carefully worded examples to arrive logically at categories of permitted, mandated and ethically forbidden actions in intelligence activity. All are carefully worded to be as unambiguous as possible. This method creates a clearly defined ethical universe but one that relies upon a necessarily simplified model of activity in the messy, complex and deeply uncertain real world of international relations. The obvious question presents itself, how far can the conclusions derived from such a model be accepted and operationalised to guide real intelligence officers and their policy masters and to form the basis of legislation and regulation?

An example is to be found in the conclusion of Chap. 3, which states 'A political community is not justified in engaging in intelligence activities against another political community as a means to pursue an unjust policy, save in those cases in which they would, in so doing, bring about the morally weighty end of minimising rights violations (an example of only being second-best justified in so doing)'. How, in practice, in the complicated world of international relations, is one to know in most cases whether the policy being pursued is 'unjust' or not, or if judged presently unjust won't still end up doing good? Real policies have multiple ramifications, and they are likely to affect many states. A balancing judgement is bound to be needed as to the overall ethical merits of the policy. And different analysts will see that balance as lying at differing points on the scale. The intentions of the government may be honourable but there may well be room for legitimate doubt about where a policy may lead. For example, other states may react to the policy that Green adopts towards Blue in ways the government did not expect. Or steps intended to reduce persecution of a minority by Blue may fail to deter the perpetrators who may decide to up the ante and intensify the harm. And even if there were initial doubts about the ethical costs of a policy, such as economic sanctions, in the end it may unexpectedly achieve the morally weighty end of minimising rights violations. We can easily accept the general proposition that espionage in the service of an unjust foreign policy may not be justified – that is what normally distinguishes, we hope, intelligence in the service of a democracy from that of a despot. Professor Fabre also does allow for exceptions of which she gives telling examples. But how is the responsible intelligence officer, perhaps the MI6 Head of Station in a hostile capital faced with an immediate decision on giving the go-ahead to an intelligence operation, to know whether the ends to which her intelligence will be put are just? That is the reason why Prof Phythian and

I in our *Principled Spying*[8] book brought in to our discussion the aretaic tradition of moral philosophy. In the end we rely on the personal value ethics of our intelligence officers.

This question exposes what I see as a limitation of a philosophical method that creates categories such as a 'just foreign policy' and then uses them in a chain of reasoning from an initial premise or axiom, in this case that of universal rights, to a logical conclusion about the conditions under which espionage can be morally justified. But these categories are artificial constructs. Whilst the logical reasoning within the model may be impeccable we have to check carefully whether the conditions of the model can sensibly be mapped back onto reality as we experience it.

I am reminded of the economists whose models use artificial categories like gross domestic product, capital investment, savings, money supply and so on that are ex post constructs from national statistics but do not actually exist to be observed in reality in the way that some economic facts can be observed such as today's market spot price of Brent crude or the tax rates set by government or the interest rate set by the Monetary Policy Committee of the Bank of England. The outcome of modelling relationships between the constructed categories of the economist may provide a logical conclusion (such as a forecast of little growth in GDP this year given rising interest rates), but such conclusions have to then be mapped back to the reality of economic activity. Usually, the models provide insights into what is going on; but sometimes, notoriously, the artificialities of the model end up generating results adrift from the real economy.

A limitation is thus created by the inevitable imperfection of the categories used by the philosopher. As Professor Fabre acknowledges in her introduction there are blurred edges to many of these categories of relevance to intelligence: foreign and domestic threats intermingle; counter-terrorism is both a criminal and a national security matter; government and private sector actors both engage in espionage and so on. Problems arise when the human scale map of reality thus constructed for the model no longer provides a sure guide showing paths safe to walk to avoid the moral hazards of reality.

Almost all the decisions that require ethical judgements will, in practice, be shrouded in lack of knowledge both of the facts of the matter and of how events may turn out. I believe therefore that we should use the language of risk management. Some decisions concerning intelligence activity will have a higher risk of ethical problems arising than others. Following decisions to authorise specific intelligence gathering operations (or to hold back on ethical grounds), events may unfold smoothly and the sought-for information be obtained without undue difficulty. In other cases, intelligence officers will struggle to know what is the right course to follow when unexpected difficulties arise. Just how likely some ethical risk is to eventuate (say to the family of an informer if suspicions about his role arise) is a fine judgement in which prior experience is likely to be the best guide. To apply a precautionary principle of no quantifiable risk would be to rule out most intelligence gathering and could undermine the duty to collect information to avoid major rights violations. There would be no such argument of course if the information could be

[8] Omand and Phythian, op. cit.

obtained from open sources, but in the nature of the threats being faced where adversaries go to great lengths to protect their secrets mostly it is not.

In most cases there will be room for different assessments as to whether any particular intelligence gathering mission is really necessary. Those assessments in turn will be dependent on the ethical risk appetite of the authorities, and that will be dependent on just how vital is the morally justified end being sought. An example would be the case for using an ethically risky method to seek intelligence on a terrorist gang intent on mass murder. Even scraps of intelligence from a delicately placed agent may lead to the prevention of an atrocity. Compare that situation with the same type of intelligence method being contemplated to help with a long-term intelligence gathering effort to check on the use of novel technologies in the building of a new warship for an adversary navy. That would likely be itself a morally justified end but justifying a lower level of ethical risk in the latter case than the former.

Decisions to authorise intelligence activity are usually taken under significant uncertainty as to the outcome. For the intelligence officer it involves taking risks in investing in the uncertain expectation of a return. Casting bread upon the waters is unavoidable. When an intelligence officer approaches a potential agent she cannot be certain that a productive relationship will ensue or whether the approach will be reported, thus increasing the risks to the officer herself. A practical obstacle in the way of the intelligence officer who wishes to act ethically is the impossibility of defining in the abstract what should be the threshold for adequate justification for a recruitment approach. An estimated 50% chance of success? Requiring a 75% chance of success? Should the threshold be set with regard to the value of the information being sought and if so by how much? The Iraq experience shows the danger of lowering the bar when the demand for intelligence is greatest.

There is a parallel here with contemporary debates about the ethics of artificial intelligence.[9] For example, GCHQ is committed to creating and using AI in a way that supports fairness, empowerment, transparency and accountability – and to protecting the nation from AI-enabled security threats pursued by our adversaries. But, to take an example, reaching the 'right' balance between false positive and false negative results from the application of an algorithm in intelligence gathering is an ethical judgement. There can be no objective answer without taking into account the relative ethical and material costs of being wrong either way (for example in applying secret intelligence to construct a no-fly list that is bound to result in some innocent passengers being misidentified by the AI system and thus put to considerable inconvenience. Such false positives are in addition likely to result in discrimination against particular ethnic or religious groups). In probabilistic terms, too, it may be a portfolio approach that is being considered, where not all of the routes being explored will return the desired results (for example when faced with uncertainty as to which fibre optic channel will carry a desired stream of traffic; even if the risk of failure in any single channel is 50%, a portfolio of ten channels sampled would yield an overall success rate of over 99%).

[9] A debate in the intelligence word prompted by the Director of GCHQ in a paper published in February 2021, https://www.gchq.gov.uk/artificial-intelligence/index.html.

There is an important distinction here, one that Professor Fabre's book does not explore, between the creation of capability to gather secret intelligence and the use of that capability once built to try to acquire intelligence on specific targets for specific purposes. In the case of human intelligence it is expensive to set up overseas intelligence stations, with trained personnel who have acquired the necessary language skills and tradecraft to be able to operate sufficiently safely. Once such a station is established, however, it can be tasked to provide intelligence on very different subjects, some that Professor Fabre might regard as morally legitimate and others more questionable. But without the human and technical infrastructure neither would be possible.

The same distinction applies even more forcibly in the digital space. The capability of GCHQ to derive intelligence from digital methods (bulk access to data, network interference and hacking, large scale data processing and storage, data mining and so on) requires a heavy prior investment in human and technical capability. Once the capability exists the (metaphorical) aerials can be turned to point in many different directions, but only once established.

The distinction between *jus ad intelligentiam* and *jus in intelligentia* comes to mind (although it is not an exact parallel). But the analogy invites the question, under what conditions is it justified for a government to provide such capability, recognising that once established it can in theory be used for very different purposes against very different targets, for good or ill.

I imagine that a philosopher following Professor Fabre's reasoning would have no difficulty in constructing a use case involving a very serious threat to fundamental rights where a government that wishes to act ethically would have had to have so equipped the intelligence agencies in advance with these capabilities (as a practical matter that equipping can take years of effort). But there would presumably then be a duty on a democratic government to ensure these powerful capabilities were not misused for unjust purposes (as is the case with the Chinese use of digital intelligence techniques).[10]

That is the reasoning that led to the Investigatory Powers Act 2016 and the regulations and oversight conditions laid down in it. I was a member of an independent commission set up by the then coalition government that fed into the drafting of the legislation.[11] We called for the provision of the digital capabilities to be subject to the 3R test: to be under the **rule of law**; to be subject to **regulation**, both judicial and Parliamentary; and for such coercive powers to be used with **restraint**, applying the ethical principles of necessity and proportionality to the authorisation of operations (once judged to have a reasonable chance of success).

Such considerations take us beyond the scope of the book. But what I am pointing to here is the general problem of operationalising ethical rules that almost always involve balancing acts within the basket of universal rights. The best that can be

[10] As described by the UK Parliamentary Intelligence and Security Committee in their report on China of 13 July 2023, https://isc.independent.gov.uk/wp-content/uploads/2023/07/ISC-China.pdf.

[11] RUSI, *A Democratic Licence to Operate: Report of the Independent Surveillance Review*, 13 July 2015, https://rusi.org/explore-our-research/publications/whitehall-reports/a-democratic-licence-to-operate-report-of-the-independent-surveillance-review.

hoped for is that the ethicist or moral philosopher can help the decision maker disentangle the complex considerations involved in ways that simplify the task of deciding what it is in the specific circumstances to act ethically. I believe that leads us to the value of having a **code of behaviour for intelligence officers**, recognising that such an ethical code is a defining characteristic of any profession. Ethics training is now mandatory for officers in the British agencies and there are ethics counsellors to whom staff can turn for support when hard issues arise.

In her conclusion, (1.5, p. 36) Professor Fabre says that the framework of just war theory is not of decisive help as she tries to make sense of the intuition that, in some cases involving the defence of fundamental moral rights, intelligence activities are morally justified. I accept her view that the just war framework cannot be of *decisive* help in a philosophical investigation such as hers. But I do maintain that the framework is subsequently of clear practical benefit in turning theory into praxis. Set down as principles such a code helps unbundle those complex ethical considerations likely to be involved in a decision to mount an operation to gather secret intelligence and make them more tractable to ethical reasoning. Decisions that often have to be taken under time pressure and stress.

My reading of the just war tradition therefore led me to propose a short set of ethical principles that can be readily applied to everyday judgements inside the intelligence community and applied as a standard by those who oversee and regulate that community. Three of those principles are the same as those that come directly from the close reasoning that Professor Fabre has conducted by applying the single moral axiom, the foundational principle she cites, that of the principle of fundamental equality.

There is thus the principle of **proportionality** to ensure that the ethical risks of operations are in line with the harm to fundamental rights that the operations are intended to prevent. In Professor Fabre's formulation, we should only be justified in imposing a harm when pursuing a morally justified end when the good one thereby brings about (in the form of that justified end) outweighs the harm. Such proportionality judgements involve weighing up many kinds of uncertainty about whether that good will actually be delivered, for which the UK Courts have been willing to allow 'a margin of appreciation' for the decisions taken by the security and intelligence agencies.[12]

There is also an element of counter-factual thinking involved in such a balancing act when it comes to authorisation of an operation since *not* conducting the operation also involves potential ethical risk. As John Stuart Mill pointed out in the mid-nineteenth century: 'A person may cause evil to others not only by his actions but by his inaction, and in either case he is justly accountable to them for the injury'.[13] This principle is well established in medical ethics, for example in the balancing act that a clinical team may have to make in intervening in medical emergencies in childbirth or following major accidents. Another example of proportionality judgement is in the licensing of new drugs and therapies where the expected benefits to a large number of

[12] David Anderson, *A Question of Trust: Report of the Investigatory Powers Review*, 2014, https://www.gov.uk/government/publications/a-question-of-trust-report-of-the-investigatory-powers-review.

[13] John Stewart Mill, *On Liberty*, 1969, ch. 1.

sufferers from a disease have to be balanced against the risks of adverse side effects for a small number of patients and how far that can be mitigated by training of clinicians and warning leaflets with prescriptions.

There is **necessity** – acting in that way, as Professor Fabre puts it, only: 'if there is no lesser morally weighted harm the imposition of which would bring about one's morally justified ends'. In the intelligence gathering context, this means being assured there is no other reasonable way to achieve the result of the proposed information collection mission at lesser ethical risk of causing harm, such as through open sources or preferring technical to human sources. Here I use the term ethical risk rather than cost since, as mentioned earlier, there may be uncertainty at the time of authorisation of the mission whether that cost will fall to be incurred.

Professor Fabre lists **effectiveness**. In her words, we can impose harm or risk on another in pursuit of a morally justified end only if one's course of action stands a reasonable chance of succeeding. In my writing I have called this test that of a **reasonable prospect of success** – having adequate justification for the expectation that conducting specific operations in pursuit of a morally justified end will be likely to deliver results of value. Adequate justification is more than just 'try it and see' and hoping that the results will accrue. There has to be evidence – it may be from previous similar operations or it may be by applying in-depth knowledge of how the intelligence processes involved would work – that gives confidence that applying some new previously untried technique may deliver. I would, however, reject a proposition that nothing must ever be done for the first time.[14]

I would therefore endorse the conclusion Professor Fabre reaches that an agent, G, is in her view morally justified (p. 31) in harming another, B, if the course of action is necessary, effective and proportionate and of course that B is guilty of contributing to the violation of some agent's fundamental rights (or failing to protect them). But I would pause on her 'if and only if' condition since it may be that we have evidence that B is preparing such a violation – it is not just after the harm has occurred that we can reach a judgement to act.

There are solid examples given by Professor Fabre of applying such a test that I find useful. The many debates I had with political scientist, Professor Mark Phythian, on the ethics of secret intelligence that resulted in our book of dialogues *Principled Spying* (that Professor Fabre kindly commended) led us to the position that what we should be seeking is a generally accepted ethical code of practice for those who practice secret intelligence. Generally accepted, meaning that in a democracy it has acquired popular support through democratic debate. That is not to say that the code is justified *because* democratically accepted, but rather that it is enforceable because the code sets a standard against which Parliamentary oversight can be applied and relevant principles can be enshrined in law. Drawing on just war thinking I have separated out three self-standing principles that are especially useful in the digital age to complement effectiveness, necessity and proportionality.

[14] One of the maxims of a favourite book, Francis Cornford, *Microcosmographia Academica: Being a Guide for the Young Academic Politician*, 1908, Cambridge: CUP (essential reading for all intelligence officers as well as academic administrators).

The first is **discrimination** in its old fashioned sense, as used in just war writing, of having the ability to assess and manage the risk of unintended (collateral) harm, such as the shelling of an enemy position killing sheltering civilians. Military commanders are under a legal duty to seek to minimise such harm. The classic example is delaying the attack on the school whose cellar is being used as a command post until after the school has closed for the evening. A digital intelligence example would be seeking to minimise privacy intrusion by intelligence officers into the lives of those who are not the intended targets of intelligence gathering (recognising that like collateral casualties on the battlefield there can be no guarantee they can be avoided).

Discrimination is needed to manage collateral harm, in the basic sense of the ability to see the difference between classes of things or people. In the laws of war that emerged from the just war tradition, the military commander faces on the one hand legitimate military targets and on the other hand groups of people that require protection such as innocent civilians not participating on the side of the adversary or surrendering soldiers. Before a new type of weapon is introduced into the battlefield there needs to be a legal assessment that the combination of weapon and operator is capable of discriminating between them. By analogy, when a new digital intelligence gathering method is introduced there needs to be confidence that there will be the human and technical ability to assess and manage the risk of collateral harm, including the implications of privacy intrusion into the lives of those not intended to be the target of intelligence gathering. The principle of discrimination also provides the basis for ethical oversight of the artificial intelligence algorithms that are increasingly being used to question large data sets. In any practical decision system (whether conducted by humans, by humans assisted by machine intelligence or by AI algorithms themselves) what is a reasonable prospect of success has to be defined, given that the possibility of error cannot be excluded. AI applications that have a low rate of false positives are said to have high specificity. Those with a low rate of false negatives have high sensitivity. Where the cursor is set between these will depend upon the consequences of getting it wrong either way.

As Professor Fabre notes, collateral intrusion has always been an issue with interception of fixed line telephones that capture the communications of all users in a home not just the targeted suspect. Procedures for discarding irrelevant information are routinely applied in the UK. Her analysis would, however, rule out another common feature of intelligence work, that the ultimate target justifying interception may not be the intelligence target herself. She writes (p. 182) ‘Green must have evidence-based reasons to believe that the individual is acting in such a way as to be liable to being subject to observation and interception tactics’. But if it is known that a foreign politician, G, of third country Orange has a habit of chatting on a mobile telephone to the leader of a hostile state, Blue, suspected of planning an unjust attack then I believe it would in serious cases be justified to authorise interception of G (if feasible) (such as the interception of inter-war Japanese communications in the justified expectation that clues to Nazi Germany’s intentions would be revealed). In such cases what matters is the system of regulation for dealing with material not relevant to the justification for the operation, for example in terms of destruction of raw material.

An example of the difficulties into which artificial categories can lead the philosopher is in the treatment of surveillance in Chap. 9. Professor Fabre’s conclusion is

that not all aspects of mass surveillance issue in a loss of privacy. I suspect that would better read that not all intelligence operations involving bulk access to personal data issue in a loss of privacy. This is another case where the answer you get depends upon the categories onto which the philosopher tries to map reality.

'Mass surveillance' and 'bulk access to data' are orthogonal concepts. Mass surveillance is best understood in terms of the dictionary definition of being the persistent observation of the population or a sizeable part of it. As a former senior judge concluded (when acting as the UK Interception Commissioner), such mass surveillance would be comprehensively unlawful under UK law. Successive senior judges have confirmed that GCHQ does not conduct mass surveillance (although they have found other problems with some of GCHQ's procedures, for example concerning the destruction of material, and have insisted upon improvements). There is no group of analysts at GCHQ (or the other agencies) conducting persistent surveillance of the UK population.

Contrast 'mass surveillance' with the category of 'bulk access to data'. All internet communications are sent in a series of data packets (via the packet switched networks of the internet, unlike in the days of telephone interception of dedicated rented communications channels). Different packets may take different routes before being reassembled in the computer of the recipient of the message. The simplest possible example would be interception of a sought-for message from the intelligence target (supposing that the IP address of her computer is known to the intercepting authorities). That involves accessing flows of data in bulk (through fibre optic cables, microwave links and satellite links), filtering to remove and discard non-communications content, and further deep packet inspection to try to identify the packets of data from that IP address. The operation is an exercise in targeted intelligence gathering, not mass surveillance, although the machines have to sort through vast quantities of non-relevant material to acquire the wanted communications. Essentially the same position exists for the US agencies under the US Constitution.

There have to be reasons that the analyst can give (and that the senior judge acting as Interception Commissioner can accept) for accessing stored bulk data in pursuit of an intelligence investigation. Professor Fabre is much too restrictive when she says (p. 203) 'But only when *the analysis of open sources* has given intelligence agencies probable cause may they retrieve already-collected non-open information about specific individuals for analysis' (my emphasis). The probable cause may legitimately come from intelligence shared by another state, or from quite different covert sources (such as interception or eavesdropping or just the result of more intensive intelligence analysis) that provided the clue that sparks the investigation and justified the inquiry of the data.

Much more complex examples might involve other indicators as search terms (names, date/time details of a wanted message), patterns of communication known to be used by a target group (a method in constant use by the banks to detect ring fraud) and other indicators that can form the algorithmic 'magnets' that can pull the needles from the data haystacks.

Bulk access operations have to be operated under clear ethical constraints to take account of the evident fact that almost all bulk data by definition relates to those who are not and would never be the legitimate target of intelligence activity. Is it

legitimate to use the argument that there is no risk of rights violation because no human being ever sees the vast mass of material that the machines do not forward in response to a query? Professor Fabre (p. 26) adopts a non-experiential account of harm to mean: that of which I am not aware and which I cannot experience can still harm me. There is, for example, the argument that knowledge of the possibility of bulk access leads to a 'chilling effect' on interpersonal communications (the panopticon parallel of Foucault). But just because agile minds can conceive of chilling does not mean that it should necessarily have significant weight in the balance. That is an empirical matter (I am not aware of solid evidence that it is a significant factor in the public's use of the internet).

Professor Fabre rightly does see the risk that could arise if material at different stages of processing is stored and therefore personal data of innocent members of the public is in theory open to examination. That risk has to be managed by regulation. A view, therefore, that I subscribe to is that agencies engaging in bulk access operations must recognise that public privacy rights are engaged right from the outset of planning such bulk operations. But with careful design of algorithms and procedures for destroying unexamined material after a set period that ethical risk can be managed down to an acceptable level applying the necessity and proportionality tests. The parallel exists in warfare where a new weapons system should not be fielded without assurance that the combination of weapon and operator is capable of sufficiently discriminating between legitimate and illegitimate targets (the principal argument against fully autonomous armed drones is that even advanced AI is not capable of giving that assurance on a crowded battlefield).

I would add explicitly to a list of ethical principles, although implicit in Professor Fabre's premises, **right intention** – acting with integrity and having no ulterior motive or other agenda either behind the authorisation of intelligence activity or in analysis, assessment and the presentation of intelligence judgements to decision makers. The principle of right intention does not rule out deception in the course of an intelligence operation, such as inserting into the digital code of malware clues that attempt to encourage a false-flag attribution to a third country. But there must be no deception of government or Parliamentary overseers, or hidden domestic political or personal agendas lying behind the authorisation or the conduct of digital intelligence activity.

And although also a practical institutional rather than a fundamental rights consideration, I would add for completeness **right authority** – establishing the level of decision making and independence of scrutiny before decisions are taken, appropriate to the ethical risks that may be run. Having right authority properly recorded allows for legal and Parliamentary accountability for decisions and oversight of activities where there may be legitimate grounds for questioning whether actions should have been ethically permitted.

The just war approach of course carries the implication that actions taken under the code (and how the code itself is drawn up) can be changed in accordance with changing current circumstances (for example when a nation is engaged in existential armed conflict) and with historical developments in moral reasoning. I see that as an advantage since allows for different judgements to be properly arrived at in times of war from in times of peace, or from greater and lesser threats. The model ethical

universe that Professor Fabre sets out in her book appears to be absolutist, setting down for all time when secret intelligence may be authorised. Would we claim that future generations, perhaps faced with circumstances of planetary survival we cannot imagine today, cannot, if they are to act ethically, choose some other path? Should we judge the morality of decisions taken by many generations long preceding the Second World War by the standard of the UN Declaration of Universal Rights that followed that experience of total war?

And is Professor Fabre pushing the concept of universal rights too far when she concludes (p. 50) 'to the extent that citizens of a democratic polity have a democratic right to pursue a particular foreign policy end and that they have good reasons for wishing to keep its details secret, they have a democratic right *that third parties not seek to appropriate and disclose the relevant information*' (my emphasis). I can see that in those circumstances the democratic state has a right to take significant steps to protect its secrets. But what does it mean to say that there is a 'right' that other states not try to spy out those secrets? Can there be a right not to be spied on by other states – enforceable by whom? That I suspect is another example where the ethical theory cannot be turned into praxis.

Professor Fabre maintains (p. 58): 'to take rights seriously is to commit oneself to the view that one may harm the interests which they protect *only in response to* rights violations or justified rights infringements' (my emphasis). But to reiterate the point, must states wait for the violations to have occurred and the damage suffered? Much intelligence activity is precautionary, including the large effort devoted to acquiring information about military capabilities and weapons procurement, just in case. To which we could add the effort to try and understand political developments in the Chanceries of overseas states (we could wish there had been better intelligence gathering in the years leading up to 1914). I would rather rely on the form of the argument Professor Fabre deploys (p. 196) that *investigative* cyber-counter-intelligence can be justified.

At the risk of being whimsical, in conclusion I am reminded of the failure of the monumental work of Hilbert to derive the foundations of all mathematics from the smallest number of axioms and of logic rules. As Kurt Gödel demonstrated, there will nevertheless be valid mathematical propositions that cannot be proved in any such consistent formal system. Nor can the system demonstrate its own consistency. Alan Turing, well known to intelligence scholars from his starring role at Bletchley Park, likewise demonstrated with Alonso Church that there is no algorithm capable of solving the halting problem that lies at the heart of modern computing.

In conclusion, readers will have to judge for themselves whether the book as well as adding very significantly to sound thinking on the ethics of intelligence may not also illustrate the limitations of philosophical method as a guide to how intelligence officers should act ethically, when they come to consider whether to stay their hand from some espionage operation or, one of the most interesting aspects of the book, on the contrary to regard themselves as being under a reinforced duty in some circumstances to gather secret intelligence even at some risk to those involved.