

Washington University in St. Louis

Washington University Open Scholarship

All Theses and Dissertations (ETDs)

Spring 4-29-2013

Spying: A Normative Account of the Second Oldest Profession

Ronald E. Watson

Washington University in St. Louis

Follow this and additional works at: <https://openscholarship.wustl.edu/etd>



Part of the [Political Science Commons](#)

Recommended Citation

Watson, Ronald E., "Spying: A Normative Account of the Second Oldest Profession" (2013). *All Theses and Dissertations (ETDs)*. 1106.

<https://openscholarship.wustl.edu/etd/1106>

This Dissertation is brought to you for free and open access by Washington University Open Scholarship. It has been accepted for inclusion in All Theses and Dissertations (ETDs) by an authorized administrator of Washington University Open Scholarship. For more information, please contact digital@wumail.wustl.edu.

WASHINGTON UNIVERSITY IN ST. LOUIS

Department of Political Science

Dissertation Examination Committee:

Andrew Rehfeld, Chair

Clarissa Hayward

Jack Knight

Frank Lovett

Ian MacMullen

Neil Richards

Spying: A Normative Account of the Second Oldest Profession

by

Ronald Edward Watson

A dissertation presented to the
Graduate School of Arts and Sciences
of Washington University in
partial fulfillment of the
requirements for the degree
of Doctor of Philosophy

May 2013

St. Louis, Missouri

copyright by

Ronald Watson

2013

Contents

List of Tables	v
List of Figures	vi
Acknowledgements	vii
1 Introduction	1
I The Concept of Spying	10
2 The Concept of Spying	11
2.1 A Few Words on Method	13
2.2 Towards a Conception of Spying	16
2.3 Objections	24
2.4 Spying and Espionage	36
2.5 Conclusion	39
2.6 Appendix: Previous Definitions of “Spying”	40
II The Ethics of Government Spying	42
3 Principles for Domestic Government Spying	43

3.1	The Presumption against Domestic Government Spying	45
3.2	Overriding the Presumption	52
3.3	Spying on Innocents	55
3.4	Harms - Proportionate and Minimal	58
3.5	Institutionalizing the Principles	61
3.6	Conclusion	63
4	Two-Level Utilitarianism	65
4.1	Classical Utilitarianism	67
4.2	Two-Level Utilitarianism	77
4.3	The Instability Objection	86
4.4	Conclusion	89
4.5	Appendix: Utilitarianism's Commitments	90
5	Some Consequences of Spying	91
5.1	Successfully Concealed Spying	94
5.2	Suspected Spying	112
5.3	Conclusion	128
5.4	Appendix: The Consequences of Spying	129
6	Utilitarian Principles for Domestic Spying	131
6.1	General Considerations for Utilitarian Principles	132
6.2	Developing Utilitarian Principles	133
6.3	Ancillary Principles	153
6.4	Objections	154
6.5	Conclusion	163

7	Principles for Foreign Spying	165
7.1	Spying on Foreigners	169
7.2	Spying on Foreign States	189
7.3	Institutionalizing Principles for Foreign Spying	205
7.4	Conclusion	209
III	The Institutions of Government Spying	211
8	Controlling Government Spies: The American Model	212
8.1	Legislative Oversight	214
8.2	Judicial Review	233
8.3	Conclusion	242
9	Reforming the Control of Government Spies	244
9.1	Mechanisms of Control	246
9.2	Reforming American Mechanisms of Control	261
9.3	Conclusion	280
	Bibliography	282

List of Tables

2.1	Four Types of Collective Intentionality	30
2.2	Some Previous (Non-Dictionary) Definitions of “Spy” or “Spying” . .	40
5.1	The Consequences of Spying	93
6.1	Just Causes for Spying	139
7.1	The Targets of Government Spying	167
9.1	Overseers for <i>Ex Ante</i> Review	274

List of Figures

4.1	The Commitments of Utilitarianism	90
5.1	The Consequences of Successfully Concealed Spying	129
5.2	The Consequences of Suspected Spying	130

ACKNOWLEDGEMENTS

I am grateful to many friends and colleagues for their comments, questions and support. The four political theorists at Washington University – Clarissa Hayward, Frank Lovett, Ian MacMullen, and Andrew Rehfeld – were rocks. They all read and reread drafts of my chapters and patiently helped me see things more clearly. They were not just careful readers, but also warm advisors and good friends. Jack Knight twice flew me to North Carolina, offered insightful criticisms on my chapters, and introduced me to some of Raleigh-Durham’s best restaurants.

I would like to thank many others who read drafts or provided me with helpful suggestions. Neil Richards, Julia Driver, Kit Wellman, Carl Wellman, John Inazu, Nate Adams, David Speetzen, Jill Delston, Keith Schnakenberg, Gordon Arsenoff, Ian Ostrander, Noel Pereyra-Johnston, Chris Claassen, Greg Whitfield, Matt Chick, Cristian Perez, Diana O’Brien, Kate Jensen, Brandon Nelson, Santiago Olivella, Anthony Stenger, Randy Calvert, Molly Scudder, Lorraine Krall.

I would also like to express my gratitude to my family in St. Louis: Mom, Leanne, Lance, Caden, and Zach. Thank you all for your love and support.

Finally, I owe the largest debt to my wife. Sarah treated this project like it was her own. She was always encouraging, patient, and supportive. I could not have asked for a better partner.

Chapter 1

Introduction

“Reliable intelligence agents,” Hobbes (1998, 144-145) says, “are to those who exercise sovereign power as rays of light to the human soul.” He continues:

... they are as necessary to the safety of a commonwealth as rays of light to the safety of a man. Or we may use the analogy of spiders’ webs, whose incredibly fine threads spread out in all directions and convey outside movements to the spiders sitting in their little cavities inside. Without intelligence agents sovereigns have no more idea what orders need to be given for the defence of their subjects than spiders can know when to emerge and where to make for without the threads of their webs.

Hobbes’s metaphors of rays of light and spiders’ webs illustrate an important purpose of government spying. Spying has the potential to secure governments from any number of threats – foreign and domestic. A government with skillful spies is alert to the gravest threats facing it. It has a window into the intentions of criminals, terrorists, revolutionaries, hostile states and others who seek to do it harm. Without competent spies, a government is in the dark and therefore vulnerable.

The protective role that government spying can play is in part why governments

have spied since the beginning of recorded history and why governments expend so many resources to build and maintain the capacity to spy today. Records of spying stretch so far into antiquity that spying is often called the second oldest profession. (Knightly, 1986) In the Old Testament, Moses sends twelve spies into Canaan to investigate the Promised Land. They were instructed to go to the mountain and:

see the land, what it is; and the people that dwelleth therein, whether they be strong or weak, few or many; and what the land is that they dwell in, whether it be good or bad; and what cities they be that they dwell in, whether in tents or in strongholds; and what the land is, whether it be fat or lean, whether there be wood therein, or not. And be ye of good courage, and bring of the fruit of the land. (Numbers 13:18-20, original emphasis)

Thucydides (1972, Book VI, Ch. XLV) in his *History of the Peloponnesian War* mentions that the Syracusans learn of the coming attack by the Athenian fleet from spies. Sun Tzu in his *Art of War* offers a typology of five different kinds of spies. When all are at work, he says, the sovereign has “the divine manipulation of threads.” The sovereign’s network of spies is his “most precious faculty.”

Governments have always been spying, then. But the resources that they put to spying today is unprecedented. A recent study by Hippner (2009) estimates that the world’s governments spend over \$106 Billion each year on foreign intelligence alone. In the United States, spending on foreign intelligence peaked in 2010 at \$80.1 Billion, but in 2012 America still spent \$75.4 Billion, more than all other countries combined, and approximately \$250 per American citizen.¹

¹The Federation of American Scientists (FAS) aggregates intelligence budget data at: <http://www.fas.org/irp/budget/index.html>

Although states can hardly do without spies if they wish to remain secure, there is, of course, a darker, more sinister picture of government spying. This darker picture is one of the government using spying to control every aspect of its citizens' lives, compelling them to act and think only in ways sanctioned by the state. The picture is best represented by the Thought Police in Orwell's *1984*, in which Winston, the protagonist, remarks:

There was of course no way of knowing whether you were being watched at any given moment. How often, or on what system the Thought Police plugged in on any wire was guesswork. It was even conceivable that they watched everybody all the time. . . . You had to live – did live, from habit that became instinct – in the assumption that every sound you made was overheard, and, except in darkness, every movement scrutinized. (Orwell, 2007, 3)

This darker picture of spying cannot be dismissed as unrealistic, merely a dystopian nightmare. States have often controlled and continue to control their citizens with spying. Nowhere was this kind of control more complete than in the German Democratic Republic (GDR) during the Cold War, however. Historian Hubertus Knabe summarizes the control exacted by the GDR's secret police (Stasi) as follows: "Precisely the hidden, but for every citizen tangible omni-presence of the Stasi, damaged the very basic conditions for individual and societal creativity and development: sense of one's self, trust, spontaneity." (Bruce, 2010, 12) At the height of the Stasi's power, it employed one secret police officer for every 180 of the GDR's citizens. (Ibid, 11) Throughout the course of the GDR, it is estimated that nearly 600,000 citizens served as informants, a startling number given the population of East Germany never exceeded 17 million. (Dennis, 2003, 90) Since the fall of the Berlin Wall, numerous East

German citizens have visited the Stasi's behemoth archives to view their own files, only to learn that they were informed on by their spouse, by their parents, or by their children.

It is tempting to view the Stasi as a relic of a bygone era. But even in liberal democratic states spying is sometimes employed to intimidate, harass, discredit, or gain an unfair advantage on political opponents. America's Federal Bureau of Investigation (FBI) with its counterintelligence program (COINTELPRO) in the 1960s, for example, attempted to subvert the civil rights movement by intimidating and discrediting its leadership.²

Spying is also sometimes bluntly employed by liberal democracies against groups unreasonably thought to be dangerous. During the first World War, for example, the American military secretly monitored blacks, fearing that they might sabotage American infrastructure. Similarly, during the Vietnam war, the FBI spied on anti-war protesters. This time the fear was that anti-war protests would turn violent, or that they would undermine popular support for the conflict. Finally, after September 11, 2001 (9/11) America's intelligence agencies and some local police forces intensely spied on Muslims, without in many cases any reason to think they were linked to terrorist plots.

Hence although government spying has the potential to secure states against threats, it also has the potential to subvert cherished liberal and democratic rights. The ethics of government spying thus requires that governments skillfully tie their own hands. If a government too tightly binds its hands and prohibits itself from too much spying, it leaves itself vulnerable to any number of grave threats. But, in contrast, if it ties its hands too loosely and spies overzealously, it threatens individuals'

²I discuss this case in more detail in chapter 5.

rights, and it risks subverting the liberal democratic culture of free thought, speech and action.

Binding the hands of government agents is not just profoundly important ethically, it is also urgent. A century ago, when relatively little information was collected about individuals, and government spying meant eavesdropping or stealing documents on site, government spying presented a limited threat to citizens and promised utility for policymakers in only a small set of cases. Now we live in an information society. Virtually every aspect of people's lives are recorded, stored, bought, sold, and distributed, and governments have unprecedented capabilities to covertly collect and compile information.

The growth in information collection and storage over the past two and a half decades is breathtaking. According to Hilbert and Lopez (2011), the world's technological installed capacity to store information, which includes everything from books, newsprint, and film to memory cards, DVDs, and hard-drives, increased from 2.8 (optimally compressed) exabytes (i.e. 2.8 quintillion bytes) in 1986, to 15.8 exabytes in 1993, to over 54.5 exabytes in 2000, and to 295 exabytes in 2007. In other words, in 2007 the world's capacity to store information was 105 times larger than it was just 21 years prior. Of course much of this information is neither personal nor secret. But the quantity of stored personal information and stored secrets has likely kept pace with the overall trend. Consider, for example, that according to the Information Security Oversight Office the number of pages that have been classified in the United States jumped from 5.8 million in 1996 to 54.8 million in 2009.³

As the quantity of information stored has expanded, so too have governments' capacities to collect information. The news is awash with stories of the intelligence

³The report is available online at: <http://nsarchive.files.wordpress.com/2010/04/2009-annual-report.pdf>.

community's new high-tech collection methods, such as ECHELON (a partnership among five Western countries to collect telephone, fax, email, and other data traffic by intercepting satellite communications) or Trailblazer (a program that collected and analyzed data en masse from the internet). America's National Security Agency (NSA), the inventor of Trailblazer, has become such a voracious collector of data that it is currently building a \$2 billion data warehouse of almost unimaginable size in the middle of the Utah desert. (Bamford 2009) Over forty years ago, the philosopher H.J. McCloskey (1971, 305) imagined what he called a "brain-mind-bug" that could collect a person's every thought and feeling without their knowledge. Today such a technology may not be far outside of some intelligence agency's grasp.

So the information revolution and the rapid technological leaps governments have made in collection technology increase the urgency of spying as an ethical issue. So too does the emergence of non-state threats. These threats, such as terrorism and organized crime, cannot be as easily ascertained and contained as the rogue states of past generations. They do not announce their intentions, and they are difficult to deter with stockpiles of weapons. As Robert Baer (2002, 271) remarks at the end of his memoir recounting his experience in the CIA's Directorate of Operations, the enemy today is

an enemy with no infrastructure to attack, with no planes to shoot out of the sky, with no boats to sink to the bottom of the sea and precious few tanks to blow up for the amusement of viewers on CNN. The only way to defeat such an enemy is by intelligence, by knowing what they plan to do next, and by being ready for them when they arrive.

Hence for a government to be in the dark today – to be without reliable spies – is more dangerous than ever.

Yet despite raising important and urgent ethical and institutional problems, government spying has not yet been tackled head-on by political theorists or by philosophers. A variety of helpful associated literatures of course exist. Legal scholars and philosophers have devoted considerable attention to privacy, sociologists have pioneered “surveillance studies,” and a host of intelligence practitioners have provided their own reflections on the ethics of spying.⁴ But those who study political ethics most closely have surprisingly not yet had their say on government spying.

The purpose of my dissertation is to take the urgent ethical problems raised by government spying seriously. In what follows I ask and answer three principal questions: What is spying? What principles should regulate government spying? And how can government agents be constrained to follow these principles?

The first chapter takes up the conceptual question. I defend the following definition of spying: agent A spies on agent B, if and only if she collects information that relates to B and intends to conceal her information collection from B. The main challenge any conception of spying faces is to cover a relatively wide range of agents often thought to be spies – e.g. defectors, moles, and informants – without including agents not often thought of as spies. This challenge is best met, I argue, by drawing on the concept of collective intentionality. Aldrich Ames, for example, although he did not intend to conceal his information collection from the U.S. government, nevertheless spied on the U.S. government, since he participated in a collectivity, which included his handlers at the KGB, that met both of the conditions stipulated in my

⁴For an overview of the philosophical literature on privacy, see Schoeman (1984). Lyon (2001) is a good introduction to the surveillance studies literature. Goldman (2006, 2010) are the best collections of essays by intelligence practitioners.

definition.

Chapters three through six examine the ethics of domestic government spying, i.e. governments spying on their own citizens within their own territories. I make two main arguments. The first is that domestic government spying should be regulated by five principles: just cause, proportionality, necessity, minimization, and discrimination. The second argument is that the law-enforcement and intelligence officials who employ these principles should not alone determine how they apply in particular cases – the principles should be institutionalized.

In chapter three I demonstrate that the five principles are supported by widespread intuitions about government spying in liberal democracies. In chapters four through six, I show that the same principles are supported by the moral theory that I think is most plausible: two-level utilitarianism. Since utilitarianism is often thought to strongly conflict with people’s ordinary moral intuitions, if I am correct and the same principles can be derived both from widespread intuitions and utilitarianism, then the principles are on strong ground.

In chapter seven, I shift my focus from domestic government spying to foreign spying. I focus on two kinds of foreign spying in particular: government spying on foreign individuals and government spying on foreign states. I argue that government spying on foreign individuals and on foreign states should be institutionalized and that both should follow principles similar but not identical to those that governments should follow in the domestic context.

In the final two chapters I turn to the institutional question. In chapter eight I examine the two primary American institutions employed to control intelligence agencies: legislative oversight and judicial review. Both I argue employ biased principals and suffer from informational asymmetries. In chapter nine I step back from American institutions and characterize the universe of possible mechanisms to constrain

intelligence agencies. Drawing on some of the more promising of these mechanisms, I propose a set of reforms for American institutions. At the heart of my proposal is an elected panel that reviews day-to-day requests to spy and performs longer-term strategic oversight. In order to allay the panel's informational disadvantages compared to intelligence agencies I recommend, among other things, including a devil's advocate in the panel's review procedures and equipping the panel with a small intelligence agency to "spy on the spies."

My hope is that this dissertation contributes not just to identifying and providing carefully reasoned answers to the ethical questions raised by government spying, but also to creatively envisioning the institutional arrangements that will best ensure that governments spy ethically. To realize these hopes, however, I first need to determine precisely what it means to spy. This is the task of the next chapter.

Part I

The Concept of Spying

Chapter 2

The Concept of Spying

What do you think spies are; priests, saints and martyrs? They're a squalid procession of vain fools, traitors too; yes; pansies, sadists, and drunkards, people who play cowboys and Indians to brighten their rotten lives. (Le Carré 1963, 203)

Spies are an inherently vicious bunch, according to Le Carré's famous protagonist Alec Leamas. But are they? Are the vices Leamas mentions necessary features of spies? Must one who spies also lie and cheat? Must the spy be morally wicked? Which behaviors and intentions, if any, are inherent features of the spy?

In this chapter, I attempt to answer these questions by defending a conception of spying. I intend for my conception to be useful for normative analysis but also to track very closely the way the word is used in ordinary language. Accordingly, the purpose of the chapter is to work through both why my definition is useful and how it corresponds to the way people normally use the word. I tend carefully to the concept of spying here because in the chapters that follow, when I turn to the normative analysis of spying, it will be crucial to have one stable conception. To my knowledge,

this is the first attempt by a political theorist or a philosopher to thoroughly defend a conception of spying.¹

I argue for the following conception of spying:

A person or collective agent, A, spies on a person or collective agent, B, if and only if she collects information that relates to B and intends to conceal her information collection from B.

I reject prominent definitions of spying that insist that the spy's intentions must be hostile to those she observes. I also repudiate accounts of spying that require that the spy violates the reasonable expectations of those she observes. The only requisite intention for a spy is that she intends to conceal her information collection from those she observes. Finally, I argue that by drawing on the idea of collective intentionality, my conception can account for a wide range of agents often thought to be spies, such as defectors and moles.

The chapter proceeds in four parts. I first say a few words about my method of conceptual analysis. I then examine how "spying" is used in ordinary language and put forward my own conception. In section three I take up five objections to my definition: it cannot account for defectors and moles, it is too capacious because it counts secret observation as spying even when people's reasonable expectations of

¹Allen (2008) proposes a conception but does not defend it at any length. I shall draw on her proposal below. Many conceptions have been proposed by international legal scholars, in international treaties, and in government documents. But these tend to be very narrow conceptions, tailored to international relations and the purpose of separating soldiers from spies. The latter purpose is of the utmost importance because it has been generally thought acceptable to execute captured spies (See e.g. Chesterman (2011, 26-27). Here is an example from the British War Office (1894, 313): A spy, in the military sense, is a person who is found in a district occupied by the enemy, collecting secretly and in disguise, information respecting his conditions and designs, with a view to communicating such information to the opposing force. Secrecy and disguise are the essential characteristics of a spy in a military sense. An officer in a uniform, however nearly he approaches to the enemy, or however closely he observes his motions, is not a spy, and if taken must be treated as a prisoner of war." (313) A definition like this one may be fine for its particular purposes, but it cannot serve as a general conception of spying. More examples can be found in the Appendix.

privacy are not being violated, it is too narrow because it fails to count instances of intrusive overt observation as spying, it wrongly includes as spying instances of secret observation that do not result in the collection of secretive information, and finally it fails to include concealment as a necessary condition of spying. The fourth section compares spying to the related concept, espionage, and the final section is a short conclusion.

2.1 A Few Words on Method

Since Rawls's (1971, 5-11) *A Theory of Justice*, it has become common among political theorists and philosophers to distinguish between a concept and a conception.² A *conception* of spying provides a set of rules that indicate when a particular action counts or does not count as spying. Thus, any conception of spying would indicate whether a cow jumping over a stream, a dog barking, or a wife hacking into her husband's email account counts as spying. Conceptions seek to provide complete accounts of which acts in the set of all actual and possible acts count as spying. The *concept* of spying, alternatively, offers no rules or counting principles; it is simply the abstract idea of spying; it is the "thing" to which conceptions refer. The relationship between a conception of spying and the concept of spying is thus that the conception provides one possible explanation or interpretation of the concept.

There can, of course, be many interpretations of the concept of spying, meaning there are usually many conceptions. But there is only one concept of spying.³ So,

²The method that I defend in this section fits squarely in what philosophers typically call either the "traditional" or "classical" theory of concepts. For a more complete account of the classical theory and criticisms, see Laurence and Margolis (1999, Ch. 1)

³Or so my arguments below suggest. This is not to say that there are not multiple senses of the word "spy," only that some of these senses map to different concepts.

when someone analyzes the concept of spying, she defends a particular conception of spying against other conceptions, claiming, e.g., that one conception accords better with ordinary language or that it is better suited for normative or empirical analysis.

Conceptual analysis is distinct from most scientific inquiry. In one sense, however, it appears similar: while doing conceptual analysis one typically proposes conceptions which operate much like hypotheses and then one tests her conception by attempting to reject it with counterexamples. But conceptual analysis is distinct from scientific inquiry since in the paradigmatic scientific test, the scientist fixes a set of relevant concepts in order to investigate the relationship between concepts rather than the nature of the concepts themselves. By fixing her concepts, a scientist can rely on observation as the final arbiter of whether her hypothesis survived a particular test.

Although conceptual analysis does rely on observation, observation cannot be the final arbiter. If one wants to understand, for example, the concept peach, an important first step toward this understanding is observing to what people refer when they use the word “peach.” But examining usage alone is rarely sufficient, since usage is notoriously misleading. Sometimes people use one word to refer to two concepts, for example. In the case of “peach” if one observed how people use the word, most of the time one would note they refer to a fuzzy, pale orange, juicy fruit, but other times they refer to a particularly attractive person or an excellent thing (as in “my car is a peach”). Even when we can distinguish multiple concepts at play, usage rarely provides a determinate conception, i.e. an account of necessary and/or sufficient conditions that determine whether a particular piece of fruit counts as a peach rather than an apple or a pear.⁴ We are left with competing conceptions, i.e. competing

⁴Some (e.g. Stich and Weinberg (2001, 638-640)) think that searching for necessary and sufficient conditions may be a fool’s errand. There may simply be paradigmatic cases of the concept and then cases being more or less similar to the paradigm. Even if this turns out to be true for “spying,” we

ways to interpret the concept.

The way to defend a conception is typically by showing that what it does and does not count fits with widely held intuitions.⁵ If it happens to fail this test, however, this failure may not be grounds for rejecting it, since it may capture widely shared intuitions better than competing definitions. Usage is helpful for the task of fitting a conception to widely shared intuitions – given that people never use the word “peach” to refer to a striped animal or a breezy day, these are unlikely to be defensible conceptions – but it must be supplemented by testing conceptions against actual cases (ideally a very large and diverse set of actual cases). In other words competing definitions can be tested to see which of them captures best widely held intuitions about what counts (or does not count) as unmistakable instances of the concept. Eventually the analysis homes in on the best of the competing conceptions, but assuming it does not count perfectly all of the relevant intuitions, more tinkering can be done with the winning conception to include even more relevant intuitions. In this process the analyst also reexamines her assumptions about what counts as an unmistakable case of the concept. Ideally, the process repeats until the analyst reaches what Rawls (1971, 46-53) called “reflective equilibrium.”⁶

can only know that it is the case if we first try (and fail) to characterize the concept with necessary and sufficient conditions.

⁵Many have criticized the reliance of conceptual analysis on intuition. Empirical evidence (e.g. Swain, Alexander and Weinberg (2008) suggests that people’s intuitions are liable to err and can be influenced by the order in which cases are presented. Our intuitions also seem to be conditioned by our moral beliefs. See Knobe (2006). Melnyk (2008, 267) concludes that intuitions cannot give us “a priori knowledge of necessary truths.” However, as Jackson (1998) argues, people tend to have pretty solid intuitions about the way words are used; otherwise they would have trouble communicating with one another. Similarly, Austin (1956, 8) argues that attempting to describe ordinary language may be a way to generate useful concepts; as Fallis (2009) says “it would be very surprising if humans had developed terms like ‘knowledge’ and ‘lying’ [and ‘spying’] but these terms were not getting at important phenomena in the real world.”

⁶Also see Daniels’s (2003) very helpful article in the *Stanford Encyclopedia of Philosophy*.

In the next section I follow this testing process to generate a conception of spying. I first examine the ordinary usage of the verb and noun form of “spy” by looking at a number of definitions in the Oxford English Dictionary (OED). Then by probing, cutting, and amending the dictionary definitions, I present a conception of spying. Finally, I test this conception against a variety of actual and hypothetical cases and compare how my conception performs against other plausible conceptions. My purpose to briefly recapitulate is to present a conception that is consistent with ordinary language, but that is also useful for normative analysis.

2.2 Towards a Conception of Spying

According to the OED there are a number of different senses in which people typically use the verb and noun forms of “spy.”⁷

Spy, v. (trans.)

1. To watch (a person, etc.) in a secret or stealthy manner;⁸
2. To keep under observation with hostile intent;
3. To make stealthy observations in (a country or place) from hostile motives;

⁷In the discussion below I use “sense” and “definition” interchangeably to refer to the numbered components of the OED definition below. I rely heavily on the dictionary definition of spying as a foil in this chapter for two reasons. The first and most important reason is my desire that my conception accord essentially with the way “spying” is used in ordinary language. The second reason is that most previous definitions have been of the narrow sort discussed in footnote 1. I have tried to engage the definitions that this is not true for throughout the text.

⁸I have renumbered the definitions for convenience of exposition. I rely heavily on the dictionary definition of spying as a foil in this chapter for two reasons. The first and most important reason is my desire that my conception accord essentially with the way “spying” is used in ordinary language. The second reason is that most previous definitions have been of the narrow sort discussed in footnote 1.

4. To (seek to) discover or ascertain by stealthy observation.

Spy, n.

5. One who spies upon or watches a person or persons secretly; a secret agent whose business it is to keep a person, place, etc., under close observation; esp. one employed by a government in order to obtain information relating to the military or naval affairs of other countries, or to collect intelligence of any kind.

It is helpful, both to further the understanding of the concept of spying and to develop a plausible conception, to examine the components of these senses. There appear to be four distinct components: (a) the subject (who or what is spying), (b) the object (upon what/whom the spying is done), (c) the action (what the subject does), and (d) the intentions of the subject (the purpose of the subject). I shall attend to each of these.

None of the senses above indicate who or what can engage in spying. Individuals obviously can spy, and we might plausibly think that collective agents can spy too. Certainly animals and perhaps even machines might be added to the list of possible spies, but at the outset of this chapter, I mentioned that while I do intend my analysis to track closely to ordinary language, I also intend for it to be useful for normative analysis. Even if bringing in non-human agents would make my account more consistent with all of the ways “spying” is used, it would come at a cost for the normative analysis in later chapters.⁹

⁹What costs, one might ask? Consider the claim that spying violates the target’s autonomy thereby undermining her status. This claim raises the question: if the spy were a non-human animal or a machine, would the target’s status still be harmed? Answering this question may be an interesting intellectual exercise, but there is considerable opportunity cost involved with answering it (and other questions like it). Thus, for the remainder of my argument, I shall consider the subjects of spying to be humans or collective agents.

So individual and collective agents can spy. Who or what can be the object of spying? One should first note that spying is always done *on* someone. But there is some confusion about this point in ordinary language. According to the five OED senses above, persons and agents can be the objects of spying, as one would expect, but the definitions also allow for many other kinds of objects. The fifth definition, for instance, suggests that a place could be the object of spying (“keeping a place under secret observation”); while the other four senses of the verb seem to leave the object of spying open entirely.

I think the confusion originates from the fact that there are often two distinct objects that are rarely distinguished: the “object of observation” and the “object of spying.” The object of observation is the object the spy directly observes. In the statement, “The American agent secretly took pictures of Soviet missiles,” the objects of observation are Soviet missiles, since the missiles are what the agent looked at and took pictures of.

Contrast the object of observation with the object of spying. If the relevant intention of the spy is to conceal her information collection (a point I argue for below), then the object of spying is the person or collective agent from whom the spy intends to conceal her information collection.¹⁰ In the case of the American agent taking pictures of Soviet Missiles, the objects of spying are the Soviets (or perhaps the Soviet Government).

It is sometimes the case that the object of observation and the object of spying are the same. When the FBI tracks the movement of a drug dealer, the drug dealer is both the object of observation and the object of spying. She is the object of observation because the FBI observes her directly to collect information, and she is the object

¹⁰In the remainder of the dissertation, I also refer to the object of spying as the “target” or the “target of spying.”

of spying because the FBI intends to conceal from her the fact that it is tracking her movements. But often the two objects are not the same. When the FBI examines the drug dealer's credit card statements, the credit card statements are the object of observation and the drug dealer is the object of spying.

It should be clear from the examples offered so far that the object of observation can be nearly anything – human or non-human. The object of spying, however, must be an agent. This claim is supported by a simple argument. When A spies on B, A intends to conceal her information collection from B. If spying requires intentional concealment, it follows that inanimate objects cannot be objects of spying. The reason is that it makes no sense to talk about concealing the collection of information from a mountain, a meadow, or the plans for a bomb. One can only conceal something from someone or something that has the capacity to observe.

When the object of observation and the object of spying are not the same, they must nevertheless bear a certain relationship. To understand the importance of this relationship, imagine a boy who watches his neighbor play video games every night without intending to conceal this from his neighbor. So far there is no reason to think that the boy is spying. But now suppose that the boy methodically conceals his watching from his parents (by locking his door, etc.), perhaps because his parents disapprove of video games. The components of spying now seem to be met: the boy collects information and he intends to conceal this fact. But these facts are not sufficient to constitute spying because it is not clear how, if at all, the object of observation (the neighbor) and the object of spying (the parents) are related. If there is no relation, as can be stipulated in this case, then it seems no spying has occurred. Thus observation with the intent to conceal is consistent with, but insufficient to constitute, spying. Spying requires that the object of observation and the object of spying be related in the right way.

But how does one know when the two objects are related in the right way?¹¹ Here it is helpful to draw on Anita Allen’s (2008, 3) definition of spying. Allen says, “To ‘spy’ is secretly to monitor or to investigate another’s beliefs, intentions, actions, omissions, or capacities, especially as revealed in otherwise concealed or confidential, communications and documents.” Allen’s conception suggests that the right relation between the two objects is that the object of observation relates to the “beliefs, intentions, actions, omissions, or capacities” of the object of spying.

As a statement of the right relationship, this is plausible. In the case of the American agent taking pictures of Soviet missiles, for example, the objects of observation (the missiles) relate to the capacities of the objects of spying (the Soviets). Similarly, in the case of the FBI and the drug dealer, the object of observation (the credit card statements) relates to the actions of the object of spying (the drug dealer).

But there is one concern with Allen’s idea of the right relationship between the two objects. The concern is with information relating to the physical condition of the object of spying. Not all information relating to the physical condition of a person fits easily under any of Allen’s categories. But there is reason to think that intending to secretly collect information about a person’s physical condition should count as spying. To see this, consider the following example. If Andrew secretly acquires Brandon’s past blood tests and learns that Brandon is HIV positive, has Andrew spied on Brandon? Most people (myself included) are inclined to say “yes.” Yet Andrew does not seem to be secretly monitoring or investigating Brandon’s beliefs, intentions, actions, or omissions. Nor does he seem to be secretly monitoring or investigating

¹¹The task here is analogous to attempts to define “personal” information in the privacy literature. Many, e.g. Fried (1970), conceptualize privacy as the control one has over her personal information. But critics, e.g. Schoeman (1984, 3) often claim that the notion of the personal loads into the idea of privacy a normative content. “It presumes privacy is something to be protected at the discretion of the individual to whom the information relates.”

Brandon’s capacities: HIV when properly managed does not hinder in any significant way most of a person’s capacities. Thus, it seems the right relationship derived from Allen’s list is too narrow. Fortunately Allen’s list can be fixed with a simple addition. Hence the object of observation must relate to the beliefs, intentions, actions, omissions, capacities, *or* the physical condition of the object of spying.

So spying must be done *by* human agents (individual or collective) and *on* human agents (individual or collective). Further, when the objects of observation and spying are distinct, they must bear a certain kind of relationship. Let us now consider the acts that constitute spying.

The root verb in OED (1) is “to watch” while the root verb in OED (2)-(5) is “to observe.” Watching with one’s eyes alone is obviously too narrow to capture all the instances people think of as spying. So, the word “watch” must be being used colloquially in (1) to mean something more like “observe.” Since one can spy without watching with one’s eyes, an observation, regardless of the sense organ from which it originates, could be an instance of spying.¹²

Observation can be done either directly or indirectly, i.e. the spy can observe her target directly or she can (indirectly) observe information related to her target, by for example reading her files, databases, diaries, or medical records, or by interrogating her friends or relatives. To make this point is simply to reiterate an idea just developed. Spying is direct when the object of observation and the object of spying are identical, otherwise it is indirect.

Since observation always results in information collection, one need not talk about

¹²In the literature on national intelligence, information collection is often divided into methods of collection, or “disciplines.” These disciplines typically range over all of the senses and include human intelligence (HUMINT), signals intelligence (SIGINT), measures and signature intelligence (MASINT), open source intelligence (OSINT), and geospatial intelligence (GEOINT). Any of these disciplines could count under my conception of spying if they are carried out with the requisite intentions. For an overview of the disciplines, see Lowenthal (2009) and Richelson (2012).

spying as a disjunctive of two activities (observing *or* collecting information). Yet while the spy always gathers information, she need not collect *new* information, that is information previously unknown to her.

Thus far I have assumed that the spy must intend to conceal her observation from the observed. It is now time to make good on this claim. To do so, I must show that A's intention to conceal her observation from B is the only necessary intention to constitute spying. I first need to make two points of clarification, however.

The first is that intentional concealment does not entail successful concealment. If Carl, the postman, reads Diane's mail and intends to conceal this fact from her, it does not follow necessarily that Diane has no knowledge of Carl's act. Carl may be very sloppy concealing his acts. The relevant fact, on my conception, is whether Carl *intends* to conceal his reading of Diane's mail from Diane, not whether Carl has actually concealed his act from Diane.

A second point of clarification is that the spy need not intend to conceal her information collection indefinitely. The spy may (and often will) intend at some time in the future to reveal his spying to the observed. So if at t_1 Carl intends to conceal from Diane that he is reading her mail, and at t_2 he intends to continue to conceal that he read her mail, but at t_3 he intends to reveal to Diane that he read her mail, Carl is still spying on Diane. The fact that Carl intends at some future time (t_1+n) to reveal to Diane his information collection has no impact on whether Carl spies at t_1 . Carl only needs to intend to conceal his information collection at t_1 .

Notice that it would be deeply counterintuitive to suggest that, were Carl to open and read Dian's mail in front of her, without in any way intending to conceal his actions from her, that those actions would constitute an instance of spying. Spying requires an intention to conceal the relevant observations, then. But is this intention the only necessary intention to constitute spying?

A number of the OED definitions suggest that spying requires hostile intent.¹³ The second OED definition, for example, defines spying as “observing with hostile intentions;” and the third definition echoes this suggestion when it mentions “hostile motives.” Further, according to the Hague Regulations, one of the first international treaties codifying the laws of war and war crimes, “A person can only be considered a spy when, acting clandestinely or on false pretenses, he obtains or endeavors to obtain information, in the zone of operations of a belligerent, with the intention of communicating it to a hostile party.” Is hostile intent necessary for spying?

It does not seem necessary. Consider the suspicious mother who secretly listens to the calls of her son. Her intentions need not be hostile. On the contrary, her intentions may be noble: she may suspect that he is buying drugs or planning to skip school and her motive may be simply to prevent him from harming himself or others. It would be odd to say that the mother did not spy on her son, simply because her intentions were not hostile. Intending to conceal that she is listening to her son’s calls seems sufficient in this case to call her acts spying.

I have now examined the four components of spying – subject, object, action, and intention – and have attempted to crystallize a few of the inherent features of spying. I argued (1) that spies can only be persons or collective agents; (2) that there are two different objects of spying, the object of observation and the object of spying and (3) the latter must be a person or a collective agent; (4) a subject engages in spying when she collects information related to the object of spying, and (5) the subject has the intention of concealing her information collection from the object of spying.

These conclusions lead to the following conception:

¹³A number of definitions also require passing information to an enemy or hostile party. See the Appendix.

A person or collective agent, A, spies on a person or collective agent, B, if and only if she collects information that relates to B, *and* she intends to conceal her information collection from B.¹⁴

On my conception, then, spying is the conjunction of an action and an intention. The action is collecting information related to an agent, and the intention is to conceal this collection from the same agent.

2.3 Objections

In 1994 federal officials arrested Aldrich Ames and charged him with spying for the Soviet Union and later Russia. Ames, who quickly became one of America's most notorious spies, was eventually found guilty of espionage. Because the information that he passed to the Soviet government led ultimately to the death of a number of covert agents, Ames could have been charged with the death penalty. But he eventually received lifetime imprisonment.

Those, like Ames, who occupy positions with access to sensitive information and who then volunteer to or are recruited to secretly pass information to third parties (usually rival governments), are typically called "defectors in place" in the literature on national intelligence.(Richelson, 2012). Defectors in place can be distinguished from moles, who are typically recruited to secretly pass information before they occupy positions with access to sensitive information. Kim Philby, who was recruited by Soviet agents in the 1930s long before he ascended to the highest echelons of Britain's MI6, is perhaps the best known mole.

¹⁴My conception is similar to those of other intentional acts. For example, popular definitions of lying often include the condition that the speaker intends to deceive the listener. See Bok (1999, 13) and Williams (2002, 96).

So Ames and Philby are two of history's most infamous spies, yet at first glance it seems that my conception of spying cannot properly account for either of them being a spy.¹⁵ The reason is that, for the most part, neither Ames nor Philby intended to conceal from the United States or Britain respectively that they were collecting information. Both had privileged access to the information they gathered and both acquired most of the information that they eventually passed to the Soviets in plain view. If spying requires that A intends to conceal her observation from B, then it seems that on my conception neither Ames nor Philby should be called a spy, since most of their observation was done overtly.

I shall ultimately show that this objection is mistaken and that my conception deals with these cases without difficulty. But I first want to assume that the objection is decisive and consider a potential solution to it. It may seem bizarre to assume that an objection is decisive when it isn't. But my own thinking vacillated for some time between the alternative conception I am about to introduce and the position I defend. Accordingly I think it is useful to alert the reader to the strengths and flaws of both approaches.

How can one ensure that Ames and Philby are rightly counted as spies? One way is to add to my proposed conception another set of acts. As it is stated, my conception counts as spying only those acts which feature A collecting information related to B and A intentionally concealing her information collection from B, but perhaps one should also include those acts when A intentionally conceals the *purposes* of her information collection from B. The revised conception would thus read as follows:

¹⁵On Vattel's (1883, 375) definition of spying, it seems, only moles are spies. He says, "Spies are those who introduce themselves among the enemy to discover the conditions of his affairs, penetrate his designs, and communicate them to his employers."

Alternative Conception: A person or collective agent, A, spies on a person or collective agent, B, if and only if she collects information that relates to B and intends to conceal her information collection from B or the purposes of her information collection from B.¹⁶

How would this alternative conception help deal with the cases of Ames and Philby? By adding the concealment of purpose to the conception, both the Ames case and the Philby case count quite obviously as spying. Both Ames and Philby openly gathered information about their home governments, including those governments' intelligence assets and capabilities, but both concealed their purposes with the information – to aid or advantage the Soviet Union – from their home governments.

So the alternative conception has the advantage of easily accounting for some of the best known cases of spying. But let me raise what I think is a decisive objection to this alternative conception.

The problem with the revised conception is that it counts as spying a fairly large set of cases that do not seem to fit with the way people use “spying” in ordinary language. The cases all feature A intending to conceal the purpose of her information collection from B, when this fact alone does not seem sufficient to count the cases as spying according to the way people tend to use the word. Consider the following case:

Bob is conducting a poll for the Congressperson for whom he works. But he fears that the content of some of the questions will rub some of the Congressperson's constituents the wrong way, so he invents a pretext. He

¹⁶Bentwich (1910, 243), a legal scholar working over a century ago proposed the following definition of “spy.” “The necessary differentiation of a spy is that by clandestine acts or false pretenses he obtains information with the purpose of communicating it to the enemy.” (My emphasis) Bentwich's use of “false pretenses” here makes his conception similar, in important respects, to the alternative conception. Bentwich's definition is very similar to Article 29 of the (1907) Hague Regulations.

tells his respondents that he works for a local news agency and the poll is for a set of articles they plan to publish.

Has Bob spied on his respondents? There are a number of ways that one might describe this case. One could say that Bob deceived his respondents, that he manipulated them, and perhaps, if one takes Bob's obligations to his respondents very seriously, that he betrayed them. But most people would not say that Bob spied on his respondents. Nevertheless, according to the alternative conception all of the conditions seem to be met. Bob collects information from his respondents and he intends to conceal the purpose of his information collection from them.

Accepting that Bob is a spy is, I think, an unacceptable entailment for a conception of spying. The revised conception is thus too capacious. A better way is required to count defectors in place and moles, then, one that doesn't also count cunning pollsters.

The best way, I think, is to stick with my original conception, but I shall have to do some work to show why the original conception accounts for these cases. The claim I want to advance is that moles and defectors in place count as spies under the original conception of spying because they are members of a collectivity that is spying. Accordingly, both Ames and Philby are rightly called spies because they were members of a Soviet collectivity spying on the United States. Neither Ames nor Philby intended to conceal *his* information collection, but both participated in a collectivity that did intend to conceal *its* information collection. Developing this claim requires a short digression on collective intentionality.

One acts intentionally when one acts purposefully, i.e. when one directs one's act

toward some objective or purpose.¹⁷ I act intentionally, then, when I ride my bike to school or when I type this essay. But not all intentionality is individual intentionality. Collective intentionality obtains when a set of people act with a shared objective or purpose. For example, every Monday during the summers I meet with a few of my colleagues to discuss recent works in political theory. We don't assemble by accident; we share the aim of meeting to discuss a particular article at a particular time in a particular place.

Not all groups possess collective intentionality. Sometimes people are related for reasons that have nothing to do with shared purposes. There is typically no collective intentionality, for example, in a group of motorists stuck in traffic on the highway.¹⁸

It is useful to distinguish four types of collective intentionality, since each type reveals a way that we talk about spying when the spy is a collectivity (see Table 2.1). The four kinds emerge from two crosscutting distinctions: simple/harmonic and organizationally sanctioned/not organizationally sanctioned. Consider the first distinction. Simple collective intentionality involves individuals working together toward a single purpose, when each individual's purpose is identical to that of the group. In the case of spying one can imagine two undercover detectives staking out what they believe to be a drug house. The collective intentionality of the detectives is "We spy on the people in the drug house," while each detective's individual intentionality is "I spy on the people in the drug house." The only difference between the collective's and the individuals' intentionalities is thus that the subject of the former is plural

¹⁷Beliefs and attitudes may also be intentional. But my concerns with intentionality are almost all with intentional actions. For more complete accounts of intentionality, see Bratman (1999) and Searle (1983, 1992, 1995).

¹⁸The motorists may have similar purposes (e.g. I need to get out of traffic and get to work) but their aims are not framed collectively (e.g. "we need to get to traffic and get to work"). Going forward, I shall distinguish between a set of people with collective intentions (a "collectivity") and a set of people without collective intentions (a "group").

(“we”) while the subject of the latter is singular (“I”).

One can distinguish simple collective intentionality from what I call “harmonic” collective intentionality. Although harmonic collective intentionality also features individuals working toward a shared purpose, this time individuals play different roles in achieving their collective purpose; hence, some individuals possess dissimilar (individual) intentionalities.¹⁹ To return to the example of the two detectives, one can imagine that one of the detectives positions himself to eavesdrop on a conversation inside of the drug house, while the other detective remains in his car, in case the two have to execute a speedy getaway. Here the collective intentionality of the detectives is to spy on the people in the drug house, but only one of the detectives individually intends to secretly observe the drug house. Nevertheless, it is just as appropriate in this case to say that both of the detectives spy on the drug house as it was in the first example when they both secretly observed the house. The reason is that the detectives were engaged in a collective endeavor the purpose of which was to spy on the drug house.²⁰

Simple/harmonic is the first distinction, then. This distinction is the only one required to meet the mole/defector in place objection, but a second distinction is helpful because it makes sense of many of the claims people make in ordinary language about spying. The second distinction is whether the spying has been sanctioned by a collective agent or not.

¹⁹An early description of harmonic collective intentionality comes in Aristotle’s *Politics* (1996, 1276b: 20): “Now, sailors have different functions, for one of them is a rower, another a pilot, and a third a look-out man, a fourth is described by some similar term; and while the precise definition of each individual’s excellence applies exclusively to him, there is, at the same time, a common definition applicable to them all. For they have all of them a common object which is safety in navigation.” “Safety in navigation,” in this example, is the harmonic collective intentionality of the rowers.

²⁰I do not mean to suggest that responsibility for the spying should be shared equally. The roles that individuals play in a collective endeavor certainly matter for how responsibility is partitioned.

To illustrate this distinction, suppose that the abovementioned detectives are members of the city police department. Decisions, we'll assume, are made in this police department entirely hierarchically – the police chief calls the shots and the police officers (including the detectives) carry out her orders. Now suppose that the police chief orders the detectives to spy on the drug house, and imagine they carry out this spying just as we described in the example of harmonic collective intentionality above (one observes while the other stays in the getaway car). This is a case of harmonic collective intentionality sanctioned by a collective agent. After the deed is done, only one detective secretly observed the people in the drug house. But one can say, sensibly, that both detectives spied on the people in the drug house and furthermore that the police department spied on the drug house.

Table 2.1: Four Types of Collective Intentionality

	Simple	Harmonic
Not Collectively Sanctioned	Two detectives secretly observe the drug dealer Spies = both detectives	One detective secretly observes the drug dealer; the other drives the getaway car Spies = both detectives
Collectively Sanctioned	Police department sanctions the two detectives to spy; both detectives secretly observe the drug dealer Spies = the police department, both detectives	Police department sanctions the two detectives to spy; one detective secretly observes the drug dealer while the other drives the getaway car Spies = the police department, both detectives

My analysis of collective intentionality thus helps make sense of two important things. First, it makes sense of many of the most common claims about spying involving collective agents, for example, “The KGB spied on Chinese scientists;” “The FBI spied on civil rights activists;” or “The St. Louis police department spied on

the biker gang.” But more importantly collective intentionality illuminates just how defectors in place and moles count as spies, even when they do not individually intend to conceal their information collection. Defectors in place and moles are both members of collectivities that spy. In the case of Ames, he participated in a collectivity, which also included the KGB (and later the Russian Federal Security Service) and which intended to conceal from the United States and the CIA in particular that it was collecting information about U.S. intelligence capabilities. Ames individually did not spy on the United States, but he participated in a collectivity that did.

I shall now briefly address four more objections to my proposed conception. The first objection is that an agent’s intention to conceal her observation is not sufficient to make her a spy – the spying must also violate the target’s reasonable expectations of when and where information related to her is being collected. I shall borrow (and tweak) one of Judith Jarvis Thomson’s (1975) examples to try to make this objection clear.

Imagine a couple having a loud argument inside of their house with their windows open, and a stranger walks by their house and hears the fight. Rather than moving on, the stranger stops in the street to listen. So far by my conception, one would not call the stranger a spy. But suppose further that he positions himself so he can’t be seen by the couple through the windows. Is he spying now? The objector says “no.” Her reason is that the couple has no reasonable grounds for believing, given that their windows are open, that their fight should not be listened to. The couple has neither signaled with socially acceptable marks (like closing their windows) that their conversation is “off limits” to someone like the stranger, nor has the stranger

made any extraordinary attempts to listen in on the argument (such as training a powerful sound amplifier on their house).

I think the objector is mistaken. The moment the stranger positioned himself so he could not be seen by the couple – i.e. the moment he intended to conceal his observation – he became a spy. To say otherwise is counterintuitive, since it would suggest that it is nearly impossible to spy on people in places where they expect they may be under observation: in public places, for example. But people are frequently spied on in public places. Take the case of police officers tailing suspects. If they take measures to conceal from the suspects that they are following them, then they are spying, and often they are doing so in public places. Hence reasonable expectations about when and where a person is under observation seem irrelevant to determining whether spying has taken place.

The objection nevertheless points to one important way to distinguish which kinds of spying most people find more or less objectionable. It seems plausible that the kinds of spying that violate a person's reasonable expectations (of privacy, non-observation, etc.) are more likely to harm her interests. She is less likely to suspect that she has been watched or that her information has been collected and, as a result, she is more likely to be victims of future threats or coercion. The spy who violates a person's reasonable expectations also cheats her target twice: first by concealing her collection of information related to her target and, second, by "free riding" on a social convention that the target and others obey. For these two reasons spying that violates reasonable expectations might be called more objectionable than spying that doesn't. But the latter case still counts as spying.

The second objection holds that A's intentional concealment of her observation from B is not necessary to count observation as spying – violating the target's reasonable expectations, when the violation is egregious, can be alone sufficient. Here

one can imagine a university installing cameras in faculty offices and declaring to the faculty “we’re watching your every move.” The university in this case, one can stipulate, is not concealing its observation. But one can safely say that the university is violating the faculty’s reasonable expectations about when and where they should be observed.

Is the university spying? The objector says “yes,” but I disagree. The university’s observation is observation of a particularly intrusive kind, a kind that has become all too familiar in a world in which closed circuit cameras, tracking software, and digital voice recorders are cheap and easily acquired. But it is not spying. Admittedly, intrusive observation of this sort is sometimes called “spying.” But I think this usage is inapt, since one can think of all kinds of observation that are deeply intrusive which most people would never call spying – the teacher staring at her student writing his exam, the creepy old man ogling the young girl or boy, the clerk glaring accusatorily at her customer. What makes a particular kind of observation spying is not that it is intrusive, or that it violates people’s reasonable expectations of privacy, or even that it is normatively objectionable; what makes it spying is that the person conducting the observation intends to conceal her observation from the person she is observing.

The third objection to my conception of spying concerns the kind of information collected by the observer. It holds that the information collected must be secret. Kim Philby (in Hitz (2008, 15)), for example, once defined spying as the collection of “secret information from foreign countries by illegal means.” Recall that my conception does not stipulate anything about the character of the information collected – the information could be found in the library, the phonebook, or other public sources. To dispel this objection, consider the case of a police officer on a stakeout. He has thoroughly bugged the house of a suspect and he spends long days listening to the suspect’s conversations. But suppose that while the officer listens to the subject he

hears only superficial conversations about sports or music. The suspect reveals nothing that he wouldn't willingly reveal publicly. Could one really say in this case the police officer has not spied on the subject, simply because he has not collected any secret information? I do not think one could, at least not without greatly altering the way people use "spying" in ordinary language.

This objection, much like the reasonable expectations objection, seems to arise from mistakenly thinking aspects of spying that are normatively relevant are also conceptually relevant. Spying to collect more secretive information is probably more objectionable than spying to collect less secretive information because of the likely harms that ensue.²¹ But the secretiveness of the information collected has no bearing on whether an act counts as spying.

A final objection to my conception claims that there are no requisite intentions for spying. Rather, what matters is that the observation is *in fact* concealed from the person observed.²² This is the weakest of the four objections considered in this section, because it clearly includes cases one would never describe as spying. Suppose Brian walks by Carter's office and overhears him hatching a plan to subvert Darnel's upcoming promotion. Brian then mentions this to Darnel. Darnel confronts Carter, revealing that he learned about Carter's plan from Brian. If Carter accuses Brian of spying, is the charge a correct one? It is not. People regularly observe all kinds of things unintentionally and often these observations are not apparent to those whom

²¹Anita Allen (2008: 3) seems to make the point that collecting more secretive information is more objectionable with her definition of spying. Spying, again, for Allen, is "to secretly monitor or investigate another's beliefs, intentions, actions, omissions, or capacities, especially as revealed in otherwise concealed or confidential conduct, communications, and documents." If the last clause means anything in this definition, it seems to mean that spying is more objectionable when the spying concerns "concealed or confidential" facts.

²²Allen's (2008) conception, which requires that the monitoring be done secretly, could be read as suggesting this objection.

the information we collect relates. The mere fact that people's observations are concealed does not make them spies. If it did, then many ordinary people spy repeatedly on a daily basis, an implication that does not seem plausible, given the way that people tend to use the word "spy" in ordinary language.

A more compelling version of this objection is the claim that while A's successful concealment of her observation from B is not a sufficient condition of spying, it is nonetheless a necessary one. When B knows that A is observing her, the objector claims, it is not correct to say that A is spying on her, even if A intends to conceal her observation from B. The objector might further claim that in this case one would say that A is attempting to spy, but because her concealment is not successful, she is not in fact spying. This objection is a stronger one; and it is not obvious that it is incorrect. Nonetheless, I think excluding successful concealment as a necessary component of spying better captures the way people use "spying" in ordinary language.

Consider a case in which Emily observes Fred and intends to conceal her observation from him. If Fred somehow discovers Emily's observation, most people would not think it inapt for Fred to say to a third party that "Emily is spying on me." Furthermore, accepting the opposite implication (i.e. accepting successful concealment as a necessary condition) leaves a host of imprecise alternatives to "spying." If Emily is not spying, what exactly is she doing? One might say Emily is "observing" or "surveilling" Fred, but both of these words leave a relevant detail out of the story: Emily *intends* to conceal her observation from Fred. Only the term "spying" seems to capture this detail. Even to say that "Emily is attempting to spy on Fred" does not seem an apt description of what Emily is doing because doing so suggests that Emily somehow cannot successfully observe Fred.

There is also the issue of partial concealment, which cuts against this alternative conception. Often a target knows (or suspects with a relatively high degree of

certainty) that she is being secretly observed, but she does not know how, when, or where she is being secretly observed. Imagine, for example, a police officer credibly saying to a suspect, “We’re watching you!” Here the suspect knows (or suspects with a high degree of certainty) that she is being secretly observed, but she does not know how, when, or where the police observe her. Now suppose that the police wiretap the suspect’s phone. Is it reasonable to conclude that she is not being spied on by the police because she knows (or suspects with a high degree of certainty) that the police are secretly monitoring her? I do not think it is. The conclusion that secret observation is not spying when it is partially revealed to the target has a host of implausible implications. For example, it suggests that it would be nearly impossible to spy inside of a totalitarian regime in which secret observation is widespread. Since inside such a regime nearly everyone would (reasonably) suspect that they were being secretly observed by the regime, the regime would be incapable of spying.

2.4 Spying and Espionage

Before concluding, I want to briefly reflect on the connection between spying and a related concept, espionage.

The OED provides the following definition of espionage:

Espionage, n.

1. The practice of playing the spy, or of employing spies.

Consider a few other proposed definitions:

2. To commit espionage, one must take several steps: procure National Defense Information (NDI), which is usually but not necessarily classified, either by stealing it oneself or by prevailing on an accomplice with access to steal it;

then make contact with a recipient of the information; and lastly transfer the information to the recipient. (Herbig, Wiskoff and Riedel, 2002).

3. ...gathering, transmitting, or losing...[information related to the national defense]. (Garner and Campbell Black, 1991).
4. The act of obtaining, delivering, transmitting, communicating, or receiving information about the national defense with an intent, or reason to believe, that the information may be used to the injury of the United States or to the advantage of any foreign nation. Espionage is a violation of 18 United States Code 792-798 and Article 106, Uniform Code of Military Justice. (US Department of Defense, 2007).
5. Industrial espionage refers to the clandestine acquisition of desirable business practices and/or technology by one company from another. By contrast, economic espionage refers to the clandestine acquisition of desirable business or government practices and/or technology by a foreign government or a company with the assistance of a foreign government. (Staples, 2007).

There are at least two senses in which “espionage” is used ordinarily, corresponding roughly to the two disjuncts in (1). The first sense is identical to spying as it has been conceived above. This claim might seem peculiar, given the many conditions definitions (2)-(5) feature which my conception of spying did not, but let me briefly suggest why I think it is nonetheless the case.

Definitions (2)-(4) all assert that the information collected, stolen, disseminated, etc. must be information relating to the national defense. Such a condition may be convenient for the purposes of characterizing a particularly serious crime, namely espionage against a state, but it renders any conception of espionage excessively nar-

row and thereby implausible. No conception with this condition could be plausible because the condition rules out all kinds of activities that people obviously count as espionage in ordinary language. If all espionage concerns national defense information, then one can make no sense of the idea of economic or industrial espionage (as we see in the fifth definition above). Nor can one make sense of espionage directed toward nongovernmental organizations (NGOs) or individuals. These are unacceptable entailments for a conception, given the way people use the word in ordinary language.

A more serious challenge to the idea that one sense of “espionage” is synonymous with “spying” is the claim – suggested by definitions (3)&(4) – that the concept of espionage includes more activities than those included in the proposed conception of spying in the previous section. Both definitions mention “transmitting” information relating to the national defense and definition (4) references “delivering,” “communicating,” or “receiving.” Do these activities rightfully belong in the concept of espionage? I do not think they do. The reason people tend to associate these activities with espionage is that they are commonly what individuals intend when they are a part of a collectivity engaging in espionage. Return, again, to the case of Aldrich Ames. Ames individually intended to transmit, deliver, and communicate information to Soviet (and later Russian) intelligence agencies; he also intended to receive information from these agencies. But none of these actions are sufficient to label Ames a spy or to say he engaged in espionage. Many American diplomats, for example, perform these actions without ever being considered spies. What made Ames a spy was that he was a member of a collectivity that intended to conceal its collection of information on American intelligence assets and capabilities from the CIA.

The second sense of “espionage” is using or hiring spies. Is this sense of espionage covered by my conception of spying? Usually it is. This claim may seem bizarre, since

it does not in general follow that one who uses or hires those who X thereby performs X. That I occasionally hire a painter, a landscaper, or a dancer does not seem to mean that I paint, landscape, or dance. But the case of spying is a peculiar one, since the one who employs spies nearly always does so to collect information about an agent with the intention of concealing her information collection from that agent. In other words, employing spies and spying tend to go together. Of course it does not follow that one who employs spies necessarily spies. One may employ spies not to gather information secretly but rather because it makes her feel powerful or wily, or because she wants to signal to others she is powerful or willing to break rules of fair play. But for her to be merely an employer of spies and not a spy herself she would have to refrain from examining the information she is paying her spies to collect. In these rare cases the second sense of espionage separates from my conception of spying.

2.5 Conclusion

I began the chapter with a reflection by John Le Carré's character Alex Leahmas, claiming that spies "are a squalid procession of vain fools, traitors too; yes; pansies, sadists, and drunkards, people who play cowboys and Indians to brighten their rotten lives." Although there may be a grain of empirical truth to this claim, I have shown in this chapter that there is no reason to think that the spy must be vicious in these ways. Indeed, my analysis shows there is no reason to think that the term "virtuous spy" is necessarily oxymoronic. Spying, again, requires only that some individual or collective agent, A, collects information related to B (another individual or collective agent) and intends to conceal this information collection from B.

I hope the conceptual brush has now been cleared, opening a clear path for the normative and institutional analyses that follow.

2.6 Appendix: Previous Definitions of “Spying”

Table 2.2: Some Previous (Non-Dictionary) Definitions of “Spy” or “Spying”

Source	Definition
Allen (2008, 3)	“To ‘spy’ is secretly to monitor or to investigate another’s beliefs, intentions, actions, omissions, or capacities, especially as revealed in otherwise concealed or confidential, communications and documents.”
Bailey (Cited in The Trial of John Beal 1865)	“One who clandestinely searches into the state and places of affairs.”
Bentwich (1910, 243)	“The necessary differentiation of a spy is that by clandestine acts or false pretenses he obtains information with the purpose of communicating it to the enemy.”
British Manual of Military Law (1894, 313)	“A spy, in the military sense, is a person who is found in a district occupied by the enemy, collecting secretly and in disguise, information respecting his conditions and designs, with a view to communicating such information to the opposing force. Secrecy and disguise are the essential characteristics of a spy in a military sense. An officer in a uniform, however nearly he approaches to the enemy, or however closely he observes his motions, is not a spy, and if taken must be treated as a prisoner of war.”
Bouvier’s Law Dictionary (Cited in The Trial of John Beal 1865)	“One who goes into a place for the purpose of ascertaining the best way of doing an injury there. The term is mostly applied to an enemy who comes into the camp for the purposes of ascertaining its situation in order to make an attack upon it.”

Source	Definition
Brussels Declaration (1874)	“By a spy is to be understood he who clandestinely or by illicit pretences enters or attempts to enter into places in the possession of the enemy with the intention of obtaining information to be brought to the knowledge of the other side.”
Hague Regulations (Art. 29, Par. 1)	“A person can only be considered a spy when, acting clandestinely or on false pretenses, he obtains or endeavors to obtain information, in the zone of operations of a belligerent, with the intention of communicating it to a hostile party.”
Leiber (cited in Geneva conventions)	“A person who secretly, in disguise or under false pretense, seeks information with the intention of communicating it to the enemy.”
Kim Philby (in Hitz 2009, 15)	One who collects “secret information from foreign countries by illegal means.”
Vattel (1883, 375)	“Spies are those who introduce themselves among the enemy to discover the conditions of his affairs, penetrate his designs, and communicate them to his employers.”

Part II

The Ethics of Government Spying

Chapter 3

Principles for Domestic Government Spying

My aim in this chapter is to derive a set of principles to regulate domestic government spying (that is a government spying on its own citizens in its own territory) from widespread intuitions about spying and about moral decision-making more generally. The method in this chapter is therefore both intuitionist and deontological: it is intuitionist because it appeals to widely shared intuitions to ground the moral rules it develops, and it is deontological because it fashions rules to regulate spying without examining in detail the likely consequences of spying.

Since my own philosophical commitments are consequentialist, I have strong misgivings about appealing to substantive moral intuitions to ground moral principles. But I recognize that many political theorists are more comfortable dealing in the currency of substantive moral intuitions than consequentialists are, and since my aim is to persuade a broader audience than just those that share my foundational commitments, in this chapter I make the best intuitive case for rules to regulate domestic government spying. In subsequent chapters I argue that the same principles that I

derive here can also be derived from utilitarianism (and no doubt other consequentialist theories), giving the principles a sort of double support, since utilitarianism is commonly thought to conflict with widespread intuitions.

I have two theses in this chapter. The first is that domestic government spying should be regulated by five principles: just cause, proportionality, minimization, necessity, and discrimination. The second is that these five principles should be institutionalized in such a way that those agents doing the spying (namely law enforcement and intelligence officials) do not alone determine how the principles apply.

A small set of widespread intuitions about spying and moral decision-making more generally support these theses: Spying is presumptively wrong, and it is particularly troublesome when it is on innocents. The presumption against spying, which varies in strength depending on the type of spying, can be overridden when there are strong countervailing reasons. Further, spying is only permissible when it is proportionate, and when it is the least harmful alternative likely to secure the good ends aimed at by the spy. Finally, the principles that government agents follow should be public, and citizens should be provided with reasonable assurances that government officials will follow them.

The chapter proceeds as follows. I begin with the presumption against spying and argue that it is best accounted for by widespread intuitions about respect for persons. In section two, I argue that this presumption against spying is not always of the same strength, but even when it is at its strongest – when the form of spying has the most potential for harm – it can still be overridden to prevent the violation of basic rights. When the presumption against spying is weaker, and thus the form of spying has less potential for harm, the presumption can be overridden to prevent violations of less basic rights. My arguments in the first two sections, I claim, are sufficient to justify the principle of just cause. In the third section I suggest that the widespread intuition

that innocents should have special protections against spying justifies a principle of discrimination, which holds that it is wrong to target innocents with spying. In section four I argue, from intuitions about avoiding excessive harms, that spying should be further conditioned by principles of proportionality, minimization, and necessity. In the final substantive section, I argue from the widespread intuition that the principles that regulate government officials should be public and from the intuition that the public should be given reasonable assurance that government officials will follow the principles they endorse, that the five principles of domestic government spying should be institutionalized. The final section is a short conclusion, which connects the analyses here to those in later chapters.

3.1 The Presumption against Domestic Government Spying

I first want to suggest that widespread intuitions support the idea that spying is presumptively wrong. When I say that it is “presumptively wrong” to spy I mean that it would be wrong to spy without a relatively strong moral reason. Spying is not a morally neutral activity. Government agents do not have a liberty (in the Hohfeldian sense) to spy when and where they please.¹ Unless they can give a moral reason to spy, they have an obligation not to spy.

Cases like the following, I think, provide support for this intuition.

You spot a police officer furtively peering through your neighbor’s window.

You approach the officer and you ask her why she is secretly investigating

¹For Hohfeld one has a liberty to X when one has no duty to not-X. See Hohfeld (1919).

your neighbor. She shrugs and says “you can’t prevent crimes you don’t know about.”

Most people would conclude (on the assumption that the officer is sincere) that her spying is wrong since she has no good reason for it. Peering into the windows of more or less randomly chosen people is unlikely to prevent crime – at least not efficiently. But notice that if we alter the case slightly so the officer provides a moral reason for her spying, it is less clear that her spying is wrong. Suppose, for instance, that she says, “we have good reason to believe a killer is hiding on this property.” Now the officer has provided a plausible reason for her spying; hence it may be justified.

This case provides support, then, for the intuition that spying is presumptively wrong. The officer must have a plausible moral reason for her spying, otherwise it is not morally permissible. What is the best way to explain this intuition?

One strategy is to point to the harms or injuries that follow from spying. The problem, however, is that we can imagine cases of harmless spying, especially if we stipulate that the spying will never be discovered. For instance, imagine that the officer in the above case snoops around the neighborhood, but his spying is never discovered. Suppose further that he finds nothing noteworthy, and he does nothing with the information he collects. Hence no one is harmed or injured by his spying. Using the harm/injury strategy, one is forced to conclude that the officer’s spying is not wrong. Nevertheless, I think many people still have the intuition that the officer’s spying *is* wrong, in virtue of the fact that he has no good reason for it.

So we need an explanation for the presumptive wrongness of spying that covers not just cases in which spying leads to harms or injuries. I think the best explanation that meets this criterion draws on the notion of respect for persons. But this explanation requires unpacking.

Let us first examine what is meant by “person.” As Rawls (2001, 19) notes, “The conception of the person itself is meant as both normative and political, not metaphysical or psychological.” The conception, in other words, is not meant to be merely descriptive. It is meant to pick out elements of human beings that are *normatively* relevant. It is meant, then, to serve as a building block for a certain kind of ethical and/or political theory.

For Rawls, there are two elements that are normatively relevant for personhood, what he calls the two “moral powers.” He says,

Moral persons are distinguished by two features: first they are capable of having (and are assumed to have) a conception of their good (as expressed by a rational plan of life); and second they are capable of having (and are assumed to acquire) a sense of justice, a normally effective desire to apply and to act upon the principles of justice, at least to a minimal degree. (Rawls, 1999, 442)

In short, moral persons can rationally pursue ends they deem valuable and they can cooperate on fair terms with others.²

It is still not obvious what it means to respect a person, or to show a person respect. If the principal normatively relevant feature of personhood is the capacity for rational agency, then respect for a person no doubt requires respecting her rational agency. Yet even this idea requires explanation.

Rawls suggests that showing a person respect boils down to giving her (or being

²Rawls’s conception of the person is clearly indebted to Kant, who held that the crucial normative feature of persons is their “humanity,” which for Kant meant their capacity to rationally determine their own ends. See Kant (1990, 56) and Korsgaard (1986, 330). Central to both Rawls’s and Kant’s conceptions is the person’s capacity to decide how to live and structure her life rationally.

prepared to give her) sincere reasons for actions that materially affect her, reasons that could be accepted from *her* point of view. He says,

Mutual respect is shown in several ways: in our willingness to see the situation of others from their point of view, from the perspective of their conception of the good; and in our being prepared to give reasons for our actions whenever the interests of others are materially affected . . . When called for, reasons are to be addressed to those concerned; they are to be offered in good faith, in the belief that they are sound reasons as defined by a mutually acceptable conception of justice which takes the good of everyone into account. (Rawls, 1999, 297)

As this passage suggests, for Rawls the test for whether the reasons that we offer could be accepted from another's point of view is whether they conform to "a mutually acceptable conception of justice." Rawls, of course, famously argued that the way to determine what constitutes such a conception is by way of a thought experiment called the "original position."

Throughout the history of moral and political thought, there have been a variety of competing suggestions about how to model impartiality, and thus a variety of different views about what kinds of reasons would be acceptable to others.³ I do not need to defend any particular method for modeling impartiality here. I will merely assume that one of these methods succeeds, and thus that it is possible for us to give others reasons that they could accept from their point of view. What I want to focus on is whether it is *possible* to give people reasons for spying on them that they could accept.

³Kant alone had three proposals for modeling impartiality (corresponding to his three formulations of the categorical imperative). Smith's (1982) impartial spectator and Hare's (1981) formulation of the golden rule provide further alternatives for modeling impartiality.

Two of the most prominent commentators on Kant – Christine Korsgaard and Onora O’Neill – have argued that according to Kant’s formula of humanity, “coercion and deception are the most fundamental forms of wrongdoing to others – the roots of all evil.”⁴ The reason is that when we coerce or deceive others, we treat them *merely* as means, not as ends.⁵ We treat another merely as a means when she cannot possibly assent to our way of acting toward her or she “cannot contain the end of this action in [her]self.”⁶(Ibid, 331)

If Korsgaard and O’Neill are correct and people cannot possibly assent to coercive or deceptive actions, then coercive and deceptive actions would not fail just Kant’s formula of humanity, they would fail *any* test that like Rawls’s test for respecting persons asked us to give (or be prepared to give) sincere reasons to others that they could accept for actions that materially affect them. If it is impossible for people to assent to coercive or deceptive actions, then it cannot be possible to give people reasons for our coercive or deceptive actions that they could accept.

I want to argue not only that Korsgaard and O’Neill are correct that people cannot assent to coercive or deceptive actions, but also that people cannot assent to being spied upon, although I will qualify these claims shortly based on two different kinds of assent.

Korsgaard points out that “[p]eople cannot assent to a way of acting when they are given no chance to do so.” When I coerce someone, she is by definition made

⁴Kant’s Formula of Humanity requires that actions be according to maxims that “treat humanity whether in your own person or in the person of any other never simply as a means, but always at the same time as an end.” (Kant, 1990) (Korsgaard, 1986, 333)

⁵In O’Neill’s 1985 words, “It is plausible to think that when we act in ways that would *always* preclude genuine dissent, we will have used others. For example, if we coerce or deceive others, their dissent, and so their genuine consent, is in principle ruled out.” (259)

⁶Korsgaard stresses that it is not enough that she would not or does not assent to our action, but that she *cannot*; it is impossible for her to assent.

to do something against her will; therefore she cannot assent to my action. When I deceive someone, in contrast, she may appear to assent to my way of acting, but she does so on false pretenses. She does not – indeed *cannot* – assent to my real way of acting because she does not know what that way of acting is. Suppose, for example, you agree to buy my car, in part because I tell you it has brand new brakes and a rebuilt transmission, when in fact the brakes are worn and the transmission shot. Here, you agree to purchase my car, but you do not assent to purchase it as it is *actually* offered, since you do not know of its actual state.

But suppose you do know of the car’s sorry state, and you accept the deal anyway. Even then it does not follow that you have assented to my offer. As Korsgaard notes, one of two things follows in such a case. Either you call my bluff and say “look I know you’re lying about the brakes and transmission, but I’m going to buy the car anyway,” or you *pretend* to accept my sketchy offer, knowing full well that the car is not in the state that I have claimed. Your knowledge of the car’s actual state, in other words, makes it impossible for you to accept my deceptive offer.

So, since people cannot assent to being coerced or deceived, coercive and deceptive acts are not compatible with respect for persons. Stanley Benn (1971, 10) extends this Kantian argument to include acts when a person “knowingly and deliberately alters [another’s] conditions of action, concealing this fact from him.” Hence suppose I want to prevent you from voting in a local election. Rather than lying to you about your polling place, I instead successfully petition to have your polling place changed and conceal this fact from you. Benn, I think rightly, suggests that you could not assent to my conduct and therefore I have not respected you as a person by acting this way.

Benn further argues that spying is precisely the kind of act that involves the spy changing her target’s conditions of action and concealing this fact from her. An

unknowing target of spying “is wronged,” Benn argues, “because the significance to him of his enterprise, assumed unobserved, is deliberately falsified” by the spy. “He may be in a fool’s paradise or a fool’s hell; either way [the spy] is making a fool of him.” (Ibid.) Since the spy intends to falsify his target’s reality, he does not respect the target as a person. Spying is therefore presumptively wrong.

The conclusion that we cannot assent to being spied upon may strike some as counterintuitive, however. They might point to cases in which someone agrees in advance to be spied on, but the agreement does not specify when, where, or how the spying will take place, leaving open the possibility that the spy can still successfully conceal her observation from her target.⁷ Suppose, for example, a homeowner in a crime ridden neighborhood consents to being spied on by his local police department. Officers, he agrees, can observe him and his property and conceal this fact from him and others. This case and others like it suggest that one can assent to be spied upon.

In response to this objection it is helpful to make a distinction between two kinds of assent: particular and general. *Particular* assent obtains when a person agrees to a specific treatment, immediately before it happens. The person is fully situated and she is asked to assess an action that is about to be performed, which will materially affect her. *General* assent, on the other hand, does not happen immediately before a specific act. Nor does it apply to one situated action. Instead it applies to a range of possible actions, and it occurs some distance in time before the actions that it concerns take place.

In principle it is possible to give general assent to *any* act. I might, for example, assent to be coerced, deceived, or spied upon if acting in these ways will save my life, or the lives of my family members. The arguments that I have been making in this

⁷Korsgaard (1986) takes up this objection for deception in footnote 6, but her response has not persuaded many commentators. See, e.g., Applbaum (1998, fn 25).

section, however, rely not on general but on particular assent. It is not possible, I have claimed, to give particular assent to being coerced, deceived, or spied upon, and this suggests that coercion, deception, and spying are incompatible with respect for persons. But someone might wonder why it is particular rather than general assent that matters.

The reason has to do with respect for persons. General assent may play a helpful role in determining when the presumptions against deception, coercion, and spying should be overridden. But it does not explain these presumptions. Particular assent, in contrast, provides one elegant explanation for why these three kinds of acts are morally problematic. Further, it is plausible to think both that to respect a person is to treat her in ways to which she can give particular assent and that we ought to treat people with respect.⁸ Kantians and many others have long argued that a society in which people treat one another with respect is an attractive normative vision, since in a society of mutual respect, people deal openly and honestly with one another; they take one another's interests and desires seriously; and they do not attempt to impose their own beliefs by force or fraud. Particular assent matters, then, since when we act only in ways to which people can give particular assent, we show them the respect to which they are morally entitled.

3.2 Overriding the Presumption

Thus far I have argued that there is a moral presumption against domestic government spying, deriving from widespread intuitions about respect for persons. This

⁸As O'Neill (1985, 259) says, "The morally significant aspect of treating others as persons may lie in making their *possible* consent, rather than in what they actually consent to or would hypothetically consent to if fully rational."

presumption is not absolute, however. Sometimes circumstances are such that it should be overridden. It is easy to imagine cases in which spying prevents catastrophic harms, for example. In these cases it seems intuitively obvious that government spying could be justified. The challenge is to specify the conditions under which the presumption should be overridden.

In contemporary moral and political discourse, the language of rights typically indicates strong moral reasons. If there are moral reasons strong enough to obligate others not to violate my privacy, for instance, then we say that I have a *right* to privacy. Further, the term “basic rights” typically indicates the strongest moral reasons.⁹ Finally, there is a fairly widespread but controversial intuition that basic rights are stronger than all other reasons. Basic rights can only be violated, many think, to prevent even graver or more widespread rights violations. As Rawls (1971) famously says, “Each person possesses an inviolability founded on justice that even the welfare of society as a whole cannot override . . . the rights secured by justice are not subject to political bargaining or to the calculus of social interests.”

Hence basic rights are the first place to look for reasons to override the presumption against spying, followed by ordinary rights. Before examining which rights – basic or otherwise – can override the presumption against spying, however, it is important to note that the presumption against government spying is not always of the same strength. Wiretapping a suspect’s cellular phone, for instance, intuitively requires a stronger reason than eavesdropping in a public place. But how does one explain the intuition that the strength of the presumption against spying should vary?

One way would be to make reference to more and less disrespectful forms of spying. Since I have already argued that spying is disrespectful, it seems natural to think

⁹I will not attempt to distinguish basic and non-basic rights here, since this has been done by a number of theorists. See, e.g., Rawls (1971) and Shue (1996).

that some types of spying are more disrespectful than others. This path may be a plausible one to follow, but I think a better way, given the way people tend to think about the wrongs of spying, is to bring in the potential harms of a particular type of spying. The presumption against wiretapping a person's cellular phone is stronger than the presumption against eavesdropping on a person in a public place, then, because, in most cases, the potential harms of wiretapping a person's private line far outstrip the potential harms of eavesdropping on them in public.

Later, in chapter 6, after I consider the consequences of spying in chapter 5, I offer a way of distinguishing more and less harmful forms of spying. Here, I will proceed more impressionistically, and merely assume that some forms of spying are more harmful than others, and that we can readily distinguish these forms.

It seems reasonable to think that the presumption against even the most harmful forms of spying should be overridden when spying could prevent the violation of a basic right. If wiretapping a person's cellular phone could, for instance, prevent a person from being wrongfully maimed, killed, or kidnapped, then the wiretapping seems (*prima facie*) justified.¹⁰ Wiretapping a person's cellular phone does not seem justifiable, however, to prevent violations of rights that are not basic. Landowners, for example, have a right not to have their land trespassed upon. But it does not seem reasonable to think wiretapping would be justified to prevent trespassing.

When the form of spying under consideration is relatively harmless, the reasons for overriding the presumption against spying can be weaker than basic rights. Nothing approaching the prevention of a basic rights violation is required, for example, to justify a government agent eavesdropping or covertly watching citizens in public places. Hence officers may be justified spying on suspected shoplifters in a department store.

¹⁰I add "*prima facie*" parenthetically here, since, as we shall see, justified spying must meet more conditions than just a strong reason.

For every form of spying, then, there will be a set of reasons weighty enough to justify overriding the presumption against spying. I will borrow a term from just war theory and call these reasons (to override the presumption against spying) “just causes.” A government agent can be said to have a just cause for spying when she has a set of reasons weighty enough to justify overriding the presumption against spying.

3.3 Spying on Innocents

It is not just the invasiveness of a particular instance of spying that affects its justifiability, however. Many people have the intuition that innocents require special moral treatment. Some think, for instance, that it is worse to spy on innocents than non-innocents. Others think that people forfeit or waive their right not to be spied upon when they engage in wrongful activity. How should we make sense of these intuitions?

It is helpful to begin by saying a few words about what I mean by the term “innocent.” By “innocent” I mean someone who is neither planning, engaged in, assisting, nor planning to assist the wrongful act that government agents seek to prevent with spying. Hence one does not lose one’s innocence merely by committing (planning, assisting, etc.) a wrong. The wrong has to be the particular wrong that government agents seek to prevent in their pursuit of a just cause.¹¹

It is implausible to think that innocents should have *absolute* protection against spying, since we can envision realistic cases in which spying on innocents could pre-

¹¹The line between innocents and non-innocents is not always a bright one. This point has been developed extensively in the just war literature, where theorists have long argued about just how much a person must assist a war effort in order to be considered a combatant (non-innocent). Should those who manufacture munitions count as combatants, for example? What about those who supply soldiers with food or other seemingly innocuous supplies? But, despite the existence of difficult cases, in most cases it will be clear who counts as an innocent and who does not.

vent catastrophic events. Consider, for example, the following case. Law enforcement agents have reason to believe that a terrorist is about to bomb a public square, but they do not know which public square. They also believe that the terrorist may contact his family members in the hours before his attack revealing his plans. Officers, further, have no reason to think the terrorist's family will provide them information voluntarily.

It is difficult to see, in a case like this one, how it would be wrong for the law enforcement agents to secretly listen in on the terrorist's communications with his family members, even though the family members are not engaged in wrongdoing. Government spying, in this case, has a nontrivial chance to prevent a disastrous loss of life. Absolute protection for innocents against spying therefore cannot pass basic intuitive tests, since it insists that sometimes the heavens must fall in order for justice to be done.

To rule out absolute protections for innocents against spying, however, is not to rule out *some* extra protections for innocents. There are a variety of strategies one could employ to craft these protections. The one that I want to explore here is to prohibit the *targeting* of innocents. In order to develop this strategy, it is helpful to draw on a popular principle in deontology, the doctrine of double effect (DDE).

DDE was first formulated by Thomas Aquinas in his *Summa Theologica* (II-II, Qu. 64, Art.7) to explain why it is permissible to kill an attacker in self-defense. One can permissibly kill an attacker, Aquinas argued, only if one does not *intend to kill* the attacker. Thrusting a sword at an attacker, for example, likely has two effects: thwarting the assault and killing the assailant. But the thrusting is only morally permissible when the latter effect is not intended. In Aquinas's words, "Nothing hinders one act from having two effects, only one of which is intended, while the other is beside the intention . . . Accordingly, the act of self-defense may have two effects: one,

the saving of one's life; the other, the slaying of the aggressor.”

Deontologists draw on DDE in a number of contexts. In the context of war, for example, it is used to explain why it may be permissible to strategically bomb enemy soldiers or facilities, even when the bombing is likely to kill innocents, and why it is impermissible to intentionally bomb civilians, for instance, in order to terrorize or weaken the resoluteness of the enemy. In Hurka's (2005, 36) words, "The discrimination condition does not forbid all killing of civilians. It concerns only targeting and therefore allows the killing of noncombatants as a side effect of force directed at properly military targets, or as "collateral damage." "

It is straightforward to adapt DDE to the case of spying. Government agents sometimes covertly observe conversations or communications with or records containing the information of both innocent and non-innocent parties in order to attain their just causes. According to DDE, this spying is permissible only when government agents do not intend to spy on the innocent parties, they intend only to spy on non-innocents. In order to clarify this point, it is helpful to recall the distinction I made in Chapter 1 between the object of observation and the object of spying. The object of observation is the person or thing that the spy directly observes, while the object of spying is the person from whom the spy intends to conceal her information collection. DDE implies that spying on innocents is permissible only when the innocents are the object of observation, but never when they are the object of spying.

Return to the case above. The intention of the law enforcement agents is to collect information on the terrorist's plans or his whereabouts. The object of spying, then, is the terrorist, and the object of observation is his family. Since the terrorist is likely to contact his family members in the run up to the attack, the law enforcement agents can justifiably monitor these communications. The law enforcement agents could not, however, justifiably target the family members. Their spying is only justified when it

is likely to produce information related to the terrorist.

So with the help of DDE we can make sense of the idea that innocents should have some – but not absolute – protection against government spying not afforded to non-innocents. Borrowing again a phrase from just war theory, I will call the resulting principle – that it is not permissible to target innocents – the principle of “discrimination.”

3.4 Harms - Proportionate and Minimal

In Aquinas’s abovementioned discussion of self-defense in which he develops the doctrine of double effect, he goes on to argue that self-defense is not unconditionally permissible. It is conditioned by proportionality, that is, one cannot use means to protect oneself that far outweigh the harms that are likely to ensue from remaining defenseless. In Aquinas’s words, “And yet, though proceeding from a good intention, an act may be rendered unlawful if it be out of proportion to the end. Wherefore, if a man in self-defense uses more than necessary violence, it will be unlawful, whereas, if he repel force with moderation, his defense will be lawful.”

It seems clear that the intuition underlying proportionality in self-defense extends to the case of government spying. Government agents may be permitted to covertly watch suspected shoplifters in a department store. They would not, however, be permitted to tap the phones, read the emails, or collect the mental health records of random shoppers to identify suspected shoplifters, since such actions would not be proportionate to the end of averting theft.

It is important to note that the determination of proportionality for spying is not the same as the complete calculation of the net benefits of spying that would be performed by a consequentialist. Proportionality, as it is typically conceived by

deontologists, demands that the costs of pursuing a just cause not exceed the benefits that flow from attaining that just cause. Benefits unrelated to the just cause do not count in the proportionality calculus. The person determining whether it is permissible to kill an attacker in self-defense, for example, does not get to take into account that the attacker is a foreman who disrespects and demeans his employees and that with his death the employees are likely to get a more compassionate manager.

We now have three principles to regulate domestic government spying: just cause, discrimination, and proportionality. I now want to argue for two more principles: necessity and minimization. Let me start with minimization. Consider the following case.

You are a police officer trying to prevent an armed robbery. You have strong reason to suspect the robbery will take place, but you do not know where or when. One way to find out, you think, is to read the suspects' emails, tap their phone calls, and hack their smart phones to track their locations. But you also think a simpler strategy is likely to be successful: placing a bug in their hideout.

I think most people have the intuition that it would be wrong for you to start reading the suspects' emails and tapping and hacking their phones before you try to bug their hideout. The first strategy would likely lead to considerable over-collection, and a good deal of the information would be about innocent parties. You have an obligation, most people think, to select the least harmful form of spying likely to secure the just cause. You have an obligation, in other words, to *minimize* the likely harms of spying.

Often, as is probably the case in the example above, selecting the least harmful form of spying means trying less harmful forms first. Once they have failed, more

harmful options can then be tried. Sometimes, however, less harmful forms of spying can be ruled out, since they are unlikely to be successful *and* their failure is likely to rule out trying more harmful forms of spying. If a government agent has reason to believe that tailing a suspect will not garner the information she requires to attain her just cause, for instance, then she may look to forms of spying that are potentially more harmful, such as hacking, wiretapping, or stealing information. But she cannot justifiably look to these more harmful forms of spying without first ruling out less harmful kinds.

The final principle – the principle of necessity – draws on the same general intuition I drew on to support minimization: that the harms of attaining a just cause should be minimized. The principle of necessity, however, does not ask government agents to compare the harms of different kinds of spying likely to attain the just cause, instead it asks government agents to compare the least harmful form of spying likely to attain the just cause against alternatives to spying that could also successfully secure the just cause. Government agents are asked, in other words, to consider whether spying is *necessary* to attain the just cause they seek. Since most people think it would be wrong for a government agent to spy, when she could secure the same information by consulting her files, by reading the newspaper, by searching the internet, or by conducting interviews, this principle is on strong intuitive ground.

The principles that I have argued for in this chapter can now be summarized as follows:

1. Just Cause = Spy only when there are reasons sufficiently strong to override the presumption against spying.
2. Proportionality = Spy only when the harms from spying are not likely to outweigh the benefits from the just cause sought.

3. Necessity = spy only when less harmful alternatives have been tried or they can be ruled out as unlikely to be successful.
4. Minimization = Spy only when the tactics utilized are the least harmful tactics likely to secure the just cause and hence all reasonable precautions have been taken to minimize harms.
5. Discrimination = Spy only when the principal target of the spying is reasonably believed to be engaged in or assisting the harm that government agents aim to prevent in securing the just cause.

3.5 Institutionalizing the Principles

What I want to establish in this final section is that the five principles that I just argued for should be institutionalized. By “institutionalized” I mean that those agents doing the spying, namely law enforcement and intelligence officials, should not alone determine how the principles apply. The power to determine when, where, and how spying takes place should be institutionally distributed. Note that my argument in this section is not for *how* spying should be institutionalized, but merely *that* it should be institutionalized.

Institutionalization is supported by two wide-spread intuitions. The first is that the rules that govern a society’s central institutions, what Rawls called the “basic structure,” should be public. In Rawls’s (1999, 48) words, publicity demands that “[a] person taking part in an institution knows what the rules demand of him and of the others. He also knows that the others know this and that they know that he knows this, and so on.”

The grounds for publicity are numerous.¹² Rather than developing them in detail, I want to highlight the pivotal role publicity plays in the arguments I made above concerning respect for persons. Recall that respect for persons demands that we give (or that we are prepared to give) reasons for our actions that affect others that *they* could accept. In a society in which persons respect one another, people deal openly and honestly with one another. They act only in ways to which others can give particular assent.

Government agents in such a society clearly could not operate according to secret principles, since it is not possible for citizens to give particular assent to principles that are concealed from them. Respect for persons thus demands open governance. Of course not all of the details of the government's business must be public knowledge, but the principles by which the government operates must be open to all. (Bok, 1989; Luban, 1996) Citizens must be able to challenge or defend the principles by which they are governed.

The five principles to regulate domestic government spying that I argued for above should be public, then. Citizens should know (or at least be capable of knowing) the principles government agents use to determine when, where, how, and on whom to spy. But to say that the principles should be public is not to say that they should be institutionalized in the sense that I explained it above, since we can imagine public principles that are not institutionally distributed.

The argument for institutionalization requires a second step, then, a step which relies on the intuition that not only should citizens know the principles that govern their government's conduct, but they should also have reasonable assurances that the government will in fact follow these principles. It is one thing for the government

¹²See Gosseries (2005) and de Lazari-Radek and Singer (2010, Section 4).

to publicly commit to a set of principles. It is a whole other thing, however, for the government to make its commitment credible, that is to take actions to assure citizens that it will in fact conduct itself as promised.

For some issues in which policymakers have no prudential or political reasons to deviate from a policy as circumstances change, a simple declaration of a policy can serve as a credible commitment. But spying is not such an issue. If government officials promised to follow the five principles outlined above, even if they were trustworthy, citizens would have reason to doubt that they would faithfully follow them. Law enforcement officials, intelligence agents, and the government officials that orchestrate spy operations have strong incentives – personal and political – to sometimes deviate from the five principles. Hence, in order for citizens to be assured that government agents will follow the five principles, government agents must somehow tie their hands with institutions.

3.6 Conclusion

In this chapter I derived a set of principles to regulate domestic government spying from widely shared intuitions people have about government spying and about moral decision making more generally. I argued that spying should follow principles of just cause, proportionality, necessity, minimization and discrimination, and further that these principles should be institutionalized, meaning those agents doing the spying do not alone determine how the principles apply.

As I mentioned at the chapter's outset, however, I have strong misgivings about deriving ethical principles from substantive moral intuitions. One worry about such an approach is that the moral intuitions employed may conflict with other intuitions people have on different moral questions. Widespread moral intuitions are exceed-

ingly unlikely to fit into a coherent whole, at least not without a considerable degree of trimming, eliminating, and altering, and I have not indicated why the intuitions that I utilized in this chapter would survive this modification process. A deeper worry is that substantive moral intuitions cannot be trusted. Since these intuitions are no doubt strongly influenced by our upbringing, our socio-economic status, our education, etc., many theorists think they should not be given evidentiary credence. (Brandt, 1998, 1990; Hare, 1973)

Allaying these worries requires not just developing and defending a moral theory and then showing that the five principles above derive from it, but developing and defending a theory that, although it may not be able to avoid drawing on moral intuitions entirely, relies chiefly on non-moral intuitions and on abstract intuitions we have about moral deliberation. These tasks are far too demanding for a book, let alone a chapter. Hence my aim in the next chapter is considerably more modest.

Chapter 4

Two-Level Utilitarianism

My aim in this chapter is to describe, in some detail, the moral theory that I take to be the most plausible: two-level utilitarianism. Utilitarianism is commonly thought by philosophers and political theorists to cut strongly against widespread moral intuitions. Hence, if the same principles that I developed in the previous chapter can be derived from a utilitarian theory, then these principles stand on solid ground. One might say borrowing a phrase from Parfit (2011) that they are “doubly justified.”

Since fewer people are familiar with two-level utilitarianism than its simpler sibling, classical utilitarianism, I contrast the two theories, and attempt to show why the former escapes the most important criticisms of the latter. Since some of these criticisms cannot be rebutted until more work is done (in Chapters 5 & 6) to determine just what the two-level view demands, I postpone my responses until the end of Chapter 6.

By the conclusion of Chapter 6 I hope to persuade my readers that the principles I developed in the previous chapter are the ones government agents ought to follow for domestic spying. I also hope to persuade them that the two-level view is a plausible

one, but I recognize that some of my readers will not be persuaded that it is the best view. For the skeptics, I hope to have, at the very least, clearly identified areas of disagreement – the forks in the road, so to speak. I also hope, rather more boldly, to guide my skeptical readers toward those moral theories most likely to be reasonable alternatives to the two-level utilitarian view. Hence, regardless of my reader’s moral intuitions, I hope to clarify her thinking about the ethics of spying and carry it a few steps forward.

The chapter has four parts. In the next section, I outline the classical utilitarian view and raise a few important objections to it. One of these objections – that classical utilitarianism is self-defeating – is decisive, at least under one common but possibly mistaken interpretation of classical utilitarianism. In section three, I present the two-level view of utilitarianism, which separates the criterion of right and wrong from the decision procedures agents employ to promote utility. This separation makes utilitarianism a more psychologically sophisticated theory and permits two-level utilitarians to escape a variety of familiar objections against utilitarianism, including importantly the charges that utilitarianism is self-defeating, that it undermines the enjoyment of private goods, and that it fails to properly respect persons. The fourth section of the paper addresses a general challenge to two-level utilitarianism: that it is unstable. As I mentioned, I postpone addressing a few of the objections I raise in the first section of the chapter until the end of Chapter 6, after more has been said about the particular principles to regulate spying that two-level utilitarians should endorse.

4.1 Classical Utilitarianism

Classical utilitarianism is an elegant and plausible normative theory. In principle, it has a complete and consistent set of answers to all normative questions, from what kind of government we ought to adopt to whether and how much we ought to give to charity. Yet it relies on a relatively small number of assumptions - perhaps fewer assumptions than any theory that approaches its level of plausibility.

Utilitarianism emerged as a systematic theory in the 19th century with the work of Jeremy Bentham, although elements of the view can be traced back through the history of ethics. Some, for example, have wondered whether Plato's Republic is at root utilitarian. (Mabbott, 1937) Others find many of the elements of utilitarianism in the ancient Chinese philosopher Mo-Tzu.¹ Classical utilitarianism received its most careful articulation and defense in Sidgwick's *The Methods of Ethics*. But the core of the view did not change radically as it was passed from Bentham to Mill and then from Mill to Sidgwick. Moore (1903), a student of Sidgwick, was perhaps the first utilitarian to leave the classical tradition.² But it was not until the last half of the twentieth century, as the many varieties of consequentialism were distinguished, that classical utilitarianism took its place as only one member of a large family of consequentialist theories.

The core of classical utilitarianism (CU) is the claim that the morally right action is the one that produces the most good. But this formulation is deceptively simple, obscuring the many philosophical commitments embedded in CU.

Here are the formulations of CU by Bentham, Mill, and Sidgwick respectively:

¹See Scarre (1996, 27-33). Also quoted in Driver (2007, 41).

²He did so, of course, by questioning the hedonism of Bentham, Mill and Sidgwick.

By the principle of utility is meant that principle which approves or disapproves of every action whatsoever, according to the tendency which it appears to augment or diminish the happiness of the party whose interest is in question...I say of every action whatsoever; and therefore not only every action of a private individual, but of every measure of government. (Bentham, 1996, Chapter 1)

The creed which accepts as the foundation of morals, Utility, or the Greatest Happiness Principle, holds that actions are right in proportion as they tend to promote happiness, wrong as they tend to produce the reverse of happiness. By happiness is intended pleasure, and the absence of pain; by unhappiness, pain, and the privation of pleasure. (Mill, 1979, 278)

By Utilitarianism is here meant the ethical theory, that the conduct which, under any given circumstances, is objectively right, is that which will produce the greatest amount of happiness on the whole; that is, taking into account all whose happiness is affected by the conduct. (Sidgwick, 1981, 411)

It should be clear to careful readers that these three formulations are not logically equivalent and thus have slightly different implications in particular cases. On Sidgwick's formulation, for example, conduct is either objectively right because it produces "the greatest amount of happiness on the whole" or it is wrong simpliciter, whereas for Mill and Bentham judgments of right and wrong seem to be scalar in nature – acts are right "in proportion as they tend to promote happiness." Besides being inconsistent, these formulations are also rife with ambiguities. In order to avoid some of these inconsistencies and ambiguities, I shall define CU as follows:

CU is the view that the morally right action is the one that maximizes the sum total happiness minus unhappiness (“utility”) for all people counted equally.

This formulation is, I believe, more or less consistent with the spirit of the classical utilitarians. Although it is not entirely free from ambiguity, it should be precise enough for my purposes.³

What does CU entail about spying in particular cases? First, it entails that it is never right to spy if the same amount of happiness can be obtained with less costly means. If one can collect the same information and achieve the same results with less intrusive means than spying, for example, by simply requesting the information or by reading the news carefully, then one ought not to spy. Hence, the rightness of spying will often hinge on whether useful information can be collected without spying or whether people can be dissuaded from engaging in socially undesirable activities without the threat of being spied on.

Second, because CU is a consequentialist theory, only the consequences of spying matter for determining whether spying is wrong. For CU, the fact (developed in the next chapter) that spying entails that the spy intends to deceive her target does not necessarily mean that spying is wrong. The intention to deceive may lead to bad consequences, such as the spy developing a bad character or bad dispositions, but it need not. Indeed, sometimes the intention to deceive may be a *requirement* of the right act. Similar things could be said for other (allegedly) intrinsic features of spying, such as that it violates the target’s autonomy or privacy.

Third, because CU is a hedonistic theory, the costs and benefits of spying matter

³I do not discuss all of the commitments of CU, only those I take to be most important. In Appendix 1, I try to make it clear how many different philosophical commitments are implied by my articulation of CU. The figure illustrates eight commitments, but there are no doubt more. Sinnott-Armstrong (2011), for example, counts eleven distinct commitments of CU.

only insofar as they affect the calculation of sum total happiness minus unhappiness. Frustrating someone's autonomy (a typical consequence of spying) is bad, according to CU, only if and because it undermines their happiness. Similarly, saving someone's life is good only if and because it permits the person saved (and those whose lives she affects) to enjoy more happiness.

Finally, because CU is a cosmopolitan theory, everyone's net happiness matters, and each person's happiness matters equally. According to CU, it is thus wrong for spies to ignore or discount the interests of others; it is wrong to give special consideration (at the level of moral deliberation) to the interests of their colleagues, their family members, or their compatriots. Every consequence of spying, if it affects someone's happiness, must be counted and counted equally, regardless of whether the affected person is a stranger, relative, or friend.⁴

Precisely what CU requires in particular cases will often be exceedingly difficult to determine with any precision. But, in principle, the theory always provides an exact answer, regardless of our capacity to discover it. Further, whether CU entails that particular cases of spying are wrong will often hinge on indirect and/or remote effects. Indirect effects are those that affect individuals beyond the spy and her target, what are often called "externalities" in economics. Remote effects are those that occur with low probabilities. Finally, even when indirect and remote effects can be identified it is often difficult to affix precise likelihoods to their occurrences.

One might worry that the complexity of determining what CU requires in particular cases is a strike against it. But in many – perhaps most – cases, determining what CU requires will not demand a complete and precise rendering of costs and benefits.

⁴It does not follow from CU's cosmopolitanism that we cannot have special obligations to our friends, relatives, etc. I explore one strategy for deriving special obligations in a utilitarian framework in Chapter 7.

It will be clear that the balance tips in one direction or the other. It is also worth noting that the complexity worry does not plague just CU, but rather any theory that has a pivotal role for consequences, and arguably *all* plausible moral theories have some role for consequences. As Rawls (1971, 30) wrote, “All ethical doctrines worth our attention take consequences into account in judging rightness. One which did not would simply be irrational, crazy.”⁵

Without examining all of the implications of CU in particular cases, we can note that CU has a surface level plausibility as an account of the ethics of spying, since CU accounts for what are arguably the two most strongly held general intuitions that most people have about spying: that a good deal of spying is wrong and that some spying is permissible if not obligatory to keep people safe. A good deal of spying is wrong, according to CU, because it destroys happiness. Among other things, as I show in the next chapter, spying can undermine people’s liberal and democratic rights. It can lead, further, to people being embarrassed, humiliated, or disrespected. But, spying can also save lives and prevent a variety of other grave consequences. So, some spying is permissible if not obligatory to keep people and their interests secure.

But even if CU turns out not to account for all of our substantive moral intuitions, this need not be a decisive reason to reject it. Other theories may be no better at coherently accounting for these intuitions, and, more importantly, CU may fit better with our non-moral intuitions than other theories. The extent to which a moral theory coheres with our substantive moral intuitions is only one dimension on which to evaluate a moral theory, then, and it may be far from the most important dimension.

Nevertheless many people think that CU leads to deeply counterintuitive – if not downright implausible – conclusions in a range of cases. The seven most important

⁵See also 159-162, and Hardin (1990, 24).

objections are: (1) CU is self defeating; (2) it undermines the enjoyment of private goods; (3) it is not consistent with respect for persons; (4) it sometimes prescribes harming innocent parties; (5) it does not take autonomy seriously enough; (6) it often hinges the wrongness of an act on whether it is discovered; and (7) it condones highly counterintuitive social arrangements such as slavery or an extensive spy state in certain imagined cases.

In the rest of this section, I develop the first of these objections and show why the objection is damning *if* CU is taken to be a theory that identifies not just the criterion of right and wrong but also the decision procedure that particular agents should follow in order to maximize utility.

Let us turn to the objection that TLU is self-defeating. Here is the argument in standard form:

The Self-Defeating Objection:

1. The best ethical theory is the one that maximizes utility.
2. CU requires that agents calculate the costs and benefits of all of their available options to determine whether a particular action maximizes utility.
3. Agents frequently fail to correctly calculate costs and benefits, leading them to perform suboptimal actions.
4. Agents can perform fewer suboptimal actions by following a theory relying less on calculation than CU does.
5. Therefore, CU fails to maximize utility.
6. Therefore, CU is not the best ethical theory.

The argument is valid. Let us examine whether it is sound. The first premise is a deep commitment of CU, so we can suppose for the sake of argument that it is true.

If CU turns out not to maximize utility, then the theory fails on its own terms. It is a self-defeating theory. The second premise is true if we take CU both to indicate the criterion of right and wrong *and* to provide agents with the best decision procedure. For now, let's assume that this is the appropriate way to interpret CU. I will reconsider this interpretation below.

The third premise seems obviously true. Human beings are notoriously imperfect calculating machines. Nevertheless it is helpful to collect some of the reasons why this is the case. (1) We often do not have the time to make considered decisions; (2) we frequently lack the relevant data to make fully informed decisions; (3) we are notoriously bad at reasoning probabilistically; (4) we are prone to bend our calculations in favor of ourselves or those for whom we care most about; (5) we sometimes fail because of weakness of will to give our choices thorough consideration; (6) we tend to ignore or discount powerful precedential effects decisions have on our characters or dispositions; (7) we sometimes fail to follow the conclusions of our calculations because they turn out to be demanding psychologically or otherwise; and finally (8) we often struggle to predict what others will do, thereby missing opportunities to coordinate and cooperate.⁶

The fourth premise is the most controversial. To be true, it must be the case that agents can maximize utility by following some theory other than CU, a theory that relies less on agents calculating the utility that follows from the options available to them. Note that the alternative theory need not abandon calculating entirely, it only needs to limit it: it is perfectly fine if the theory has agents calculating some of the time so long as other times agents are restricted from making calculations.

It seems obvious that such alternative theories exist and plausible that one of

⁶Variants of these reasons can be found in Hare (1981), Mackie (1984), Parfit (1984), Brandt (1998).

them better promotes utility than CU. To see why this is the case, it is helpful to draw on a non-moral example: playing poker. I am not much of a poker player. I tend to get caught up in conversation, losing track of the game, or I have one too many drinks thereby losing the required lucidity to play well. When I do play poker well, I do not play like a savant. Poker is a complex game, and with my focus at the poker table and my rather mediocre capacity to do mental math, I rarely have the time to calculate with precision whether I should fold, check, or raise. If I decide that I should raise, I almost never have the time to calculate how large my raise should be. In order to play poker well, I need to follow a few rules that keep my calculations relatively simple. For example, don't consume more than a couple of drinks. Don't talk politics or philosophy (for either topic is sure to distract me from the game). Roughly calculate pot odds.⁷ Attempting to calculate precisely the expected value of every one of my actions would lead rather predictably to me losing all of my money, since I tend to botch calculations, especially when the numbers aren't round. Further, when calculations get complex – and when there are fascinating conversations to participate in – I tend to experience weakness of will, abandoning difficult calculations. Worst of all, when I abandon difficult calculations, I tend to engage in wishful thinking – betting when the odds are against me. Hence, to play my best poker, I must avoid the strategy of thoroughgoing calculation. I'm better served by following a few rules, featuring simplified calculations.

How is the non-moral case of playing poker analogous to deciding whether to spy (or to other moral decisions)? In both cases it is difficult to calculate reliable expected values, and time is often scarce, making the challenge of complex calculations beyond

⁷In poker, pot odds are the current size of the pot (i.e. what could be taken with the winning hand) divided by the cost of the call under consideration. Calculating pot odds is a shortcut for estimating the expected value of the call.

the reach of most people. If anything, the calculations facing the spy are considerably more complex than those facing the poker player. Unlike the calculations by the poker player, the calculations by the spy often rely, as I mentioned, on indirect or remote consequences. They also often rely on probabilities calculated from limited data.

Perhaps despite the two activities both being calculatively complex, we are nonetheless more likely to calculate the consequences of spying correctly than those of a decision in poker (because the stakes are higher, etc.), or perhaps there are no rules to guide our spying better than running continuous calculations as there are for poker. I see no reason to seriously entertain either of these possibilities. Moral calculation in difficult cases is just as likely if not more likely to go awry as decisions about difficult hands in poker. We follow rules because we are far from perfect calculating machines. When it comes to moral decision-making, as I mentioned, we have a tendency to bias calculations in our own favor – a problem that affects most of us less intensely when we play games like poker, since the effects of such self-deception are typically relatively immediate and undesirable. For spying, many of us follow a simple rule such as, “Don’t spy,” and there are good reasons to think that this rule (or other similar) rules lead us more often to do the right thing than if we continuously calculated whether to spy when the option was available. At the very least the rule protects against a certain degree of self-serving calculation, and it is simple enough to apply in a pinch.

So, it seems that premise four is true, which means the argument that CU is self-defeating is sound *if* my assumption about premise two is true. Let us return to this assumption. Is the best interpretation of CU one that sees it as providing both the criterion of right and wrong *and* a decision procedure? I think the answer is “no.” All three of the classical utilitarians seemed to appreciate that utility will often not be maximized by agents attempting to maximize utility. Bentham wrote, “It is not

to be expected that this process [of calculation] should be strictly pursued previously to every moral judgment.”⁸ Similarly, Mill wrote that it is a mistake to

confound the rule of action with the motive of it. It is the business of ethics to tell us what are our duties, or by what test we may know them; but no system of ethics requires that the sole motive of all we do shall be a feeling of duty; on the contrary ninety-nine hundredths of all our actions are done from other motives, and rightly so done, if the rule of duty does not condemn them. (1979, 289)

Finally, Sidgwick states the view most precisely:

The doctrine that Universal Happiness is the ultimate standard must not be understood to imply that Universal Benevolence is the only right or always best motive of action. For...it is not necessary that the end to which gives the criterion of rightness should always be the end at which we consciously aim. (1981, 413)

If CU does not require that agents always behave as utility maximizers, then premise two is false and the objection that CU is self-defeating does not succeed. But because my purpose is to explain what I take to be the most plausible moral theory, I do not want to pin my presentation to controversial interpretive claims about what the classical utilitarians believed. Thus, rather than continuing with CU, for the rest of the chapter I employ a neo-utilitarian theory that *explicitly* separates the criterion of right and wrong from the decision procedure that agents should follow. The new the-

⁸Bentham also writes, somewhat comically, “The principle of utility, (I have heard it said) is a dangerous principle: it is dangerous on certain occasions to consult it. This is as much as to say, what? That it is not consonant to utility, to consult utility: in short, that it is not consulting it, to consult it.” (1996, Chapter IV, Sec.VI).

ory, what I call “two-level utilitarianism,” following Hare (1981), I will now consider in more detail.⁹

4.2 Two-Level Utilitarianism

Two-level utilitarianism (TLU), like CU, holds that the maximization of utility is the criterion (or “standard”) of what is morally right. Acts are morally right if and only if they maximize the utility of all people counted equally. But TLU rejects the inference that because utility maximization is the correct criterion of right and wrong that it is also (always) the best decision procedure to guide agents’ conduct. Bales (1971) was the first contemporary philosopher to clearly articulate the two-level view. He wrote,

...an assumption apparently shared by act-utilitarians and critics alike, is that acceptance of the act-utilitarian account of right-making characteristics somehow commits one a priori to a particular decision-making procedure: the procedure of estimating and comparing probable consequences of alternative acts. This is an erroneous assumption...the account itself places no a priori restrictions whatever on the procedures we use to isolate that [utility maximizing] alternative. (263)

But TLU has roots reaching far beyond Bales. We have already seen that TLU was anticipated by if not endorsed by the classical utilitarians. Reaching even further back, Brink (1986) sees traces of TLU in Butler’s *Fifteen Sermons*, and Hare (1981:

⁹The theory goes by many different names. Pettit and Brennan (1986) refer to it as “restrictive consequentialism,” Railton (1984) as “sophisticated consequentialism;” Hooker (2008) calls it “partial rule consequentialism” and still others e.g. Alexander (1985) and Mason (1998) “indirect consequentialism.”

25) argues that the distinction between two levels stretches back to antiquity – to Plato’s distinction between knowledge and right opinion and to Aristotle’s distinctions between right motivation and practical wisdom, virtues of character and virtues of the intellect, and the “that” and the “why.”

TLU gets the label “two-level” because it distinguishes two levels of ethical analysis. Ethical thinking at the first level, what Hare (1981, 25-26) calls the “intuitive level,” is done by appealing to a decision procedure, composed of a set of rules that state “how agents should deliberate, reason, and make moral decisions.” (Brink, 1986, 424) Many of these rules are likely to be familiar, resembling the principles most of us are brought up to follow. Principles, for example, prohibiting (most) lying, cheating, stealing, breaking promises, and killing are almost certainly included in our optimal decision procedures.

At the intuitive level, agents faced with ethical decisions decide not by carefully calculating the consequences of available options but rather by following more or less uncritically the rules in their decision procedures. The rules composing decision procedures are thus not rules of thumb: they are not, that is, meant to assist agents in utilitarian deliberation.¹⁰ They are meant to be adhered to strictly. Hence, at the intuitive level two-level utilitarians behave just like deontologists who endorse an identical set of rules.

But two-level utilitarians have another level to appeal to in order to justify their decision procedures. This second level of analysis, Hare calls the “critical level.” Beyond justifying decision procedures, the critical level is also sometimes required to adjudicate conflicts between rules.¹¹ If my decision procedure, for example, includes

¹⁰See Hare (1981, 38) and Brink (1986, 425).

¹¹The critical level could also be required in cases when we face decisions where none of our rules apply (if such cases exist).

a rule prohibiting spying but also includes a rule instructing me to protect others from harm when doing so is relatively costless to me, then in some range of cases these two rules will come into conflict. Sometimes I may have a third rule in my decision procedure to appeal to, permitting me to quickly resolve this tension, other times, however, I may have to ascend to the critical level in order to adjudicate the conflict.

At the critical level, two-level utilitarians attempt to reason as it is often mistakenly thought utilitarians must always reason: they try to reason like ideal utility maximizers (like Smith's impartial spectator or what Hare calls "archangels"), attempting to foresee and weigh all of the consequences either of their available options in cases in which rules conflict at the intuitive level, or the consequences of adopting a particular rule when deliberating about the optimal decision procedure. All decisions at the critical level are made according to TLU's criterion of right and wrong: acts are right if and only if they maximize utility. Similarly, decision procedures are "optimal" if they guide agents to "make the highest proportion of right decisions in actual cases where their decisions make a difference to what happens – weighted, of course, for the importance of the cases, that is, the amount of difference the decisions make to the resulting good or harm." (Hare, 1979, 115)

For some agents the optimal decision procedures will determine nearly all of the moral questions that confront them. Optimal decision procedures, in other words, will be designed to prevent these agents from ascending to the critical level, where they are likely to make poor judgments. Nearly all of their ethical thinking will be performed at the intuitive level. For other agents, ascending to the critical level may be less rare. But even for the most sophisticated agents – those most adept at engaging in utilitarian calculus – ascent to the critical level will likely be uncommon and probably should be confined to moments of calm contemplation, since utilitarian

calculation, as we've seen, is a dangerous method.

Determining the optimal decision procedure for each agent is thus a complex empirical matter – a project suited as much to the social scientist as it is to the ethicist. We need to know, among many other things, a person's cognitive capabilities, her dispositions to fudge or abandon calculations, and the unique circumstances in which she is placed. *Ceteris paribus*, those who reason well probabilistically and who have strong focus and self-discipline will have more complex decision procedures.

But there are a few features that all agents' optimal decision procedures are likely to share. Agents' decision procedures, first, will likely be composed of simple rules. One reason for simple rules is that as rules become more complicated, they become more difficult and thus more costly to learn. As exceptions and qualifications are affixed to a principle, eventually it becomes impossible for anyone to remember the principle, let alone to employ it efficiently in real world circumstances.¹² Another reason is that simple principles permit the development of stable dispositions – dispositions crucial for making good decisions when time is scarce, and for lowering the likelihood that agents will engage in special pleading.

This reason suggests a second feature of optimal decision procedures. The principles that compose them will be deeply engrained: they will be stable dispositions, more or less fixed elements of agents' characters. The reason is made plain by Hare (1981, 47):

If we wish to ensure the greatest possible conformity to what an archangel would pronounce, we have to try to implant in ourselves and in others whom we influence a set of dispositions, intuitions, prima facie principles

¹²Hare (1981, 35) points out that we may be able to learn principles in a Rylean sense (Ryle, 2009), permitting us to “know” them without being able to articulate them precisely. But there are surely limits to what we can know, even in the Rylean sense.

(call them what we will) which will have this effect. We are on the whole more likely to succeed in this way than by aiming to think like archangels on occasions when we have neither the time nor the capacity for it.

The last few decades of psychological research lend credibility to Hare’s claims. Research increasingly supports a “dual process” model of human cognition, which posits (like TLU) that humans reason in two very different ways, one intuitive and automatic and the other critical and rational.¹³ Because a great majority of our decisions are made at the intuitive level, if we want to make good decisions most of the time, it is crucial that we “implant” good dispositions. Rules of thumb to aid utilitarian calculation at the critical level may be useful, but these rules will be employed with much less frequency than the deeply ingrained principles at the intuitive level.

A consequence of our principles being deeply ingrained is that when we violate them – even when we have decided that this violation is for the best – we will experience psychological discomfort in the form of regret, remorse, or guilt. This psychological discomfort is an aspect of TLU that has not been fully understood or appreciated by most political theorists. Walzer (1973, 171), for example, in his famous article on dirty hands, thinks the only reason that a utilitarian agent should feel guilty for breaking a moral rule when doing so maximizes utility is because such a feeling is useful for future utility promotion. But then, he wonders, how a utilitarian agent could feel guilty for what he knows is not wrong. He concludes, “It is best to say only that the more fully they [utilitarian agents] accept the utilitarian account, the less likely they are to feel that (useful) feeling.” (172)

But Walzer’s conclusion is mistaken. Two-level utilitarian agents feel psychological discomfort for breaking moral rules to the extent that these moral rules are deeply

¹³See Chaiken and Trope (1999). Thaler and Sunstein (2009, 19-21) provide a simple but compelling explanation of the view.

ingrained in their characters, not to the extent that they accept or doubt utilitarianism.¹⁴ It is true that the two-level utilitarian believes she has the dispositions and convictions she has just so she can better maximize utility, but this does not imply that she can turn her dispositions and convictions on and off at her whim, nor would she want to if she could. When she is faced with a conflict between two deeply held principles, she must choose. Two-level utilitarianism provides her a *standard* for making this choice; it does not give her absolution. Hence the theory explains both why she is uncomfortable about her choice and why it is a good thing that she feels this discomfort.

So, the principles that constitute our decision procedures will be simple and deeply ingrained, and this latter fact can explain certain elements of our moral experience that are often thought to be in tension with utilitarianism. Finally, intuitive principles will be imperfect, meaning they will sometimes lead us to perform wrong acts. But how, one might wonder, could a set of rules, which will sometimes get it wrong, be better than utilitarian calculation, which can in principle always get it right? The reason is that while *in principle* utilitarian calculation can identify the right action, *in practice* it will not – in practice utilitarian calculation will produce the wrong answer more often than a set of strictly followed rules, for the many reasons canvassed above. Maximizing utility is thus best served by strictly following the optimal decision procedure, even when this decision procedure will sometimes get it wrong.

It is on this point that two-level utilitarianism diverges most obviously from the more familiar rule utilitarianism. The rule utilitarian says an action is wrong if and

¹⁴Walzer mentions numerous times that utilitarianism will play almost no role in moral education. If by this he means the utilitarian criterion of right will not be the primary focus of most moral education, then he is almost certainly correct. But utilitarianism plays numerous other roles in moral education beyond inculcating its criterion of right. Most notably it guides the selection of optimal decision procedures, which are the centerpiece of moral education.

only if it is forbidden by the rules that if complied with would maximize utility. Following the rules, in other words, is *always* the right thing to do for rule utilitarians. For two-level utilitarians, in contrast, rules are merely decision procedures designed to assist us in performing the right actions as often as possible and particularly when it matters most. They do not *determine* the right actions.¹⁵

With TLU made more or less clear, let us return to one of the objections raised above – that utilitarianism undermines the enjoyment of private goods – and examine whether TLU helps rebut the objection. Here is the objection in a simple form:

The Private Goods Objection:

1. The full enjoyment of many goods rely on stable expectations that others are not spying.
2. TLU obliges agents to calculate the costs and benefits of spying and other available options to determine whether spying is the right action.
3. When agents decide whether to spy by calculating costs and benefits, others cannot develop stable expectations about being spied on.
4. Therefore, TLU undermines the enjoyment of “private” goods.
5. The best ethical theory will not undermine the enjoyment of “private” goods.
6. Therefore, TLU is not the best ethical theory.

This argument should look familiar to those acquainted with the legal scholarship on privacy, where “reasonable expectations” often play a vital role.¹⁶ The main thrust of the argument is that we cannot develop stable expectations about other people’s

¹⁵See Hooker (2008) for a very nice discussion of the differences.

¹⁶See, e.g., *Katz v. United States*, 389 U.S. 347 (1967)

spying when we cannot predict when and where they will spy. Further, knowing that others use the utilitarian calculus is not sufficient for us to predict their behavior, since there are so many considerations that go into utilitarian calculations and since people have such a wide range of proficiency identifying and applying these considerations.

It should be clear from the discussion thus far that the second premise of the argument is false. TLU *does not* oblige agents to calculate the costs and benefits of all of the options available to them. On the contrary, agents following TLU adhere to a set of rules justified on utilitarian grounds. Since two-level utilitarians recognize the value of the enjoyment of private goods, the optimal decision procedures will be designed to protect this value.

One could object that TLU still does not provide anyone with stable expectations. Because each person has different cognitive capabilities, takes on different roles, etc., some will have more permissive rules for spying than others. Stable expectations will not develop when we all follow different rules. But this objection is not persuasive. Many of us already have relatively stable expectations about spying, despite people in our society following different rules. We know, for example, that the rules police officers follow are more permissive than those doctors follow. Yet this difference does not impede us from developing stable expectations about when and where we are likely to be spied on. It is also not obvious that we would all follow different rules for spying, as the objection suggests. Often the optimal rules in our decision procedures will be optimal in part because they reflect social norms or laws that coordinate mass behavior.

Notice that the private goods objection hints at a reason for moving to a two-level theory that goes beyond the fact that we are fallible utility calculators. Indeed, there is a further *set* of reasons supporting TLU, developed in a series of papers by

Philip Pettit (Pettit and Brennan, 1986; Pettit, 1988, 1989). Pettit draws attention to certain *consequences* of consequentialist deliberation that tend to destroy important benefits. The most important of these consequences is others becoming aware of an agent's deliberation. Certain benefits can be realized, Pettit argues, *only* when agents do not know or suspect that others are engaging in consequentialist calculation. The enjoyment of private goods may be one such benefit. Pettit and Brennan mention further "the security which lovers or friends produce in one another by being guided, and being seen to be guided, by maxims of virtually unconditional fidelity." (Ibid: 450) But more important for the analysis of the ethics of spying they point to respect for persons. Respecting a person, they argue, involves acknowledging his rights, i.e. showing that you "regard certain of the claims he makes as privileged." (451) But behaving as a utility maximizer, and being seen as behaving like a utility maximizer, is antithetical to this acknowledgment.¹⁷ Hence a utilitarian concerned with realizing the value of respect "must forswear calculation and calculative monitoring in favour of the commitments...distinctive of respect for persons." (Ibid.)

If Pettit's reasoning is sound, then the objection to TLU on the grounds of respect for persons fails. If making decisions about whether to spy by calculating costs and benefits leads to disrespecting people, and respect is a value of significance, then two-level utilitarians have strong reasons – beyond the practical reasons against calculating canvassed above – to follow principles at the intuitive level that prohibit or limit calculation.

I have now addressed three of the objections raised against CU above – that CU is self-defeating, that it is not compatible with the enjoyment of private goods, and that it cannot show appropriate respect for persons – and I have shown that by

¹⁷The obvious objection is that the utilitarian should not acknowledge rights, he should (by using concealment and deception) act as if he acknowledges rights. See Pettit (1988, 53-55) for a response.

moving to a two-level theory, these objections can be avoided. In my explanation of two level-utilitarianism, I also suggested that the two-level view circumvents certain phenomenological problems with utilitarianism associated with the problem of dirty hands. Agents who endorse TLU *can* experience moral dilemmas when two intuitive principles come into conflict, and they will often experience guilt, regret, or remorse for breaking with their intuitive moral principles.

I have not yet addressed four objections to CU that I raised above. Since more needs to be said about precisely what principles a two-level utilitarian would endorse for spying, in order to meet these objections, I have postponed their consideration until the end of Chapter 6.

The objections I have raised thus far have all been to classical utilitarianism, and my strategy for meeting these objections has been to move to a two-level view. Whether this strategy is successful depends not just on the plausibility of the principles for spying that the theory produces, it also depends on the plausibility of the two-level theory more generally. Since some have doubted the two-level theory's plausibility, I respond to the most oft-repeated of these doubts below.

4.3 The Instability Objection

Some have argued that TLU is unstable.¹⁸ In Williams' (1988, 189-190) words the two-level model is unstable, since

it represents the intuitive responses as deeply entrenched, surrounded by strong moral emotions, sufficiently robust to see the agent through situations in which sophisticated reflection might lead him astray, and so on;

¹⁸Williams (1988), Alexander (1985, 1989), and Alexander and Moore (2012). See also Levy (1994).

and yet at the same time explains those responses as devices to secure utilitarian outcomes.

It is not plausible psychologically, Williams thinks, that people could view their intuitive principles as being merely instrumental. For to evaluate one's intuitive principles using utilitarian calculation would lead people to doubt the depth of their commitments to the principles. Once a person evaluates her principles through a utilitarian lens, her principles cease to be deeply ingrained, they cease to be intuitive. They become more like rules of thumb. As Alexander (1989, 824) puts it, "The central issue for the indirect consequentialist is whether it is psychologically possible for us to know *both* the justifications and the motives [for our non-consequentialist principles] and that they are different without undermining the rules and dispositions and thus the indirect strategy."¹⁹

But the instability objection, I want to suggest, is not persuasive. First, using an analogy with prudential reasoning, it seems clear to me that when I reflect on my intuitive (prudential) principles, I am not led necessarily to doubt them. On the contrary, my faith in my intuitive principles is often buttressed. When I reflect, for example, on the rules to guide my poker playing (as I often do in the days leading up to a big game) I often conclude that there is little I can do by way of tweaking my rules to improve my performance. Hence, if anything, my reflection deepens my commitment to my intuitive principles. If I want to play well, I need to stay true to my principles.

Admittedly, critical reflection *could* weaken my commitment to a principle, but

¹⁹Williams later (190-192) reformulates his objection, saying it is not merely a psychological claim but also a philosophical claim. But many have been confused by his reformulation. In his comments on Williams' essay, Hare says, "I cannot understand why Williams makes such heavy weather . . . of the combination of critical with intuitive thinking." Shaw (1999, 163) similarly comments, "it is difficult to pinpoint exactly what his [Williams'] objection is."

often this a good thing, since it is an indication that my intuitive principles could be improved. Even if critical reflection led me to doubt the optimality of my intuitive principle and it failed to lead me to a better principle, I am still unsure that my reflection would uproot my convictions, since I can still take comfort in knowing that following a suboptimal rule is likely to be better than both following no rule at all and than engaging in prudential calculation.

One need not rely on the analogy to prudential reasoning, however, to show that the instability objection is unpersuasive. For most people the practice of trying to live ethically involves not just instilling in themselves certain habits of action and mind, it also involves careful thinking about exactly *which* habits of action and mind they ought to be inculcating. As Hare writes,

It has always seemed to me that this [instability] objection . . . will not be sustained by anyone who has experience even of *trying* to live a morally good life. It is perfectly possible at the intuitive level to treat moral duty or virtue as ultimate . . . while at the same time to recognize that in order to establish that those traits of character really do constitute virtue, and that those moral principles really are the ones we should observe, requires more thought than the mere intuition that this is so.

Hence I conclude that the instability objection does not succeed. We not only can reason as TLU demands, many of us do. As Shaw (1999, 163) says, “Harean agents believe that they are fully justified in having the principles and moral feelings they do...[nothing is] unstable, incoherent, or inauthentic about the two-level model in either theory or practice.”

4.4 Conclusion

In this chapter I have tried to explain, and to a lesser extent defend, the moral theory that I find most plausible. Two-level utilitarianism holds that the morally right action is the one that maximizes the sum total happiness minus unhappiness (“utility”) for all people counted equally. But it rejects the utilitarian calculus as a decision procedure in favor of deeply ingrained, strictly followed principles. Agents should not make ethical decisions by calculating the costs and benefits of all of the options available to them, then. Instead they should strictly follow a set of well-trying rules justified by the utilitarian calculus.

Whether the principles that two-level utilitarianism entails in the case of government spying correspond to those I developed from widespread intuitions, however, is not yet clear, since I have not yet done the critical thinking required to generate intuitive principles for spying. The first step toward elucidating these principles is a careful reckoning with the consequences of spying. Accordingly, I spend the next chapter identifying some of the more important consequences that often result from spying.

4.5 Appendix: Utilitarianism's Commitments



Figure 4.1: The Commitments of Utilitarianism

Chapter 5

Some Consequences of Spying

The consequences of spying are incompletely understood in discussions in political and legal theory and are even less well accounted for in popular discussion. The harms of spying, for example, are often cashed out in terms of violations of people's privacy. (cf. Allen, 2008) But there is nothing nearing agreement in the literature on the definition of privacy. Consider just a few of the better known conceptions of privacy. For Warren and Brandeis (1890) privacy is "the right to be left alone;" for Fried (1970), Parker (1974), Moore (1998, 2003), and Westin (1968), it is control over (usually "personal") information; for Freund (1971) and Pound (1915), privacy is an extension of one's personality; for Feinberg (1983), privacy is having autonomy over personal concerns; for Parent (1983) privacy is the condition of "not having undocumented personal knowledge about one possessed by others;" finally, many (e.g. Thomson (1975) and Posner (1981)) have argued that there is no encompassing concept of privacy; rather there is merely a bunch of distinct ideas that have been (mistakenly or not) thrown together under the privacy rubric.

Further, even if theorists did settle on one conception of privacy, it is not clear that this conception would always count spying as a privacy violation. Does spying in

public places, for example, violate a person's privacy? What about spying that does not gather secretive or sensitive information? If the answer to questions like these turn out to be negative, then theorists must either accept that these kinds of spying are harmless, or they need to search for other consequences of spying.

It is useful to distinguish between two sorts of spying the consequences of which are often very different: successfully concealed spying and suspected spying. Suspicion here can be thought of probabilistically.¹ If the target has no reason to believe that she is being spied on, then she assigns a zero probability. Similarly, if the target has compelling reasons to believe she is being spied on, then she assigns a number near one to her suspicion. Of course suspicion of spying can be aroused even when there is no spying, but these cases are not my concern here.

Although spies can take precautions to ensure concealment, they can never be certain that their spying will escape detection and avoid arousing suspicion. As Epicurus is said to have argued:

It is impossible for someone who secretly does something which men agreed not to do in order to avoid harming one another or being harmed to be confident that he will escape detection, even if in current circumstances he escapes detection 10,000 times. For until his death it will be uncertain whether he will continue to escape detection. (1997, 35)

To spy is thus always to risk being suspected or exposed.

The chapter is organized into two substantive sections. I first identify the consequences of successfully concealed spying – to the spy, to the target of spying, and

¹The person doing the calculation will most likely be the target. However, as I will show below, spying will tend to affect the beliefs of people other than the target. I assume that those calculating suspicions follow basic rules of rationality. I am not concerned with the paranoid, i.e. those who believe they are being spied on but have no evidence to support this belief.

to others – and examine under which conditions these consequences are harms or benefits. Successfully concealed spying, I argue, deceives the target and it often leads to the collection of sensitive information, which can then be employed to harm or to benefit the target or others. It also tends to make the spy more likely to spy in the future.

I then turn to suspected spying, and again examine the consequences to the the spy, to the target of spying, and to others. When the target suspects she is being spied on, she will sometimes condition her behavior. Other times she will not alter her behavior, but experience a loss of enjoyment. The suspicion (or revelation) of spying can also lead to a host of emotional responses by the target and others, and in extreme cases it can undermine people’s status as equal citizens. Further it can undercut the trustworthiness of the spy, expose her to retaliation, and diminish public trust in government agents. The final section is a short conclusion.

The consequences that I discuss are summarized in the table below.

Table 5.1: The Consequences of Spying

	To the Spy	To the Target	To Others
Successfully Concealed Spying	*Spy becomes more likely to spy	*Spy uses sensitive information she collects to harm/benefit target	*Spy uses sensitive information she collects to harm/benefit others
Suspected Spying	*Spy’s trustworthiness diminished *Spy harmed by retaliation	*Target self-censors *Target’s enjoyment diminished *Target has emotional response *Target’s status is undermined	*Public trust diminished *Others self-censor

My analysis shows that the harms of spying can be considerable, but so too can the benefits. Spying can undermine cherished liberal and democratic values, even as

it can be essential for protecting these values. But calculating the harms and benefits of spying can be extraordinarily complex. Most of the consequences I discuss do not follow necessarily from spying, and to complicate matters further, nearly all of them can be harms under some conditions and benefits under others. To guide my reader through these complexities, I have appended two more detailed figures at the end of the chapter, mapping the pathways of my arguments.

5.1 Successfully Concealed Spying

In a set of recent articles, Mahon (2007, 189-190) (2008) defines “deceive” in the following way: “to intentionally cause another person to have or continue to have a false belief that is truly believed to be false by the person intentionally causing the false belief by bringing about evidence on the basis of which the other person has or continues to have that false belief.” On this rather cumbersome definition, spying does not seem to count as deception. Mahon’s definition requires that the deceiver cause the false belief by the “bringing about of evidence” to the person she intends to deceive. But the spy typically brings forward no evidence, so the spy, on Mahon’s conception, does not deceive.

But one might ask whether Mahon’s is the right account of deception. A more expansive account of “deceive” may be articulated as follows: “to intentionally cause another person to have a false belief, which is rightly believed to be false by the person intentionally causing the false belief.”² On this account, deception need not be the

²This is a minor alteration of the OED definition. Mahon (2008) proposes and rejects something very near this conception because an implication of this conception is that a variety of other (some would say) seemingly counterintuitive acts count as deception. For example, if I somehow carefully use electric shocks to alter your memories, so as to make you falsely accept some proposition, P, then I have deceived you. I’m prepared to accept that these acts involve deception. It seems to me a less undesirable implication than the implication of excluding the kinds of deception (such as

consequence of A bringing evidence to B in order to cause B to have a false belief. A could also deceive B by knowing B's belief and then changing what is the case so that B's belief becomes false.

Consider two cases involving bank robbers A and B and their loot. In the first case, A buries the loot in an overgrown lot and then tells B that he put the loot in his basement. By both Mahon's definition and the more expansive definition above, it seems correct to say that A has deceived B. Now consider a second case. In this case, A and B together bury the loot in the overgrown lot, but later the same evening, A unearths the plunder and moves it to his basement without informing B, in order to trick B about the plunder's whereabouts. By Mahon's definition no deception has occurred in this case. A has brought no evidence to B in order to cause him to have a false belief. But by the more expansive definition, A has deceived B, since A intentionally causes B to have a false belief.

Since it seems plausible that one can deceive either by intentionally bringing evidence to cause a false belief or by intentionally changing what is the case in order to cause a false belief, I think the more expansive account of deception is superior to Mahon's.

It follows from the more expansive account that spying is a kind of deception. Because the spy intends to conceal her information collection from her target, she intends for the target to go on believing that she has not had her information collected when this is not the case. The spy intentionally alters the world of the target in order to render her beliefs false and thereby to deceive her.

Consider the recent case of spying with webcams in Pennsylvania.³ In early 2010,

spying) I discuss above. Note that the conception I propose is even less expansive than Chisholm's and Feehan's (1977) classic account.

³The complete text of the lawsuit is available at: <http://media.philly.com/documents/robbins17.pdf>

the principal of Lower Merion High School accused 15-year-old Blake Robbins of taking illegal drugs. When the principal presented photographic evidence to demonstrate this “fact,” it became clear (from the location of the pictures, etc.) that some of the high school’s administrators were using the webcam affixed to Mr. Robbins’s school-issued laptop to periodically spy on him.⁴ Pictures of Mr. Robbins were taken in various places including his home and in his bedroom. Some of the pictures even featured Mr. Robbins and his high school friends scantily clad or undressed.

In this case, school administrators clearly intended Mr. Robbins to operate under the false assumption that his webcam was not capturing his private behaviors. Had Mr. Robbins known that his webcam could be turned on at any time by school officials and that school officials were in fact periodically exercising this prerogative, he no doubt would have turned off his computer, disabled the camera, or inhibited his behaviors when the computer was present. The spying was only effective because the administrators deceived Mr. Robbins about when and where he was under observation.

One consequence of successfully concealed spying, then, is that the spy deceives her target. According to some ethical theories, deception is wrong, full stop. Kant famously argued, for example, that it is never morally permissible to lie. (Korsgaard, 1986) But for most ethical theories, the issue is not so simple. The consequences of deception must be understood before judgments about right and wrong can be made. So let us go a step further and try to understand the effects that typically follow from the spy’s deception.

The most notable consequence is that the target tends to behave in a less guarded

⁴I place “fact” here in quotes because it turns out that the photographs that administrators believed to implicate Mr. Robbins for drug use only established that he was eating Mike and Ike candy.

way: she is less inhibited. She does and says things she would not do if she suspected she were being spied on. She also leaves information unsecured that she would hide away if she suspected she were being spied on. The result is that the spy tends to collect more sensitive and secretive information when her spying is successfully concealed.

But even this result is not unambiguously good or bad. “Information is power,” goes the adage, but of course power can be used for good or for ill. Recall the story of Gyges, told by Glaucon in the second book of Plato’s *Republic*. Gyges, the shepherd, discovers a ring in a tomb, which is only revealed after an earthquake. In the tomb is the corpse of a giant lying inside a sarcophagus in the shape of a great bronze horse. The giant wears a brilliant gold ring, which Gyges takes from the tomb. Later Gyges learns that his new ring has magical powers: a simple turn of the ring makes him invisible. The ring gives Gyges the power to spy at his whim, with no chance of being discovered. How does Gyges use his newfound power? He seduces his queen and murders his king, all in order to become the king himself.

Gyges’s case demonstrates the power of spying to produce bad consequences. But we can imagine an alternative story in which Gyges uses his golden ring for good. Gyges, on this alternative account, is like the comic book hero Superman, committed to justice and protecting the innocent. Rather than seeking his own material advantage, he uses his power to spy to thwart criminals and terrorists, to prevent wars, to save people from shame and humiliation, and to promote the well-being of all.

Most people would not be as purely good or evil as the two Gyges, but the cases illustrate the range of purposes to which a spy can put her informational advantages. Whether the collection of sensitive information is a cost or a benefit depends on the nature of the information – what effects revealing it will have, what activities it gives insight into, what lengths the target will go to ensure it remains concealed, etc. –

and what the spy does with the information she collects. Certain kinds of information could do grave harm to the target or others if it is revealed. Think of the great lengths countries with nuclear capabilities go to ensure that nuclear secrets are not revealed to terrorists or rogue governments. These precautions are reasonable because a nuclear bomb exploded in a large city could kill thousands if not millions of people, not to mention the devastating consequences to the world economy. Other kinds of information cannot be put to harmful purposes, and some spies would not use potentially harmful information even if they could.

The possible harms that could follow from the wrong person getting her hands on sensitive information are almost too numerous to list. The revelation of information can lead to embarrassment, physical harm, shame, termination of employment, loss of status, etc. But information need not be revealed for it to be harmful. The spy can also use the information she collects to manipulate or blackmail her target.

These possible harms of concealed spying are particularly grave in political contexts. Sensitive information can be used to fix a jury, destroy political opponents, squash dissent, convict an innocent person, pass harmful legislation, cover up fraud or malfeasance, etc. In fact, there are extensive historical precedents for such harms, even in liberal democracies. During the FBI's counterintelligence program "COIN-INTELPRO" in the 1950s-70s, federal agents used sensitive information gathered by spying on domestic political organizations and their members to, among other things, harass and discredit members of the U.S. Socialist Worker's Party, create internal conflict within groups opposed to the Vietnam war, and discredit the civil rights leader, Dr. Martin Luther King Jr.

Since it demonstrates the degree to which spying can be employed to do an agent harm, the King case is worth exploring in more detail. The FBI campaign to "neutralize" King was so extensive that William Sullivan, the man in charge of the cam-

paign, called it a “war” in which “[n]o holds were barred.” The FBI extensively spied on King, opening his mail, bugging his home and his hotel rooms, wiretapping his telephone, and tailing him and his confidants. It used the sensitive information it gathered from spying to undermine his reputation with American government officials, foreign leaders, various churches and universities, and the press. (Fain, Plant and Milloy, 1977)

At one point the FBI attempted to blackmail King. The Bureau sent him a recording of conversations from his hotel room, which “was accompanied by a note that Dr. King and his advisors interpreted as threatening to release the tape recording unless Dr. King committed suicide.” It is unclear whether the FBI actually intended to demand King’s suicide, but there is no doubt that the Bureau’s intentions were nefarious. One agent testified that the purpose of the hotel recordings were to sabotage King’s marriage. (Fain, Plant and Milloy, 1977, 104-105)

The evidence is overwhelming that the FBI sought to undermine King because of unfounded fears, including that King would abandon his message of non-violence and that he was taking orders from communists. In a memo to FBI field offices in March 1968, the FBI director explained the purpose of COINTELPRO and later revealed the fears harbored by the bureau about King. “The purpose of this new counterintelligence endeavor is to expose, disrupt, misdirect, discredit, or otherwise neutralize the activities of black nationalist, hate-type organizations and groupings, their leadership, spokesmen, membership, and supporters, and to counter their propensity for violence and civil disorder.” He later added that one of the long-term goals of the program was to, “[p]revent the rise of a “messiah” who could unify, and electrify, the militant black nationalist movement. Malcolm X might have been such a “messiah;” he is the martyr of the movement today. Martin Luther King, Stokely Carmichael and Elijah Muhammed all aspire to this position...King could be a very real contender for this

position should he abandon his supposed “obedience” to “white, liberal doctrines” (nonviolence) and embrace black nationalism.”⁵

The FBI’s fear that King would abandon nonviolence accompanied a host of other ungrounded apprehensions. In December 1963, FBI officials held an all day conference dedicated to King’s allegiance to the Communist Party. At the meeting’s outset, participants agreed that King “was knowingly, willingly, and regularly cooperating with and taking guidance from communists.” (Kotz, 2005, 83) The truth, however, was that King had almost no links to communism. Some of his advisers, e.g. Stanley Levison, had distant and tenuous links to communism, but King himself had no direct links to communism, and there was no evidence that any of his advisers remained connected to communist organizations. Only in 1976, eight years after King’s death, in testimony to the Church Committee, did the FBI admit that neither King nor his Southern Christian Leadership Conference were cooperating with communist organizations.

The King case provides a clear example of how sensitive information collected from spying can be used for ill, and it is important to keep in mind that the resources available to the FBI in the 1960s are but a sliver of those available to the Bureau today. The FBI today can compile a considerably more thorough dossier on its targets than it could in the 1960s. Then, conversations took place on the telephone or face to face. Little if any information was stored in electronic databases. If the FBI wanted to illicitly obtain King’s medical records, for example, it would have to break into a

⁵Memo from FBI Headquarters to all SACs, 3/4/68 in Fain, Plant and Milloy (1977). Snippets of this memo appear (paraphrased) in the popular protest rock band Rage Against the Machine’s song “Wake Up.” In the song’s background a track plays, “He may be a real contender for this position should he abandon his supposed obedience to the white liberal doctrine of non-violence and embrace black nationalism. Through counterintelligence it should be possible to pinpoint potential troublemakers and neutralize them.” I point out the appearance of this memo in popular culture to foreshadow a couple of points I make below: that spying can evoke widespread fear or paranoia and that it can damage the trustworthiness of the spy.

hospital. Now, nearly every aspect of our lives is collected and stored electronically. Every bit of information, no matter how sensitive or secretive, is thus vulnerable to the clever hacker spy.

But just as sensitive information can be used for ill, it can also be used for good. Consider the following example. In June 1986, two hospitals in the UK initiated a program using covert video surveillance (CVS) to monitor children suspected of having abuse-induced illnesses. (Southall et al., 1997). The monitoring occurred only in isolated hospital cubicles. Parents were told that their children's various physiologic conditions were being monitored, but they were not informed of the hidden cameras or microphones in the cubicle. They were urged to care for their children as if they were at home. Of the 39 children monitored all were suspected of being victims of abuse; 36 were referred because of recent "apparent life-threatening events," including suspected poisoning and strangulation. In 33 cases, CVS revealed abuse, including in 30 cases intentional suffocation. In one case a child's parent deliberately fractured her arm.

Southall et al. (1997) conclude from their study of the program that CVS is an invaluable diagnostic tool because the traditional methods of diagnosing child abuse are notoriously erroneous. Abusive parents are often skilled liars and can concoct plausible stories explaining the harm done to their children, and in many cases children are too young or too scared to provide case workers with verbal confirmation of abuse. When abuse is not diagnosed, children almost always return to their homes and are thereby placed at risk of serious harm, and in many cases death. Furthermore, because children often have siblings, and abusive parents tend not to target their mistreatment at only one child, the correct diagnosis of abuse can prevent harms not just to the child under investigation.

CVS has the additional benefit of making convictions of abusive parents all but

certain. Typically it is difficult to demonstrate abuse in court, particularly when there are nearly always experts willing to testify that the evidence presented by doctors and prosecutors is insufficient to infer that the causes of harm were parental mistreatment (Ibid: 750). But in the cases in which CVS identified abuse, nearly all of the abusive parents (more than 30 of 39 cases) were charged and convicted.

I am not defending the use of CVS to diagnose child abuse. Such a defense would require knowledge of all of the plausible alternative methods for diagnosing abuse, not to mention a more complete understanding of the likely costs of CVS, which is no doubt likely to discourage some parents from taking their children to the hospital to get urgent medical care.⁶ My aim is merely to point out how sensitive information can be used for good. CVS produces information that doctors, nurses, case workers, and police officers can use to prevent – or at least to make less likely – a set of deeply harmful consequences. We can all (I presume) accept that it is nearly always better if innocent children are not physically or psychologically abused or even killed, that parents who mistreat children in these ways should be held responsible and provided with professional counseling, and that those who break just laws prohibiting child abuse should be punished. CVS, when used properly by health care professionals and law enforcement officials, plausibly makes these desirable outcomes more likely.

More generally, spying, of which covert video surveillance is a subset, sometimes produces information that makes it more likely that some agent can act to prevent bad consequences, or (less frequently) makes it more likely that some agent can act to promote good consequences. Spying does not itself, in these cases, prevent bad consequences, but rather reveals information, making it possible for someone rightly motivated and appropriately placed to prevent bad consequences – it produces ac-

⁶There is a live and sophisticated debate on CVS. See, for example, Foreman and Farsides (1993); Samuels (1994); Southall and Samuels (1993, 1995); Thomas (1996).

tionable intelligence.

The benefits of spying aimed at producing actionable intelligence are thus conditional in two respects. If there exists no person both capable and willing to prevent bad consequences, or if the spying does not produce relevant information, then no bad consequences are averted and the spying does not prevent, but likely produces costs. If doctors, for example, could not pass evidence of abuse to the judicial system or to law enforcement or social services professionals, then CVS would not be an attractive tool. Similarly, if parents employed only methods of abuse that could not be caught on camera, CVS would not be useful.

An obvious example of government spying that produces actionable intelligence is the spying used by law enforcement officials to prevent crimes.⁷ Without a fairly complete picture of a criminal's intentions, crime prevention is often difficult, if not impossible. Spying often illuminates the intricate intentions of criminals: it reveals when, where, how, and even why a particular crime will take place. (Solove, 2003) On television shows and in movies, law enforcement officials routinely use spying to foil terrorist plots or prevent murders, but in the real world, spying is also employed to prevent more mundane offenses. In 2009, for example, the FBI arrested Raj Rajarantnam, the billionaire founder of the hedge fund The Galleon Group for insider trading. The U.S. attorney's office built its case against Rajarantnam (and other conspirators) primarily on evidence garnered from over 18,000 wiretapped telephone

⁷Because crime prevention is not the same thing as the prevention of wrongful acts, some caution is required before accepting crime prevention as a benefit. Not all laws are just laws. Some spying may simply protect and further entrench a morally dubious status quo. Spying in the American South last century, for example, no doubt perpetuated Jim Crow laws; and this for years would have been considered by many to be legitimate law enforcement. But virtually no one (still) thinks that spying, in these cases, prevented bad acts, nor should they. The Mississippi State Sovereignty Commission (MSSC), for example, was set up to protect the state against the enforcement of civil rights laws by the federal government. The MSSC placed secret agents inside a range of organizations thought to be advancing or sympathetic to civil rights not only to collect information but also to disrupt or subvert the organizations' goals. See Bowers (2010) and Irons (2010).

conversations.⁸ Rajaratnam was found guilty of illegally profiting from securities trades based on inside information received from executives within companies such as IBM, Intel, and the consulting firm McKinsey & Company.

Spying is not only used to prevent crimes, it is also used to enforce laws that have already been broken. Law enforcement officials, in these cases, use information gathered from spying to identify, track, and capture criminals, and in some cases, to ensure that criminals are prosecuted and convicted. Because criminals often go into hiding and because relatives and close associates often cannot be relied on for accurate information about their whereabouts, spying is sometimes expedient if not necessary for ensuring arrests and convictions.⁹

Spying also plays an important role in the international analogues to crime prevention and enforcing already broken laws: preventing surprise attacks, and tracking and bringing to justice those who have committed offenses against the state or its citizens. Furthermore, governments routinely engage in counterespionage, i.e. they spy on those suspected of being spies for other states (or non-state actors). In these cases, the outcome that governments typically seek to prevent is the unwanted disclosure of secrets – usually about national defense – that can be used to coerce, manipulate, or

⁸SEC v. Rajaratnam, 622 F.3d 159 (2nd Circuit 2010).

⁹One should also be cautious before accepting law enforcement as a benefit. Holding other things equal, one might accept that it is generally desirable in a moderately just legal system that law enforcement officials successfully arrest, prosecute, and convict criminals. If criminals are never arrested and never convicted then the law has very little deterrent effect. The risk also remains that criminals remaining on the streets will commit future crimes. But not all criminals must be caught and convicted for the law to have its deterrent effect; and many criminals pose minimal threats to society, even when they're never apprehended. Nor is it obvious that lacking the ability to spy will affect successful law enforcement in most cases. Before the twentieth century, the United States had no (non-military) federal investigative or intelligence agencies, nor was undercover police work particularly common, but, for the most part, criminals were still apprehended and the law still had a deterrent effect. So successful law enforcement sometimes requires spying, but there is no reason to think that all successful law enforcement requires spying. See Powers (1987) and Kessler (2003); Marx (1990, Ch. 1) provides statistics on the number of FBI undercover operations and the budget set aside for these operations.

otherwise harm the state or its members.

So government spying can produce actionable intelligence that prevents crimes, harms, and threats from abroad, assists in law enforcement, and/or aids in the identification or capture of foreign spies.

Thus far the argument has been that spying can lead to the deception of the target about her audience and the security of her information. The target's false beliefs may then lead her to behave in a less guarded way, thereby disclosing sensitive information to the spy. Whether the spy's collection of sensitive information turns out to be a cost or a benefit depends on the nature of the information as well as what the spy does with the information. With the right kind of information, the ill-intentioned spy can profoundly harm her target or others; similarly, the virtuous spy can avert catastrophic harms.

Let us now turn away from consequences primarily affecting the target and toward consequences mostly affecting the spy. Plutarch wrote:

The consciousness that I have done terrible deeds, like a sore in the flesh, leaves in the mind a regret which is forever wounding and piercing it. . . . Neither a costly house, nor a heap of gold, nor pride of race, nor high office, nor charm nor eloquence of speech, make life so peaceful and serene as a soul pure of evil acts and desires, having as its spring a life a nature steadfast and undefiled. From it flow noble deeds, bringing with them an inspired and joyful energy, together with loftiness of thought and a memory sweeter and more lasting than the hope which Pindar says is the support of old age. (Plutarch, 1951, 353)

Spying can be what Plutarch calls a “terrible deed,” that is, it can obviously and decisively violate a society’s established moral norms. In these cases, one effect of spying is that it can have a strong emotional effect on the spy, perhaps even “wounding and piercing” his soul as Plutarch suggests. Although I think the spy’s feelings of guilt, shame, and regret can be a consequence of spying, I want to explore a different more important implication of the passage. Plutarch claims that noble deeds flow from a soul “pure of evil acts and desires.” The inverse of this claim, which the passage also seems to support, is that wicked deeds flow from a soul that is filled with evil acts and desires. As a person diverges from moral norms, her soul loses its purity and is clouded by her bad actions, she becomes more likely to diverge from moral norms in the future.¹⁰

In her analysis of the ethics of lying, Bok (1999) relies on a similar logic to defend her claims that telling the solitary lie is rather rare, and that lying corrupts the liar. Lies tend to distort the liar’s moral and psychological barriers: later “lies seem more necessary, less reprehensible; the ability to make moral distinctions can coarsen; the liar’s perception of his chances of being caught may warp.” (25) In short, lying tends to inure the liar. Future lies seem less dangerous, less morally problematic.

These are ultimately empirical claims which Bok does not defend with evidence, psychological or otherwise. But the psychological evidence supporting Bok’s claim is compelling. Dan Ariely (2012), for example, ran an experiment in which participants performed an exercise in which they were given a minor incentive to cheat by misrepresenting data displayed on a screen. The exercise was iterated many times,

¹⁰In his (1997) novel *The Untouchable*, based loosely on the life of the Cambridge spy Anthony Blunt, John Banville has his protagonist Victor Maskell voice both of Plutarch’s worries in the context of spying, “...it was not the philosophy by which I lived, but the double life itself...that acted on me as a debilitating force. I know this has always been said of us [spies], that the lying and the secrecy inevitably corrupted us, sapped our moral strength and blinded us to the true nature of things, but I never believed it could be true.” (45)

and what Ariely wanted to know was whether people would cheat more frequently as the game progressed. He hypothesized that participants would exhibit “balanced cheating” at first, in order to maintain their beliefs about their own honesty. But at some point participants would reach an “honesty threshold,” after which they would start to think “What the hell, as long as I am a cheater, I might as well get the most of it.” (129)

Ariely’s results support his hypothesis. Not only did participants cheat more frequently as the exercise went on, but the frequency with which they cheated tended to jump abruptly. He concluded, “...when it comes to cheating, we behave pretty much the same way as we do on diets. Once we start violating our own standards...we are much more likely to abandon further attempts to control our behavior and from that point on there is a good chance that we will succumb to the temptation to further misbehave.” (130)

Beyond the psychological impact of lying, Bok argues that lies are often accompanied by more lies to fill in, cover up, or make sense of earlier lies. Most people have been ensnared by their own lies and often the only way to release themselves from these snares is to tell another lie thereby laying yet another trap. As Bok points out, sometimes the thicket of lies becomes so dense for the liar that she requires great intelligence and memory to “remain true” to her lies. One might ask whether spying presents a similar problem. Does spying corrupt the spy?

Following Bok’s logic, the answer seems to be both yes and no. On one hand, lying and spying share an important feature for the liar and the spy: both can be expedient for achieving prudential aims. It is this expedience combined with the success of previous lying and spying, one might think, that tends to distort the consideration of the next lie, the next choice to spy. This line of argument fits nicely with Bok’s broader discussion of lying from the liar’s perspective. The liar is likely, she contends,

to underestimate the costs of her lies both to herself and to what she calls “human communities.” One might think that in both cases the liar and the spy underestimate (consciously or not) so as to promote their own prudential interests.

On the other hand, spying and lying seem different because it is not always obvious how new spying might extricate a person from her previous spying. Once A is suspected by B of spying, it is difficult to see in most cases how A could allay B’s suspicion by spying more. New lies can often cover up old ones, but it doesn’t seem like new spying can cover up old spying.

Bok’s (and Plutarch’s) logic provides some reason, then, to believe that spying can corrupt the spy, but it may not be as compelling as it is in the case of lying. Three more reasons, however, I want to suggest, support the claim that spying begets spying: it tends to produce expertise, which is likely to be leveraged in the future; institutions are often created to spy and these institutions often outlast their original purposes; and finally, spying is often an option in strategic interactions in which, given each player’s calculation of her own prudential interests, the choice to spy now makes future spying more likely.

Increasing the likelihood that an agent will spy in the future does not, however, necessarily corrupt the agent. Future spying may be justified spying. But Ariely’s study does suggest that corruption is a risk: each additional time an agent “cheats” by violating moral norms, she risks reaching a threshold at which she stops concerning herself with moral norms altogether.

First, consider that spying tends to produce a kind of expertise that acts like lying do not, or do so to a lesser degree. While one can make sense of a good or bad liar and there are professions (e.g. lawyers and politicians) that are often – perhaps unfairly – accused of producing particularly adroit liars, the practice of lying is not typically

considered one in which people tend to develop skills, technological or otherwise.¹¹ Spying, on the other hand, is typically carried out by experts – experts in wiretapping, satellite surveillance, tailing a suspect, hacking into a database, etc. Those who develop these sorts of expertise for one purpose are more likely to use them in the future for another, particularly given that these skills are now so valuable on the open market.¹²

Closely related to the expertise developed by spying are the institutions designed – usually by governments, but also by corporations and other private organizations – to carry out spying. Like the expertise created to spy in one context but employed in another, one might expect that institutions built to tackle one problem will sometimes be leveraged for other purposes perhaps long after the original problem has waned as a concern. As Knightly (1986, 6) in his sweeping and insightful analysis of twentieth century spying remarks, “once established, intelligence agencies have proved very difficult to get rid of.” More generally, institutions, as has been pointed out in the literature on political economy, tend to have a kind of “stickiness.” North (1990, Ch. 11) famously argues that there is a “path dependence” to institutions: peculiar historical events can lead to one set of institutions winning out over another. The winning institutions can then long overstay their welcomes, persisting long after they are unnecessary or when there are “better” institutions to put in their place.

Henry Shue (1978) raises this worry with respect to the practice of torture. He asks

¹¹Javers (2010, Ch. 7) tells the story of a firm called BIA, created by former CIA operatives, which consulted for various corporate clients, teaching them how to detect lies. The service proved quite useful for the purposes of due diligence before a merger or evaluating the credibility of an executives statement to shareholders. But, when BIA employees attempted to expand their services by teaching their clients how to effectively lie, they found that it was almost impossible to beat their own techniques of lie detection. Expertise in lying, in other words, was very difficult to produce.

¹²A number of books have come out recently exploring the privatization of intelligence. See e.g. Shorrock (2008).

whether what he calls “terroristic torture” could be used briefly and then banned. He concludes doubtfully, pointing out that it is rare that those with the power to torture relinquish their power. More commonly, torture becomes “a routine procedure institutionalized into the method of governing. Some bureaus collect taxes, other bureaus conduct torture. First a suspect is arrested, next he or she is tortured. Torture gains the momentum of an ingrained element of a standard operating procedure.” (138)

If routinization of this sort is a serious concern for torture, then the worry is even graver for spying.¹³ Spying, like torture, is a tremendously powerful tool, which those in possession would be reluctant to renounce, and the institutionalization of spying is more likely and less controllable than torture. While some governments no doubt have bureaus of torture, all (literally *all*) governments have bureaus devoted to domestic and/or international spying.¹⁴ Further, given the profoundly secretive nature of spying, it may be even less controllable than torture. With torture there are at least identifiable victims, people with whom to sympathize and around whom to organize political resistance and change. Not so for spying. The victims of spying are often nameless and faceless. They have no idea that they have been the target of investigation until long after the fact if at all.¹⁵

So the tendency to institutionalize spying coupled with the tendency of institutions to outlive their original purposes is another reason to think that spying begets spying. A third and final reason applies to spying by rivals, competitors, and enemies,

¹³The routinization worry has not gone unarticulated in political debate. See Moynihan (1998, Chs. 7-8).

¹⁴See Knightly (1986, Ch. 1).

¹⁵Or when the victims suspect they are being spied on, it is sometimes difficult for them to provide evidence to demonstrate this fact. A recent example of this is *ACLU v. NSA*, in which the district court judge ruled that the NSA’s warrantless wiretapping was illegal and unconstitutional, but the decision was later overruled because the plaintiffs could not demonstrate legal standing because they couldn’t prove that they were wiretapped.

whether they be individuals, firms, or states. The idea is that the structure of the strategic interaction among competitors sometimes dictates that spying at t_1 makes spying at t_2 much more likely. Consider that many competitive situations involving spying can be specified by the classic prisoner's dilemma game. Imagine, for example, two firms in the same industry each choosing whether to innovate or to spy to steal their competitor's innovations. If we assume that it is less costly to spy and steal secrets than it is to innovate, then each firm plausibly has the following preference ordering: spy when the other firm innovates, innovate when the other firm innovates, spy when the other firm spies, and innovate when the other firm spies. The dominant strategy of this game for both players when they play only one round is thus spy.

But the assumption that the game is played only once is probably not correct for the situation as I have described it, given that the firms are competitors in an ongoing struggle for market share and profits. What happens when the game is iterated? The simple answer is that it depends. When the players know there is some finite number of rounds, spy remains the dominant strategy. However, when the players expect the game to continue in perpetuity, as is likely in the situation I have described, experiments suggest that cooperation (innovation in this case) is possible. Axelrod's (2006) famous study, for example, suggests that under a particular set of fairly common conditions the optimal strategy in the iterated prisoner's dilemma is tit-for-tat, or payment in kind. The player executing the tit-for-tat strategy begins with cooperation (innovation) and only defects (spies) in response to defection (spying) by her competitor. When both players use the tit-for-tat strategy, one choice of spy by either player (perhaps because the player misperceives innovate for spy) turns the game into a spiral of spying (often called the "death spiral"). Spying at t_1 produces spying at t_2 , which produces spying at t_3 , and so on. Spying begets spying.

If it is correct to think of some competitive situations involving spying as an it-

erated game of this sort, then there is another reason to think that spying will tend to produce more spying. I am, of course, not committed to saying that most competitive situations involving spying are games of this sort. When spying is difficult to detect, it might not be optimal to play tit-for-tat. If one's competitor releases a similar product or technology, for example, how is one to know with any certainty whether the product or technology was developed internally or stolen? Often there is no way of knowing (at least not without spying!).

Spying can make the spy more likely to spy in the future, then, and one risk of repeated spying is that the spy will become corrupted. But spying today can also produce institutions and skills for more effective or efficient justified spying in the future. Obviously, institutions would need to be designed very carefully in order for this to be the case, and some might doubt whether such fine grained institutional design is likely to be successful. But the possibility seems undeniable.

5.2 Suspected Spying

When spying is suspected, it often compels the target to self-censor, especially when she suspects that knowledge of her behavior will lead to sanctions or rewards.¹⁶ Hence the spy can sometimes control her target without directly interfering in her affairs.

Let me explain this more carefully. When the target suspects she is being spied on, she changes her beliefs about the relevant audience for her behaviors, augmenting the relevant audience to include at least the spy. If this new, larger relevant audience includes in it people who endorse norms that proscribe (praise) some of her behaviors

¹⁶Connolly (1993) calls this self-censorship "anticipatory surrender." More broadly self-censorship is an example of what is often called the second face of power. See, e.g., Bachrach and Baratz (1962).

and who have the capacity and will to sanction (reward) her for violating (adhering to) their norms, then suspected spying may lead her to engage in self-censorship. Spying “chills” the target’s conduct.

Notice that in some cases the target need not suspect sanctions or rewards to censor her behavior. The suspicion of spying alone is sufficient. In most cases, however, the target’s suspicion of rewards and sanctions is necessary to condition her behavior.

Consider the following example. Suppose you sit on the local school board and you are a closeted communist – closeted because in your town one cannot both openly be a communist and serve on the school board. Suppose further that I am a rabid anti-communist and that I have a track record for publicly outing communists. Now imagine that you begin to suspect that I am spying on you. Not wanting to jeopardize your seat on the school board, you refrain from reading Marx, you skip the local party meeting, and you censor yourself when you talk politics with your friends and family.

The example suggests that your beliefs about the relevant audience for your behaviors play a decisive role in determining whether you engage in self-censorship, and that my spying affects your beliefs. When you had no reason to suspect that I was spying on you, you read your Marx, attended your local party meeting, etc. But when you began to suspect that I belonged in your relevant audience, you immediately began to inhibit your behaviors.

The insight that one’s beliefs about one’s audience shape one’s behaviors has been developed previously in a number of other contexts by political theorists. Indeed, I take it to be the key insight operating in Foucault’s analysis of disciplinary power.¹⁷ Because it helps further illuminate the connection between one’s beliefs about one’s audience and anticipatory responses to power, I shall briefly consider Foucault’s (1995)

¹⁷It is also the key insight operating in Bentham’s proposal for the Panopticon. See Bentham (1838, Vol. IV). I take up Bentham’s view briefly below.

famous account.

Disciplinary power, for Foucault, is subtle and profoundly invasive, reaching into every nook and cranny of one's life. In contrast to what Foucault calls pre-modern power, it renders most coercion and force by one actor over another – and particularly the most spectacular and brutal kinds – unnecessary. As Foucault says, "...the perfection of power should tend to render its actual exercise unnecessary." (201) When describing the operation of disciplinary power, Foucault draws on Bentham's Panopticon as a metaphor. The Panopticon is an architectural form, designed principally to be a prison but adaptable to many other purposes, featuring a tall central tower circumscribed by a building divided into cells. From the tower each cell and its inhabitants are entirely visible; there is no place for prisoners to escape visibility.¹⁸ But from the cell, while the tower is visible, the occupants of the tower are not: venetian blinds are affixed to the windows of the tower, and the hallways leading to the center of the tower zig and zag, so as not to betray the presence or absence of a guard. The design is meant to give the impression to inmates that they are under intense and unremitting scrutiny; it is meant to produce the same effect that observing every prisoner would have, at a fraction of the cost.

The Panopticon, Foucault says, "assures the automatic functioning of power." (201) I think this claim is best understood by drawing on the concepts of audience, norms, rewards, and sanctions that I mentioned above. When a prisoner is isolated (i.e. she has no audience or a very limited one), as is the case when she is in a dungeon, she might attempt an escape, she might try to communicate with other prisoners, or she might pitch a fit because of her poor conditions. But when the prisoner is under the gaze of her captors, she will tend to conform to their norms. She does so because

¹⁸Bentham's plan also incorporated a series of tubes, connecting each of the cells with the tower, permitting the guards to hear the prisoners.

she knows or expects that attempting escape, communicating with other prisoners, or pitching fits comes with considerable sanctions from the guards. Because she wishes to avoid these sanctions, she conforms. At first, the prisoner's conformity is just the result of a rational calculation. The costs to her of the sanctions overwhelm the benefits of non-conformity.

Yet Foucault's analysis goes further than just this rational calculation. He claims that the constant visibility characteristic of the panopticon (and of modern society more generally) ultimately plays a role in *producing* subjects, i.e. in determining the beliefs and values of the prisoners. Over time, conformity can lead to norm internalization. The prisoners' repeated conformity can habituate them to internalize the beliefs and norms of their captors.¹⁹ Exactly how this process of internalization occurs and how long it takes to operate are not entirely clear; what is clear is that internalization does happen across many contexts.²⁰

Norm internalization is likely most common when there is no protected space where visibility does not extend. This link is perhaps why spying has been the hallmark of regimes aspiring to near total control of their citizens. To eradicate dissent and guarantee strict obedience to an ideology, it is not sufficient to convince subjects that their illicit behavior will be gravely punished. Regimes must also instill a belief in their subjects that every forbidden behavior (or thought) will be discovered and punished, that there is no protected or private space, separated from the state's eye.

¹⁹The prisoner example may not be the best one to illustrate Foucault's point, since there may be no amount of habituation that would lead many hardened criminals to incorporate the norms of their jailors. A better example comes from Susan Bordo's (2004) work on how women in the West internalize norms of the "ideal" female figure.

²⁰As Hardin (2002, 29) remarks, "Norms can evidently be internalized, so that we simply act from them without need of sanction. I have little to say here of this possibility, on which even the best arguments for how it works are not very compelling, although the claim that it works seems clearly to be correct."

Every friend or relative could be an informant; every room could be bugged; strangers could be watchers. Defiance always comes with punishment, and to speak or even to gesture in the manner of remonstrance is to risk betrayal and dire penalty. As Govier (1996, 151) points out, in these circumstances of totalitarianism, “the space between people is destroyed; there is an artificial loneliness. People cannot confidently share their feelings and beliefs.” Trust is annihilated.

Foucault mostly focuses on the implications of disciplinary power for people’s freedom. Self-censorship and norm internalization, he emphasizes, are always costly to freedom. But whether self-censorship is on balance a good or a bad consequence depends on the counterfactual condition, that is, it depends on whether an agent’s conduct would be better or worse in the absence of the particular instance of disciplinary power. Something similar can be said of norm internalization, when it follows from habitual self-censorship: whether it is on balance a cost or a benefit for agents to internalize a norm depends on how they would behave if the norm were not internalized. When agents engage in behaviors that are on balance uncontroversially beneficial to themselves and others, inhibiting these behaviors is a cost. Similarly, if norms that generally lead to beneficial conduct are replaced by norms that lead to harmful behavior, then norm internalization is a cost.

The potential magnitudes of these costs cannot be exaggerated. Self-censorship can lead to the destruction of liberal and democratic values, for example. Consider, first, the democratic values of free speech, free assembly, and political compromise. None of these values are secure when citizens and government officials lack assurances that others will not spy on them. If people suspect they are being spied on by their government when they speak out publicly against it, for example, they may censor their speech. Spying can thus have a chilling effect on public speech. But it is not only public speech that could be affected by spying. Charles Fried (1970, 143) argues that

some degree of privacy is necessary to define oneself. People need protected spaces outside of the gaze of others to try out political or ethical positions. More recently, Neil Richards (2008, 389) has argued that intellectual privacy, “the ability, whether protected by law or social circumstances, to develop ideas and beliefs away from the unwanted gaze or interference of others,” is crucial for the freedom of speech, because “[s]urveillance or interference can warp the integrity of our freedom of thought and can skew the way we think, with clear repercussions for our subsequent speech or writing.”

A similar line of argument holds for the freedom of assembly. If people suspect that their behaviors are being observed or that their movements are being tracked by their government, then they will be reluctant to attend gatherings not in accord with the values of their rulers. They will remain in places they perceive to be safe, places where they are either free from the government’s gaze or where they will not appear threatening. They will avoid what Hannah Arendt calls “spaces of appearance,” where people’s ability to see one another illuminates possibilities for collective action.²¹

The suspicion of spying can suppress the energetic contestation so vital to a flourishing democracy, then. It can also substantially impede good governance. Legislators, bureaucrats, and judges often require closed doors to reach compromises in order to make the best feasible policies or judgments.²² If public officials suspect that their private deliberations may reach broader and especially hostile audiences, then they may be unwilling to compromise. They may fear that compromise will be perceived as disloyalty, or they may use deliberations to grandstand or to score political

²¹See Arendt (2006) and Marquez (2011).

²²See Thompson’s discussion of “temporary secrets” (1999, 184-185). See also (Bok, 1989, Ch. 2) and Luban (1996).

points. Transparency might be, in most cases, like sunshine, the best antiseptic. But in some cases, secrecy and therefore assurances against spying are necessary for good governance.

Turn now to the liberal values implicated by self-censorship. Liberal societies are supposed to tolerate a wide range of views about the good life, what Rawls (2005) calls “conceptions of the good.” Further, they are supposed to tolerate people living out their diverse conceptions of the good, so long as their conduct is not unreasonable or excessively harmful to others. But even liberal societies contain intolerant people, indeed the majority is often intolerant of a range of harmless behavior. The intolerant seek to impose their own beliefs on others; they seek to reward adherents and punish apostates; they seek, in short, the destruction of liberalism.

Keenly aware of this danger, Mill (1989) famously argued that liberal societies require more than protections for citizens against their governments, they also require protections for individuals against the force of prevailing opinion. In his own words: “Protection...against the tyranny of the magistrate is not enough: there needs protection also against the tyranny of the prevailing opinion and feeling; against the tendency of society to impose, by other means than civil penalties, its own ideas and practices as rules of conduct on those who dissent from them.” (8)

One of the principal ways that liberal societies protect individuals against the force of prevailing opinion is by establishing norms (practices, laws, etc.) protecting privacy. Privacy norms specify, among other things, when, where and with what means it is appropriate for people to observe others. One very common privacy norm, for example, is that it is inappropriate to observe people in their places of residence, without their consent. Privacy norms create private places and private channels of communication, both of which permit people to escape the observation of select others. They permit people to freely engage in a range of harmless counter-normative

behavior. They enable people to behave uninhibitedly, even eccentrically, to act on their passions and proclivities, without censure or shame.

When people suspect they are being spied on, they can no longer trust that privacy norms are being respected. Without private places, people tend to bow to majority opinion.²³ People act, dress, speak and even think more like their neighbors. Those with religious, sexual, and political views not in step with prevailing opinion become vulnerable to shame, embarrassment, and even physical harm. They live in fear, or self-denial. Since it is dangerous to share the deeply felt elements of their identity, they struggle to forge bonds of friendship and intimacy.²⁴ The diversity characteristic of liberal societies begins to wane.

Hence the costs of self-censorship can be considerable. But when agents engage in behaviors that are uncontroversially harmful to themselves and others, inhibiting these behaviors can be on balance a benefit, and just like the harms of self-censorship, these benefits can be significant. Bentham exuberantly campaigned for the Panopticon because he thought, if employed appropriately, it could produce virtually innumerable benefits. He begins his collection of letters on the Panopticon (1838) saying, “Morals reformed – health preserved – industry invigorated – instruction diffused – public burthens lightened – Economy seated, as it were, upon a rock – the Gordian knot of the Poor Laws are not cut, but untied – all by a simple idea in Architecture!”

Bentham’s exuberance induces a kind a skepticism in many readers, but there is reason to think even he underestimated the full range of applications for the Panopticon. For Bentham, the Panopticon was primarily an architectural form. Only with

²³Although less likely, the result could be the reverse. If a person came to suspect that she was being spied on by a powerful minority group, she might begin to flatter members of the minority.

²⁴See Fried (1970, 142). He argues, “To be friends or lovers, persons must be intimate to some degree with each other. Intimacy is the sharing of information about one’s actions, beliefs or emotions which one does not share with all, and which one has the right not to share with anyone.”

later interpreters of Bentham’s work – and particularly with Foucault – does the Panopticon become just a metaphor. Wherever visibility (or “observability” more precisely) can reach, behavior can plausibly be altered. “Visibility is a trap,” Foucault warns. And with tiny cameras, microphones, programs with the potential to capture one’s every keystroke, etc. there are now few thoughts and even fewer behaviors that cannot conceivably be observed.

The chief potential benefit of self-censorship is its contribution to people’s basic security. Some of those willing to contemplate atrocities can be deterred by the threat of being caught and punished, and this threat can be magnified by spying. As people begin to suspect that they are being spied on, their beliefs about the likelihood that they will be caught and punished for performing harmful or illegal acts increases.

Yet the benefits of self-censorship, as Bentham suggests, stretch well beyond security. Spying can be used in neighborhoods to prevent property crimes; in stores to avert theft; in schools to ensure students are on their best behavior; in the workplace to promote productivity; or among states attempting to ensure compliance with international agreements. When agents suspect they are being spied on, they are less likely to engage in all kinds of harmful or otherwise wrong actions (especially when these actions risk punishment). They are also more likely to engage in virtuous or beneficial actions.

Psychological studies support these claims. A number of studies show that the presence of others or signals that others may be observing produce more cooperative behavior. (Andreoni and Petrie, 2004; Burnham, 2003; Dawes, McTavish and Shaklee, 1977; Hoffman et al., 1994; Kurzban, 2001) Haley and Fessler (2005) show that merely including the images of eyes on a computer screen increased cooperation in the Dictator Game. Burnham and Hare (2007) found that the presence of a robot with eyes produced a similar result for the Public Goods Game. In a set of field ex-

periments, Bateson, Nettle and Roberts (2006) and Ernest-Jones, Nettle and Bateson (2011) showed, first, that images of eyes substantially increased contributions to an “honesty box” for contributions to a building-wide coffee fund, and then that images of eyes significantly decreased littering behavior in a university cafeteria.

Sociological studies suggest similar effects. A recent meta-study of the effects of CCTV cameras on crime rates, for example, concluded that the presence of CCTV cameras reduces crime rates by approximately 4%. (Welsh and Farrington, 2003) Similarly, Newburn and Hayman (2002) found that the presence of a surveillance camera decreased violent incidents by both prisoners and their captors.

Although CCTV is usually an overt method of observation, and overt observation may, in some circumstances, have a stronger effect on people’s behavior than when they suspect spying, this need not be the case. Overt observation can be resisted in ways that spying often cannot (it is easier to avoid being filmed, for example, when you know the location of the camera), and the suspicion of spying, when it is strong, can profoundly affect people’s conduct.

So the suspicion of spying can lead people to condition their behavior. But sometimes the suspicion of spying will not lead people to alter their behavior at all; instead it will lead them to enjoy the activities they are engaged in less. In rare instances, as is the case with the exhibitionist, the suspicion of spying can *enhance* a person’s enjoyment, but I will leave these cases aside.

We can again draw on the idea of an audience to develop this point. Some activities to be fully enjoyed require audience restrictions. When spying is suspected, the target learns that her audience is not restricted in the right ways, thereby diminishing her enjoyment. Consider two close friends having a conversation when they begin to

suspect that they are being spied on. Although they continue their conversation, it turns to less intimate matters as a result of the suspicion. It becomes superficial, trite, less playful and spontaneous. The two friends still enjoy their conversation in its new shallower form, but the conversation has nonetheless been tainted. They enjoy it less.

Notice that what matters for the friends to fully enjoy their conversation is not that the audience of the conversation is actually restricted to just the two of them. Rather what matters is that the friends *believe* that the audience of the conversation is so restricted. Had the friends not come to suspect that they were being spied on, the tenor and substance of their conversation would not have changed and they would have presumably experienced the full value of their conversation.

I shall not attempt to produce an exhaustive list of activities that require restricted audiences for their full enjoyment, but here are a few of importance: conversing between friends, colleagues, siblings, etc., having sex, eating, exercising, dancing, working, singing, playing. The list could expand, no doubt, to include many pages. People often require restricted audiences to try things out not knowing if they will excel or embarrassingly fail (or just look ridiculous trying), to perform acts with a level of gusto they are not comfortable displaying in front of others, to perform acts others may find shameful, wrong, or offensive, to perform acts that deepen the bonds of friendship or love, or to do the things some believe must be done alone, such as grieving a loved one, choosing one's commitments, or sorting out one's political views.

Even if the suspicion of spying does not alter people's behavior, then, it may cut into the enjoyment they derive from some of their most cherished activities. Connected to this loss of enjoyment, but analytically separable from it, are a range of emotional responses that can follow from the suspicion of spying. A host of negative emotions

tend to follow the moment that someone learns they have been spied on. People feel anxiety and paranoia, for example. They wonder what the spy knows and how she might employ that knowledge. They wonder whether they are still secretly under observation. They look everywhere for additional signs of spying. They begin to question whether other agents are secretly collecting their information. Some hatch conspiracy theories.

Offense and humiliation can also follow from the suspicion of spying, when the spying is believed to occur exclusively or disproportionately on disadvantaged groups, groups that have been unfairly targeted in the past, or groups that are widely thought incapable of ordering their own affairs. In Muslim communities in the United States, for example, it is generally believed that Muslims have been unfairly targeted by America's intelligence agencies after 9/11. Hence, holding all other factors equal, in a suspected terrorism case in the United States spying on an imam could be more harmful than spying on a priest, since in the former case, revelation of the spying could further entrench humiliating stereotypes about Muslims being terrorists. Negative stereotypes of this sort often lead to profound harms of misrecognition.²⁵ Peaceful, law-abiding, even patriotic American Muslims are unjustly signaled out as violent supporters of terrorism, intent on America's demise. Their deeply felt identities are warped, demeaned, and demonized.

In egregious cases, the suspicion or revelation of spying, when it is disproportionately on disadvantaged groups, can go beyond causing offense to undermining the dignity of members of these groups. Waldron (2012) develops a parallel argument in his discussion of hate speech. He describes a person's dignity as follows.

²⁵Taylor (1994), for example, argues "... our identity is partly shaped by recognition or its absence, often by the misrecognition of others, and so a person or group of people can suffer real damage, real distortion, if the people or society around them mirror back to them a confining or demeaning or contemptible picture of themselves." (25)

A person's dignity is not just some Kantian aura. It is their social standing, the fundamentals of basic reputation that entitle them to be treated as equals in the ordinary operations of society. Their dignity is something they can rely on in the best case implicitly and without fuss, as they live their lives, go about their business, and raise their families. (5)

The principal problem with hate speech, Waldron argues, is that it can undermine this dignity. Hate speech aims "to compromise the dignity of those at whom it is targeted...it aims to besmirch the basics of their reputation, by associating ascriptive characteristics...with conduct or attributes that should disqualify someone from being treated as a member of society in good standing." (Ibid.)

Although spies usually do not aim to besmirch the basics of people's reputation, their spying still may have this effect. Return to the case of spying on Muslims after 9/11. Because unjustified spying on Muslims was relatively widespread, the effect was not just to offend isolated individuals. Instead, it signaled to Muslims more generally that they were no longer citizens in good standing. It undermined or contributed to undermining their dignity.

Consider the comparatively recent case of the New York Police Department's spying on Muslims, which *The New York Times* editorial page described as

widespread police spying and the creation of police records containing information on Muslim people, mosques and campus groups, as well as luncheonettes, dollar stores and other legitimate businesses owned and frequented by Muslims, with no apparent reason to think anything wrong was going on.²⁶

²⁶<http://www.nytimes.com/2012/03/04/opinion/sunday/surveillance-security-and-civil-liberties.html>

It is no surprise that this program has negatively impacted the respect and standing of Muslims in New York and elsewhere. The same article reported that Muslims were “reluctant to pray openly at mosques, join in faith-based groups, or frequent Muslim hangouts for fear of being watched and possibly tarred by “guilt by association.” It concluded pessimistically: “not just Muslims are threatened by this seemingly excessive warrantless surveillance and record-keeping. Today Muslims are the target. In the past it was protesters against the Vietnam War, civil rights activists, socialists. Tomorrow it will be another vulnerable group whose lawful behavior is blended into criminal activity.”

Of course not all of the emotional reactions that people have to learning about spying are negative ones. Spying can also lead to a sense of comfort and security. This effect, I think, is rather more common than most recognize. Many people feel more secure because they believe that certain qualified others are spying. They believe CIA, NSA, MI6, and other intelligence organizations around the world routinely spy, for example, to identify and prevent terrorist and other serious threats to the lives and well-being of people around the world, and they feel more secure because of this belief. They further feel more secure – both in their bodies and their property when they spot unmarked police cars patrolling their neighborhoods. They sigh in relief when they hear that there are undercover officers in public buildings and places scanning crowds for weapons.

Beyond evoking emotional responses, the suspicion of spying can also alter people’s beliefs about the trustworthiness of the spy. Since spying often breaks entrenched norms of appropriate observation, it prompts some to discount the degree to which they trust the spy and those by whom she is employed. The spy’s loss of trustwor-

thiness can be on balance either a benefit or a harm, however. If the spy is not trustworthy, her spying may beneficially reveal that fact to those who suspect or discover her spying. But if the spy is trustworthy, then her loss of trustworthiness may on balance be a harm, since an agent's capacity to carry out her plans often depends on others and, more specifically, on the beliefs others have that her interests and theirs are intertwined.²⁷ Hence for an agent to lose trustworthiness is typically for her to lose some power to carry out her plans.

Even more important than the spy's loss of trustworthiness is the loss of public trust that can follow from spying being suspected or discovered. If I learn that my neighbor has been spying on me, I lose trust in her, and I begin to wonder whether my *other* neighbors are spying on me. Similarly, if I learn my government is impermissibly collecting my health records, I begin to wonder whether they are also impermissibly collecting my bank records, my text messages, etc. Hence spying may not just diminish the trust people have in the spy, but also the trust people have in relevantly similar agents. Because spying is a secretive activity and not all – indeed perhaps astonishingly few – secretive activities are ultimately revealed, when revelation does occur it typically prompts one to wonder just how much more of the secretive activity is taking place of which one is unaware. Spying, in other words, induces a sort of paranoia: it feeds conspiracy theories, which erode the trust necessary for well-functioning government institutions.

Skeptics will no doubt claim that most incidents of spying will have negligible effects on the public trust. Williams (1973), for example, criticizes consequentialists for appealing to remote effects.²⁸ But this objection misses the point. My claim is not

²⁷Here I'm drawing on an "encapsulated interests" conception of trust. See Hardin (2002).

²⁸He says, "The certainty that attaches to these hypotheses about possible [remote or indirect] effects is usually pretty low; in some cases, indeed, the hypothesis invoked is so implausible that

that every incident of spying leads to an erosion of public trust, but rather that many incidents do. Watergate, for instance, had an adverse effect on the public's trust in the U.S. government. So too did revelations during George W. Bush's presidency that the NSA spied on Americans without warrants, and the revelation of the FBI's spying on Martin Luther King Jr. Nor is my claim that every incident of spying leads to a noticeable reduction in every person's trust, only that spying can lead to a loss of trust among those connected to the spying.

Historical evidence supports this claim. Consider, for example, reactions by anti-war activists' to COINTELPRO. William Bennett, a philosophy professor at Swarthmore, who learned from the revelation of the COINTELPRO files that he was under surveillance commented, "Sometimes you get the feeling that the FBI has everybody under surveillance." (Davis, 1992, 11) This reaction among activists was not unreasonable, considering paranoia may have been precisely what the FBI intended to create. Davis cites a memo from the FBI's Philadelphia field office that encouraged agents to intensify their efforts so as to promote "the paranoia endemic in these circles and to further serve to get the point across that there is an FBI agent behind every mail box" (10)

Hence spying can not only erode the trustworthiness of the spy it can also erode the public trust necessary for the proper functioning of government institutions. The final potential consequence of suspected spying that I want to consider is that it can lead to retaliation against the spy. The costs of retaliation for individuals can be considerable, including torture or even death. These dangers are no doubt most

it would scarcely pass if it were not being used to deliver the respectable moral answer, as in the standard fantasy that one of the effects of one's telling a particular lie is to weaken the disposition of the world at large to tell the truth."

apparent when the spying is between countries and the stakes are high. Nathan Hale, for example, one of America's earliest spies, was captured and hanged by the British during the Revolutionary War. (Rose, 2007) More recently, the treachery of Aldrich Ames led to the death or imprisonment of many informants in the Soviet Union spying for the United States.²⁹

5.3 Conclusion

My purposes in this chapter have been to collect some of the more significant consequences of spying and to determine under what conditions these consequences are harms and under what conditions they are benefits. There is much more to be said about the consequences of spying than that it violates people's privacy. Spying can undermine cherished liberal and democratic rights, even as it is crucial for protecting these rights.

With the consequences of spying in place, I can now determine which principles the two-level utilitarianism that I argued for in the previous chapter supports for spying. I take up this task in the next chapter.

²⁹See CIA Directory Deutch's "Statement to the Public on the Ames Damage Assessment," where he claims, that "By revealing to the Soviet Union the identities of many assets who were providing information to the United States, he not only caused their executions, but also made it much more difficult to understand what was going on in the Soviet Union at a crucial time in its history."

5.4 Appendix: The Consequences of Spying

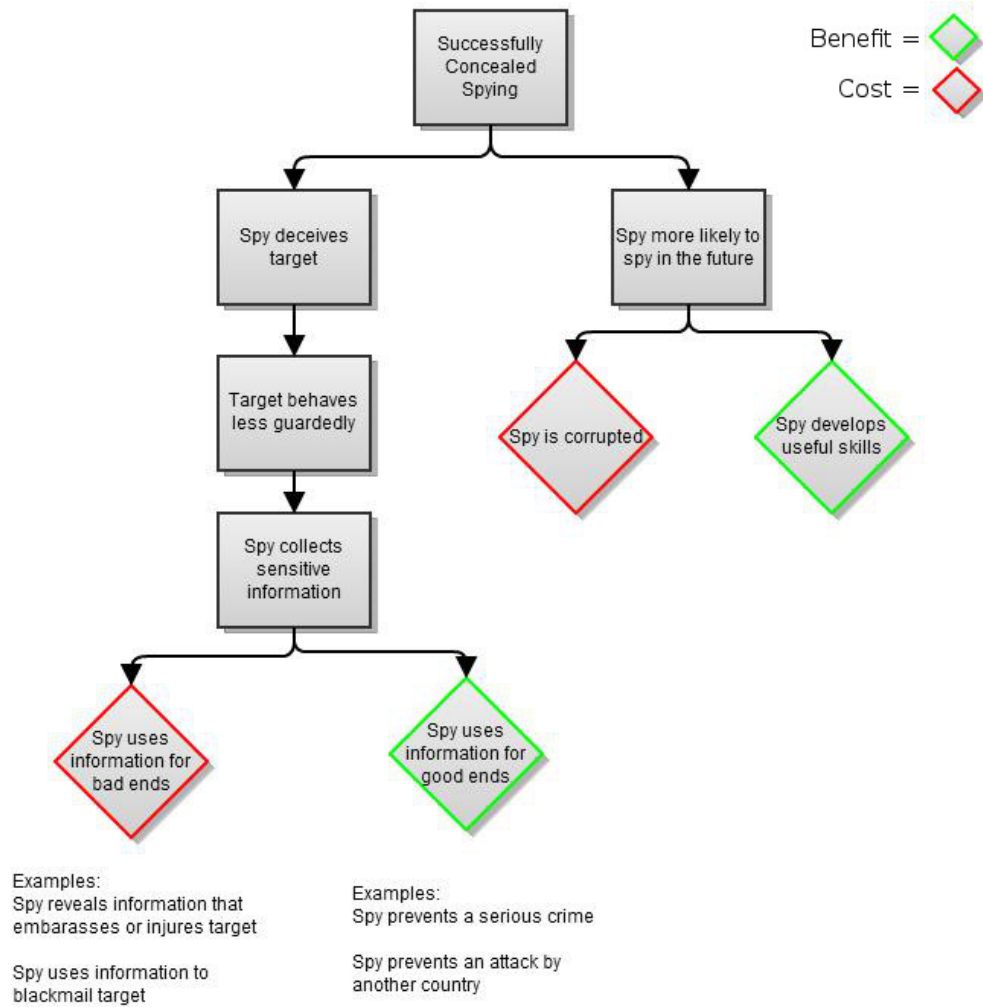


Figure 5.1: The Consequences of Successfully Concealed Spying

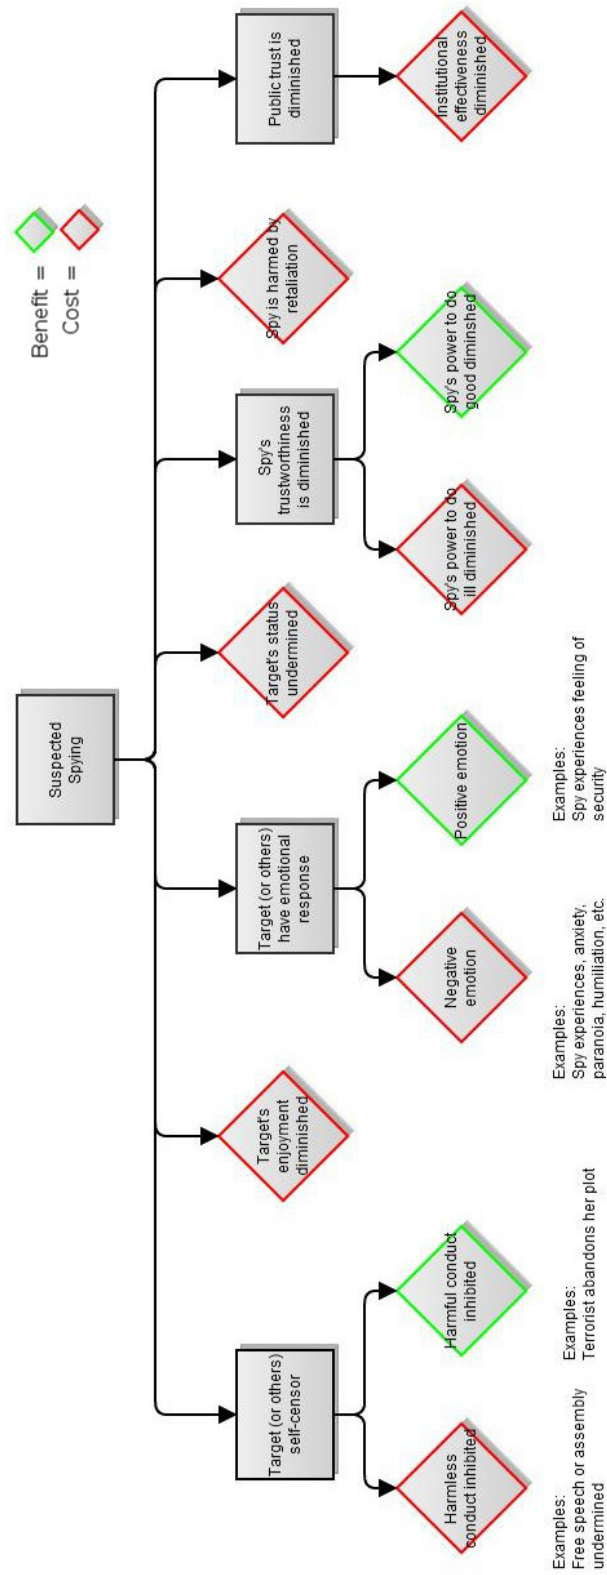


Figure 5.2: The Consequences of Suspected Spying

Chapter 6

Utilitarian Principles for Domestic Spying

My goal in this chapter is to design the intuitive principles supported by two-level utilitarianism for domestic government spying. My strategy for doing so is inductive. I begin from an outright prohibition on spying and then search for neighboring principles that are “better,” i.e. likely to produce greater net benefits.¹ When I find a better principle in this respect, I repeat the process and look for an even better neighboring principle. I continue to repeat this process until it seems further improvements are not possible.

The chapter’s main finding is that two-level utilitarianism supports the five principles developed in chapter two: just cause, proportionality, necessity, minimization, and discrimination. Since both widespread intuitions and utilitarianism, an ethical theory often thought to diverge strongly from our substantive moral intuitions, support these principles, the principles stand on strong normative foundations.

¹By “neighboring principle” I mean a principle that is identical in all respects but one.

I begin the chapter by articulating a set of general considerations for the design of intuitive principles, gleaned from the discussion of TLU in chapter 4. These general considerations helpfully guide my search for neighboring principles. Next, employing the method sketched above, I work toward the optimal utilitarian principles to regulate domestic government spying. Following the order I followed in chapter 3 I then argue that these principles should be institutionalized. The fourth section of the chapter briefly develops two ancillary principles regulating the purchasing and sharing of information that are required for the efficacy of principles of domestic government spying. Section five returns to four objections raised in chapter 4, and the final section is a short conclusion.

6.1 General Considerations for Utilitarian Principles

By collecting and synthesizing some of the insights from chapter 4, I have developed four considerations for the design of good utilitarian principles. Good principles limit the calculative demands on agents, they are simple, they protect against special pleading, and they are strategic. Although good utilitarian principles tend to meet these criteria, it is not necessarily the case that the *best* principles, from the utilitarian perspective, will be the most strategic principles, the simplest principles, etc. The four considerations merely assist the search for optimal principles, since the best principles are exceedingly unlikely to be highly complex, unstrategic, etc.

Good utilitarian principles, first, should limit the calculative demands they place on agents. As we saw in chapter 4, agents often have restricted time and information to make decisions, they tend to be poor probabilistic reasoners, they sometimes fail

to give their decisions full consideration because of weakness of will, and they tend to fudge calculations, favoring themselves or those they are in special relationships with over others. Hence the best rules will be those that do not ask agents to engage in intensive calculations.

Second, good utilitarian principles are simple. They are simple enough to be learned and deeply ingrained as dispositions or character traits. They can be employed automatically, and connected to habitual patterns of judgment and emotional response.

Good utilitarian principles, third, protect against the human tendency to engage in special pleading. To some extent this is achieved by simple principles, strictly followed. But agents can twist even simple principles to unfairly benefit themselves. Thus, good principles will tend to specify deliberative procedures, which, when followed, will demonstrate both to the agent and others that precautions are in place to minimize special pleading.

Finally, good principles are strategic. They encourage good responses from those who learn that they are being followed. For spying, strategic principles are particularly important because of the role (discussed in Chapter 5) that people's expectations about their audiences play for determining their conduct. If people believe that the government observes or is likely to learn of their conduct, then they will be more likely to condition their behavior so it conforms to norms enforced by the government and its officials.

6.2 Developing Utilitarian Principles

Perhaps the best set of principles for government spying includes only one principle prohibiting all government spying. An outright prohibition is simple, it places

almost no calculative demands on agents, and it protects against special pleading. Government agents would not be permitted to spy for selfish or otherwise morally dubious reasons if they were not permitted to spy at all.

The benefits of a complete prohibition against spying, if people generally complied with it, would be considerable. Individuals could develop stable expectations about where, when, and by whom they are being observed. They could perform harmless counter-normative activities in private places without fear of sanction. They could develop their thoughts free from the gaze of others. Safe places would exist where people could speak out against their government or against their employer, without fearing reprisal. People's information would also be relatively secure from government intrusion. They would not have to worry, for example, whether government agents monitor their purchases, the library books they read, or for whom they cast ballots.

Nevertheless, a blanket anti-spying principle is suboptimal, since it leaves too many potential benefits on the table. The world we live in includes bad or seriously misguided people who often perpetrate awful acts profoundly harming the lives and interests of many people. Government spying can often prevent these acts at a comparatively low cost. So, the simple anti-spying principle is not the best principle.

Those who deeply distrust government might worry that any set of principles that permits government spying would lead to abuse the harms of which would outweigh whatever benefits government spying might procure. If you give the government an inch, they think, it will take a mile. Despite pointing out an important concern for the consideration of utilitarian principles, namely that government spying will sometimes be misused, this objection is mistaken. It is mistaken because it ignores the experience of many liberal democracies. Government spying, in many liberal democracies, has played a crucial role in preventing unspeakable harms, while threatening the rights and welfare of citizens in comparatively minor ways. The risk of abuse is

an important concern for the design of utilitarian principles for spying, then, but it is not so great that it favors prohibiting government spying altogether.

Since the problem with the anti-spying principle is that it leaves significant benefits on the table, it seems natural to reformulate the principle by prohibiting spying except when it is likely to lead to significant benefits. A different problem, however, emerges for this reformulation. The problem is that “significant” is vague. Preventing a terrorist attack obviously seems significant, but for a whole range of other benefits it is difficult to say whether they are significant or not. Certainly reasonable people can disagree about what counts as significant, and vagueness of this sort courts both abuses and mistakes. Agents who wish to spy only to further their own political or personal interests can stretch “significant” to fit their purposes, and agents who have mistaken ideas about what counts as “significant” can carry out costly spying to little or ill effect. Vagueness also curtails accountability. One of the conditions of effectiveness for an accountability relationship is that the accountability holder has clear appropriateness standards for the accountable agent’s conduct. When standards of appropriate conduct are vague, it is difficult for accountability holders to agree or to make a clear case that the accountable agent has acted inappropriately.

The vagueness problem plausibly can be solved by admitting a short list of clearly articulated exceptions to the anti-spying principle. But then a difficult question emerges: which cases should count as exceptional? From a utilitarian perspective, exceptional cases should be those that experience suggests are likely to produce net benefits. As a rough cut, then, it might make sense to admit exceptions *only* for spying that can prevent an agent from (wrongfully) killing, torturing, raping, imprisoning, or kidnapping another agent. This rough cut suggests that exceptions to the anti-spying principle should be admitted only when they promise to deliver benefits of a particularly high level of *intensity*. Cases which approach or achieve the same

level of net benefits by delivering relatively low intensity benefits for long durations and/or to wide swaths of the population are ruled out as permissible exceptions.²

One might wonder, however, whether narrowing the range of permissible exceptions to the anti-spying principle to high intensity benefits is justifiable on utilitarian grounds. Two reasons suggest such narrowing is justified. Lower intensity more diffuse benefits are usually more difficult to accurately predict *ex ante*, making it difficult to articulate exceptions; and, more importantly, narrowing permissible exceptions based on the intensity of benefits permits a simple list of exceptions, making the identification of exceptions comparatively simple both for government agents and their accountability holders.

But arguably even this revised principle, which prohibits spying except in clearly articulated cases, still leaves too many benefits on the table. Some spying, for example eavesdropping in a public place, does not typically produce considerable harms. If these relatively harmless forms of spying promise to produce nontrivial benefits, *even if* they do not promise to prevent harms to life and limb, it seems reasonable on utilitarian grounds to permit them, so long as permitting them does not appreciably increase the risks of mistakes or abuse.³

The problem raised by relatively harmless spying can be addressed by dividing spying into more and less harmful types. How finely spying should be divided is a matter of balancing two concerns. If we divide spying too finely, then the number of categories will exceed government agents' capacities to learn the rules and to apply

²I think this is the best way to interpret the principle of just cause in just war theory. A plausible reading of McMahan's (2005) conception of just cause, for instance, is that it is concerned only with the intensity of benefits, not with their duration, scope, or likelihood.

³A similar problem does not present itself in the case of war, since war always risks killing and maiming on a grand scale. There may be *relatively* harmless wars, but even in these, the harms are of immense proportions.

them efficiently. If, on the other hand, we divide spying too coarsely, then our principles will lead to too many mistaken recommendations in particular cases.

I want to suggest these two concerns can be balanced with a simple division of spying into four categories based on two distinctions. Spying has the potential to lead to higher intensity harms when it is *aided* and when it attempts to collect *private* information. By aided spying, I mean spying that is not performed *only* with the spy's five senses, the spying is in some way assisted by technology. Bugging the neighbor's house is aided spying, then, while eavesdropping at the neighbor's window is unaided spying. By private information I mean information that is concealed or protected from observation using methods acknowledged by social norms or established by laws. Hence, information on a person's (personal) computer is private information, while information in the newspaper is public information. Further, a person's conduct and her communications, when they occur in private places (i.e. places where social norms or laws dictate restricted observation) count as private information.

Spying that seeks to collect private information has the potential to lead to greater harms because individuals tend to protect sensitive information, that is information that could harm, embarrass, or be leveraged to manipulate themselves or others, by attempting to conceal it. Since most private information is protected by social norms that we all (or at least a great majority of us) follow, the target is also more likely to feel violated, and more likely to censor her future behavior when norms about appropriate information collection have been violated.

Aided techniques, such as wiretapping, hacking, or body scanning, are more likely to successfully produce private information because they allow spies to peer into places where people typically do not expect that their information can be collected. Aided tactics, when they are discovered, also have stronger effects on people's expectations about when, where, and by whom they are being observed. Discovering an electronic

listening device in your home or office, for instance, has a more profound effect than finding someone's ear pressed to your door, since the former suggests the spy has considerably greater capabilities to penetrate into your inner world. The more you believe a spy can peer into your inner world, the more you are likely to feel violated, the more you are likely to censor your future behavior, and the more you are likely to invest in expensive countermeasures.

Based on these two distinctions, the harms of spying are likely to be most intense when it is aided and when it aims to collect private information (when it is "private" for short), and they are likely to be least intense when it is unaided and non-private (or "public"). Two categories sit between these extremes: spying that is aided and public and spying that is unaided and private. I shall assume, however, that these two categories lead to harms of the same intensity, thereby reducing the number of categories to three.⁴

So we have now divided spying into three more and less harmful types. The payoff of this division is that it permits us to circumvent the "harmless spying" objection. For spying with lower intensity harms, more exceptions can be admitted. In other words, the justificatory bar for aided/private spying will be considerably higher than the bar for unaided/public spying. The former may require the prevention of *serious* harms, while the latter may require only the prevention of relatively pedestrian offenses. The table below illustrates the three categories and suggests a tentative set of just causes for each.⁵

⁴Another simplifying assumption I am making is that these two criteria are dichotomous. In fact they are continuous variables. Places are not simply public or private, they can be recognized as being *more or less* public or private than others. For the purpose of making simple, easily learned and utilized rules, however, it is helpful to dichotomize the criteria.

⁵We can of course imagine significantly more harmful forms of spying by conjoining spying with other wrongs (e.g. breaking and entering, blackmail, etc.), and no doubt some of these more harmful

One way to interpret the list of exceptions is that they provide “right reasons”

Table 6.1: Just Causes for Spying

Type of Spying	Just Causes
Aided/Private	Violent harms (e.g. war, terrorism, murder, armed robbery, kidnapping, etc.)
Aided/Public or Unaided/Private	Non-violent but serious harms (e.g. drug crimes, fraud and theft of large amounts, etc.)
Unaided/Public	Misdemeanors (e.g. minor property damage, theft and fraud of small amounts, traffic crimes, etc.)

or “just causes” for spying. Spying is only permissible, on this interpretation, when it meets the principle of just cause. This interpretation is useful since it highlights a connection to the well established and highly sophisticated literature in just war theory. The principle of just cause is central to just war theory, though it is typically supplemented by a set of principles, such as proportionality, last resort, and discrimination. Just cause by itself is insufficient to regulate the conduct of war.

The same is true in the case of government spying: just cause by itself is not an optimal principle. The chief reason is that the right reasons that underlie the principle of just cause pick up only the intensity of the harms that follow from spying, while they hold the duration of the harm and the number of individuals harmed constant at an average or “typical” level. But since the duration of spying and the number of people harmed by spying obviously vary widely case to case, there is a danger that even the least costly form of government spying can become very costly as it grows

forms are justified to prevent catastrophic harms. I focus here strictly on spying, and assume that further analysis would be required to assess the just causes for these more harmful types.

in duration and in targets. Eavesdropping in public places, for example, if it were performed for many years on thousands or tens of thousands of people, could turn out to generate substantial harms.

A reasonable solution to this problem is to embrace a proportionality constraint. On a common interpretation, the principle of proportionality demands that the expected net benefits of spying be positive, that is it demands that the expected benefits of spying exceed the expected costs.⁶ On this interpretation, the principle of proportionality requires that government agents engage in a restricted cost/benefit calculation. The calculation is restricted because the agent does not examine the costs and benefits of *all* of her available options, only the costs and benefits of the particular form of spying under consideration.⁷ Although this is an important restriction, since agents can sometimes be faced with a seemingly infinite number of options, the principle of proportionality, on this interpretation is still too calculatively demanding, since it requires government agents, who often have to make decisions quickly, to identify and assign likelihoods to all of the relevant harms and benefits of spying. As a decision procedure, it is too calculatively complex and thus too prone to error and misuse.

The principle of proportionality is more attractive, however, when it is interpreted in a practical way, i.e. when it asks government agents to determine whether it is reasonable to think that the likely harms of the spying are *clearly* outweighed by the

⁶As Orend (2005) says (referencing the proportionality condition in just war theory), “Only if the benefits are proportional to, or “worth”, the costs may the war action proceed.”

⁷Hurka (2005, 38) suggests that the proportionality condition (in the context of just war) ought to be interpreted in a comparative way. Thus one option would be said to be more/less proportionate than others. This interpretation is elegant because it can collapse the last-resort and necessity conditions into the proportionality condition, but I think it should be resisted, since it produces a principle that is too calculatively demanding.

assistance spying likely lends toward achieving the just cause.⁸ In Sidgwick's (1891, 254) words, proportionality (on its practical interpretation) prohibits "any mischief of which the conduciveness to the end is slight in comparison with the amount of mischief." The chief role of the practical principle of proportionality, then, is not to stand in for the consequentialist calculus, it is to rule out cases of spying (not already filtered out by the principle of just cause) in which costs are clearly likely to exceed benefits.

The proportionality principle is also more attractive when we limit the goods that count in the calculation. Some think we should limit the goods that count toward proportionality on moral grounds. Hurka, for example, reaches this conclusion for just war's proportionality constraint by examining the following case.

Imagine that our nation has a just cause for war but is also in an economic recession, and that fighting the war will lift both our and the world's economies out of this recession, as World War II ended the depression of the 1930s. Although the economic benefits of war here are real, they surely cannot count toward its proportionality or make an otherwise disproportionate conflict proportionate.

I am inclined to reject Hurka's substantive moral intuition in this case: it does not seem *obviously* mistaken that economic growth could tip the scales in favor of going to war. Economic growth often profoundly and positively affects people's well-being, and it does not seem intuitively plausible that effects on well-being should be left out of moral calculation. Yet there is still the *practical* question of whether the principles government agents use to determine proportionality ought to limit the goods used in

⁸(Lackey, 1989, 40-41) seems to take a similar position in the just war context. He says, "a war for a just cause passes the test of proportionality unless it produces a *great deal* more harm than good."

the calculation.

Without such practical limitations, government agents would be asked to identify and weigh remote consequences. Yet adding remote consequences substantially adds to the complexity of the proportionality calculation and therefore renders it more prone to mistakes in application and abuse. If the calculation of proportionality is limited to those goods contained in the just causes, it is more difficult for government agents to botch or fudge calculations, since it is clearer to accountability holders which goods belong in the proportionality calculation.

This argument to limit the *goods* used in the proportionality calculation seems to extend also to the *evils* used in the proportionality calculation. But I do not think the evils used in the proportionality calculation should be limited. The reason has to do with the kinds of errors that are likely to follow from the misapplication of principles of government spying. Law enforcement and intelligence agents, historical experience suggests, tend toward faking or inflating the case to spy, not toward downplaying it. By restricting the goods in the proportionality calculation without also restricting the evils, the principles of government spying can correct for this tendency.

The addition of the principle of proportionality, then, may be sufficient to ensure that government agents do not carry out spying that is on balance harmful. It is not sufficient, however, to ensure that government agents choose the least harmful option available to them to secure the just cause, a requirement that is clearly demanded by utilitarianism.⁹ One way to require government agents to choose the least harmful available option to secure the just cause is to simply make the requirement a principle standing along side just cause and proportionality. But adopting such a principle is not optimal, since doing so would not clearly indicate to government agents what

⁹See my discussion of classical utilitarianism in Chapter 4.

is required of them to guarantee that the alternative they select is in fact the least costly alternative capable of securing the just cause.

This clarity can be achieved by adding two principles along side just cause and proportionality rather than just one. Government agents should first compare spying to alternatives that are likely to be less harmful and that could be reasonably certain to secure the just cause. In other words, they should determine whether spying is *necessary* to secure the just cause. In particular cases there may be a variety of less harmful alternatives to spying that could secure the just cause. But overt methods of collecting information, such as reading the newspaper, searching the internet, or interviewing potential sources will most often be the relevant comparison set. Government agents should either try less costly available options before spying or they should demonstrate that it is reasonable to think that these less costly alternatives would fail.

Assuming no less harmful alternatives to spying can secure the just cause, government agents should, second, ensure that the harms of spying are *minimized*. The minimization of harms entails eschewing more intrusive tactics when less intrusive tactics are likely to succeed. It is never permissible, for example, to use aided spying when unaided spying is likely to secure the just cause. It also suggests taking a range of precautions to avert harms, such as limiting the duration of spying, securing the information collected from spying, discarding information unrelated to the just cause, using the most reliable available personnel, and incentivizing agents to collect and store information according to appropriate procedures.

Thus far I have argued from a simple prohibition against government spying to four principles for government spying: just cause, proportionality, necessity and minimization. Agents following these principles are likely to capture many of the potential benefits of spying. They are unlikely to engage in unnecessary harmful spying. They

will not spy when less costly alternatives are available to secure the just cause. And, since they will tend to select the least harmful kinds of spying likely to secure the just cause, they will tend to promote the *net* benefits of spying.

Yet I want to suggest that these principles remain suboptimal in two respects: they are not strategic, and they fail to thoroughly protect against abuse and misapplication. A principle is strategic when it encourages good responses from those who learn that the principle is being followed. Hence a principle is strategic, first, when it encourages only beneficial actions and discourages only harmful ones. Principles can fail to be strategic, then, by failing to create good incentives, or even worse, by creating perverse incentives.

Which principles are strategic in this first sense depends on people's expectations about the rewards and punishments that are likely to follow if their actions are discovered. In an unjust state that punishes many beneficial actions and rewards many harmful actions, an outright prohibition against spying is probably strategic, since prohibiting spying opens up space for beneficial actions, while reducing the likelihood of people being unjustly punished. In a perfectly just state that punishes only harmful actions and rewards only beneficial actions, in contrast, strategic principles permit a wide range of spying, since spying increases the likelihood that harmful behavior will be punished, while at the same time it increases the likelihood that beneficial behavior will be rewarded.

In most liberal democracies laws and social norms tend to punish harmful actions, especially severely harmful actions. The same laws and norms less frequently reward beneficial actions. Further, every liberal democracy punishes some harmless actions, and no liberal democracy can fully guarantee that its public officials will not wrongfully punish citizens for prudential or political advantage. People as a consequence expect that if they commit serious harms and they are caught, they will be punished.

But some also worry that they will be punished for harmless but counternormative actions. Finally, people worry about the potential that powerful public officials have to punish their legitimate dissent. Hence in most liberal democracies it is reasonable to think that strategic principles will signal to people that they risk being spied on when they engage in *and only when they engage in* harmful activities. Such principles do not encourage good deeds, but they do discourage harmful conduct, and they protect citizens from unwarranted interference into harmless activities by the government.

This line of argument is supported by a further set of responses that people might have to learning how their government regulates domestic spying. When principles unfairly or unequally target certain groups, they can demean, humiliate, and disrespect members of those groups when they become public. Principles can also have these effects if they signal to people their chosen pursuits are unworthy, shameful, or depraved. People's self-respect often depends on the existence of spaces for action free from government intrusion. Further, when citizens worry that they are under covert observation by their government, there are a range of activities that can become less enjoyable because they are less private. Finally, when citizens suspect that the government spies on them, they may lose trust in their government and its institutions.

Let me now explain why the four principles that I have developed thus far to regulate domestic government spying are *not* strategic. Consider the following case. America's National Security Agency has developed a new piece of software, code-named "Dragnet." Dragnet is designed to identify and prevent domestic terrorist attacks. By collecting, storing and analyzing the internet searches, posts to social networking sites, financial transactions, text messages, and phone calls of ordinary Americans, Dragnet can sometimes identify terrorists and terrorist plots in action.

Notice that the dragnet case could plausibly meet all of the four principles dis-

cussed above. The problem with cases like Dragnet is that ordinary people reasonably fear abuse when their personal information is collected in droves; they also worry that their harmless counter-normative activities will be exposed, punished, or leveraged against them. They want law enforcement officials and intelligence agents to possess thick files on suspected terrorists and criminals, but they worry about these same government agents having dossiers on law-abiding citizens. The temptations to misuse information for prudential or political gain, they think, are too high. So too are the temptations to police norms prohibiting harmless or even beneficial behavior. Indeed, the more innocents believe that the government possesses detailed dossiers on them, the more likely they will be to self-censor, especially when the conduct challenges the government or its officials. Principles that permit operations like Dragnet, then, are not strategic, since they discourage people from engaging in a range of harmless or beneficial behaviors. They further risk humiliating or disrespecting those who engage in harmless but counternormative activities, fomenting paranoia, undermining trust in the government, and diminishing people's enjoyment of a range of private activities.

Adding a principle to protect innocents can render the four principles more strategic. It can signal to innocents that they are strongly protected against spying by their government, while signaling to non-innocents that their harmful conduct is likely to be discovered and punished. But utilitarians obviously cannot support an absolute protection for innocents, since in a plausible range of cases (as the case of the terrorist and his family in chapter 3 suggests) spying on innocents may be required to avert catastrophe. What the utilitarian requires, then, is a principle that protects innocents, but does not afford them *absolute* protection.

I want to suggest that the principle of discrimination (discussed in Chapter 3) can play this role. The principle of discrimination, recall, does not absolutely reject

spying on innocents. It merely requires that government agents not *intend* to spy on innocents. Innocents, in other words, cannot be the target of spying, though they can be the target of observation. The principle is strategic, since it signals to innocents that they cannot be targeted, while at the same time it leaves very few benefits on the table. It rightly (from the perspective of the utilitarian) permits spying in the case of the terrorist and his family and it prohibits spying in the Dragnet case.

The inevitable objection to discrimination as a utilitarian principle is that we can imagine cases in which targeting innocents leads to benefits large enough that they outweigh the harms of targeting innocents. The question, however, is whether such hypotheticals actually constitute an objection to the principle of discrimination. Every principle, when placed in the hands of real-world agents, inevitably leads to mistakes (from the perspective of the utilitarian's standard of right) in particular cases. Even the optimal set of principles – that is those that guide agents to “make the highest proportion of right decisions in actual cases where their decisions make a difference to what happens weighted, of course, for the importance of the cases, that is, the amount of difference the decisions make to the resulting good or harm” – occasionally lead those who employ them astray. (Hare, 1979, 115) Hence pointing to hypothetical cases in which a principle fails to live up to the consequentialist's standard of right is not a persuasive objection. A persuasive objection to the principle of discrimination on consequentialist grounds would have to point to a principle capable of producing larger net benefits in real-world circumstances.

I have now argued that the five principles to regulate government spying that I developed in Chapter 3 – just cause, proportionality, necessity, minimization, and discrimination – are also defensible on utilitarian grounds. But my assumption thus

far has been that government agents will faithfully apply these principles.¹⁰ This assumption, I want to suggest, is implausible. Men are not angels. Even the most upright government agents can be expected to occasionally buckle under the weight of temptation. “If men were angels,” Madison (2009, 264) argued in Federalist 51, “no government would be necessary. In framing a government which is to be administered by men over men, the great difficulty lies in this: you must first enable the government to control the governed; and in the next place oblige it to control itself.” The five principles that I have developed can guide government agents toward ethically employing spying to control the governed. But the principles cannot guarantee that government agents will control themselves. Accordingly, I want to suggest that the five principles should be institutionalized, in order to ensure that government agents conscientiously apply them. In the final two chapters, I argue for *how* these principles should be institutionalized; here my purpose is only to make the normative argument *that* they should be institutionalized.

The utilitarian argument for institutionalization has been made in a series of American legal cases connected to the Constitution’s Fourth Amendment. Justices have pointed out that the principal goal for law enforcement officials is to prevent and punish crimes. Because spying is such a useful tool for these ends, government agents are routinely tempted to spy, even when spying does not meet appropriate standards. This temptation can lead them to (consciously or unconsciously) misrepresent or even manufacture the case for spying. Hence law enforcement officers’ strong institutional interests stand in the way of them consistently applying legal principles objectively or unbiasedly. Those detached institutionally from tracking and arresting suspects,

¹⁰The question of which principles are optimal for regulating government spying cannot be separated from the institutional environments in which the principles are to be applied. This point is eloquently made in the context of just war by Buchanan (2006).

the argument goes, are better placed to faithfully apply the law. In Justice Jackson's words,

The point of the Fourth Amendment, which often is not grasped by zealous officers, is not that it denies law enforcement the support of the usual inferences which reasonable men draw from evidence. Its protection consists in requiring that those inferences be drawn by a neutral and detached magistrate, instead of being judged by the officer engaged in the often competitive enterprise of ferreting out crime. Any assumption that evidence sufficient to support a magistrate's disinterested determination to issue a search warrant will justify the officers in making a search without a warrant would reduce the Amendment to a nullity, and leave the people's homes secure only in the discretion of police officers.¹¹

Over two decades later Justice Stewart summarized Jackson's reasoning, saying "[T]he whole point of the basic rule so well expressed by Mr. Justice Jackson is that prosecutors and policemen simply cannot be asked to maintain the requisite neutrality with regard to their own investigations."¹²

Although this argument is aimed at law enforcement officials, it applies just as well for intelligence agents. The chief aim for intelligence agents is to identify and prevent foreign and domestic threats and to root out spies from other countries. These aims, just like the aims of law enforcement, no doubt place professional or institutional blinders on those who possess them. There is no reason to think, then, that the FBI agent would apply the law more faithfully than the local police officer.

This argument made in American case law plausibly could be derived from a more

¹¹ *Johnson v. United States*, 333 U. S. 10, 333 U.S. 13-14 (1948)

¹² *Coolidge v. New Hampshire*, 403 U.S. 443 (1971)

general argument with deep roots in the history of political thought: *Nemo iudex in causa sua*. Or, no one should be a judge in their own cause. As Locke (1988, 275-276) says, "...it is unreasonable for Men to be Judges in their own Cases...Self-love will make Men partial to themselves and their friends." Hobbes (1994, 98) states the case even more starkly: "seeing every man is presumed to do all things in order to his own benefit, *no man is a fit arbiter in his own cause*."

The case for institutionalization is further strengthened by a careful consideration of three of the consequences of spying discussed in the previous chapter. One worry in that chapter, recall, was that spying could lead to the corruption of the spy. Institutionalizing principles for government spying, however, mitigates this risk. Institutions not only reduce the incidences of wrongful spying, but they also serve as a concrete reminder to government agents and their accountability holders of what counts as permissible spying. A second worry was that government spying could evoke anxiety or paranoia in the citizenry, and when government spying is done disproportionately on disadvantaged groups, that government spying could undermine the status of individuals in these groups. Institutionalization can blunt this concern by signaling to citizens that the decision to spy is not just in the hands of a few agents strongly incentivized to spy. Indeed, it is difficult to see how the principles argued for above could be highly strategic without being institutionalized. It is one thing to *say* government agents will follow a set of principles; it is an altogether different thing to *compel* them to follow the principles. Finally, since spying has the potential to diminish the public's trust in governmental institutions, the best way to secure this trust is to signal to the public that measures are in place to ensure the government's power to spy is not abused.

In short, since law enforcement officials and intelligence agents have so much to gain professionally and their handlers have so much to gain politically from spying,

and since citizens have so much to lose from government agents abusing their powers to spy, these government agents should not be the only arbiters of how the five principles of government spying apply.

The conclusions reached so far in this section can now be summarized as follows.

1. Government agents should follow the following principles for domestic spying:
 - (a) Just Cause = spy only for an enumerated set of right reasons.
 - (b) Proportionality = spy only when the likely harms of the spying are *clearly* outweighed by the likely assistance spying lends toward achieving the just cause.
 - (c) Necessity = spy only when less costly alternatives have been tried or they can be ruled out as unlikely to be successful.
 - (d) Minimization = spy only when the tactics utilized are the least harmful tactics likely to secure the just cause and hence all reasonable precautions have been taken to minimize harms.
 - (e) Discrimination = spy only when the principal target of the spying is reasonably believed to be engaged in or assisting the harm that government agents aim to prevent in securing the just cause.
2. These principles ought to be institutionalized in such a way that the government agents whose activities should be regulated by them are not the only arbiters of how the principles apply.

The worry one many might have about how I have developed the utilitarian principles is that they were determined, in large part, by my starting point: an outright prohibition on spying. In order to allay this worry, I will begin the process of identifying

utilitarian principles from the opposite extreme. That is, I will begin from the hypothesis that all government spying should be permitted, and then work inductively toward optimal principles. Rather than developing principles in detail as I did above, however, my aim is merely to persuade the reader that the road leads in the same direction.

The case against permitting all government spying is hugely overdetermined. If government agents were always at liberty to spy, people could not develop stable expectations about the audience for their actions (without expending considerable amounts on countermeasures), nor could they guarantee their personal information would remain concealed. The enjoyment of private goods would rapidly diminish. People's autonomy would be gravely threatened, since the pressures to conform to social norms would be virtually unchecked. The social costs of stifling free thinking and experimental ways of life could be considerable. Further, the benefits of such a permissive policy would be minimal. One might think that permitting all government spying would lead to the prevention of many serious (and even less serious) harms. But permitting all spying is more likely to lead to a whole lot of very ineffective or even harmful spying.

The way to improve on the maximally permissive policy is to selectively prohibit certain kinds of spying. Yet once we begin considering such prohibitions, we are faced with the same considerations that we faced above. On one hand, we want to stabilize people's expectations about when, where, and how the government can spy, and we want the expectations that we stabilize to discourage harmful behavior (and to the extent possible encourage good behavior). On the other hand, we want to make it possible for law enforcement officials and intelligence agents to employ spying to prevent the most harmful actions, while making it difficult for them to misuse spying for prudential or political benefit. There are obviously many ways one could go about

designing principles to meet these objectives. But one way is to adopt the principles we reached above when we worked backward from an outright prohibition. Hence, it seems plausible to think that we will reach the same utilitarian principles for spying regardless of our starting point. Whether we work from the most or the least restrictive policies, by making the same set of assumptions, we will land on the same set of principles.

6.3 Ancillary Principles

So both widespread intuitions *and* two-level utilitarianism support the same five principles for domestic government spying. What I want to briefly point out in this section is that the five principles cannot be efficacious without governments endorsing certain ancillary principles. Two in particular come to mind.

First, the principles of domestic government spying cannot be efficacious if governments do not regulate the purchasing of information from parties that collect information covertly. No matter how well the five principles are institutionalized, if governments can circumvent them by having contractors do their dirty work, then the principles will be no more than parchment barriers. This ancillary principle is particularly important now that governments buy untold quantities of information from private corporations.

Similarly, the principles of domestic government spying cannot be efficacious if governments acquire information about their own citizens that was covertly collected by foreign governments. National intelligence agencies have increasingly agreed to monitor each other's citizens and then to exchange the information they collect, so as to circumvent domestic constraints. Sepper (2010, 173) mentions a number of such cases. For example, the German BND used the European Counter-Terrorist Intelli-

gence Center in Paris to collect information from German law enforcement agencies it would not be permitted to read according to German law. In Norway, the CIA was given free rein to investigate Muslim groups, without adhering to Norwegian law. To support the claim that the practice of circumvention is fairly widespread, Suskind (2006, 85) mentions George Tenet's remarks at a meeting of UKUSA intelligence chiefs, in which he claims the "shackles" of domestic spying laws will "at the very least, be loosened, if not in practice discarded" by international intelligence sharing. Tighter sharing agreements between state intelligence agencies, it seems, threaten considerably the efficacy of domestic constraints, and given that threats to a state's security increasingly cross international borders, one might expect that intelligence sharing agreements will only become more comprehensive.

6.4 Objections

Let us finally return to four of the objections to classical utilitarianism raised in chapter 4 and see if TLU circumvents them. Note that my defense of the above principles does not hang on my ability to rebut these objections, since I have argued that the principles have multiple sources of justification. But my view is that TLU is the most plausible moral theory, and I want to demonstrate to my reader why I think this is the case. The first objection is that classical utilitarianism does not take autonomy seriously enough. Is this claim true for TLU?

The claim that (my rendering of) TLU does not take autonomy seriously enough could be interpreted in two ways. The first is that I have misspecified the optimal principles. The right specification of intuitive principles, on utilitarian grounds, would give more protection to people's autonomy. The second interpretation is a deeper philosophical point. On this interpretation, the claim is that TLU has an

incorrect value theory. The “correct” theory of value gives a more central place to autonomy. Autonomy is not valuable merely because it promotes happiness; it has intrinsic value.

The first interpretation seems mistaken, since it is difficult to see how, practically speaking, there could be a set of intuitive principles that better protect autonomy than those that I argue for above. Of course it is possible to give citizens absolute protection against government spying, but it is a mistake to think that such protection would also promote citizens’ autonomy. A person’s autonomy is threatened not just by its own government intruding into its affairs, but also by meddling by other actors, such as private citizens, corporations, foreign governments, etc. that a citizen’s government tends to protect her against. The principles that I defend above are designed, in part, to protect citizens against harmful forms of intrusion, intrusions that would almost certainly violate a person’s autonomy, or impair her capacity to act autonomously in the future.

If I am right that from a practical perspective the principles that I defend above best promote autonomy, then even if the second interpretation of the autonomy objection is correct, it has no impact on the intuitive principles that one should endorse. If we accept a theory of value that gives autonomy a more prominent place than utilitarianism does, we will still endorse the principles that I argue for above. Hence, one might say, the principles above are *triple* justified: they are supported by widespread intuitions, utilitarianism, and moral theories that give autonomy a more privileged place than utilitarianism does.

The second remaining objection to classical utilitarianism is that it prescribes spying on innocents. Does this objection have any force against TLU? I do not think it does. I argued above that the two-level utilitarian would endorse a set of intuitive principles that gives strong protection to innocents. In particular, the two-level

utilitarian would accept a principle of discrimination that holds that it is never permissible to *target* innocents. True, this principle does not give *absolute* protection to innocents, but absolute protection for innocents, I argued, is not plausible. The reason I cited above is that absolute protection for innocents leaves too many benefits on the table. In more concrete terms, what this means is that an absolute protection for innocents would leave people more vulnerable to violent and otherwise serious harms.

It is further worth noting that the discrimination principle is not the only protection for innocents in the intuitive principles I argued for above. The minimization principle also requires that government agents take reasonable precautions to limit harms to innocents, by for example discarding information gleaned from spying not likely to be useful to securing just cause.

Let us now turn to the third objection to CU, that the wrongness of spying should not depend on whether it is discovered. This objection does not seem to be avoided by my move to a two-level theory. Consider the following case:

Crime Prevention: Imagine the police chief in your community initiates and manages a systematic spying program of those he suspects to be “threats to the community.” Threats are determined not by any objective criteria, but rather by the chief’s hunches. Those suspected to be threats are secretly monitored unremittingly. The spying results in countless arrests and the prevention of numerous serious crimes. The program saves lives and prevents many from being gravely harmed. But the program is never revealed to anyone outside the police force.

Many of us doubt that a case like Crime Prevention could ever occur in the real world. Such an extensive program of spying is exceedingly unlikely to remain secret for very

long. But let us assume that the details of the case are true. What does TLU say about the case? The costs in Crime Prevention are rather minimal and the benefits are considerable – so considerable, in fact, that it is not entirely unreasonable to think that spying was the best available alternative. So, TLU seems to be committed to concluding that the police chief acted rightly. Yet had the spying been discovered, TLU would likely reach a different verdict because the costs of the spying would likely increase exponentially. The police force’s reputation would be seriously damaged; a general paranoia would ensue; people would wonder whether they were watched in the past, whether they are now being watched, etc.

Some might take TLU’s analysis of Crime Prevention as an indication of a deeper problem with the theory. They might say that the wrongness of spying has nothing to do with whether it is or is likely to be revealed or not. The police force’s spying, they would maintain, was wrong *even when it remained a secret*.

Whether the objection fails or succeeds depends on a fine parsing of our moral intuitions, which raises the question of how trustworthy our intuitions are in cases like Crime Prevention. I want to argue that they are not trustworthy down to the level of specificity required for the objection to succeed. My argument does not deny that most of us, when we consider a case like Crime Prevention, think that something is not quite right. I certainly have a gut reaction that the police chief has not behaved in an exemplary fashion. The question is whether my gut reaction can be confidently interpreted as an indication that the police chief’s actions are wrong. I do not think that it can, since my gut reaction could be not a response to the wrongness of the case but instead a response to the bad decision procedure the chief employed or to his blameworthiness. Moral intuitions (at least *my* moral intuitions), I want to suggest, are blunt instruments. They do not permit us to parse out our responses to cases to a fine level of detail.

Consider first the possibility that my gut reaction is a response to the police chief employing a bad decision procedure. As I mentioned above, many people (myself included) doubt whether a case like Crime Prevention could really remain secret forever. In most cases, massive spying programs will be revealed, leading to prohibitive costs. The optimal set of prima facie principles for the police chief, in other words, likely prohibit such widespread spying programs. Hence, my gut reaction may be a response to a bad decision procedure, permitting me to say, “The chief got lucky this time and stumbled upon the right act. But he is unlikely to be so lucky in the future. Most of the time wide spread spying programs will be discovered, so the chief’s decision procedure was suboptimal.”

A second way my gut reaction to Crime Prevention could be interpreted is as a response to the chief’s *blameworthiness* for initiating the program. TLU, like many other consequentialist theories, is not committed to the standards of praise and blame being identical to those for right and wrong. In other words, just because the police chief acts rightly in this case, it does not follow that he should be praised or even that he is not blameworthy. Blaming the police chief may turn out to promote utility by, for example, causing him and other police officers to follow more optimal decision procedures in the future.

TLU likely tells us, then, that the chief’s action in Crime Prevention reveals a suboptimal decision procedure and that his action is blameworthy.¹³ So if TLU is true, we probably *should* have a gut reaction that something is amiss in Crime Prevention. But, again, our reaction would not necessarily indicate the wrongness of the chief’s spying. In order to determine whether cases like Crime Prevention strengthen or weaken the case for TLU, our moral intuitions would need to be much more precise

¹³On one very plausible version of TLU, these turn out to be identical because the standard of blame and praise just is whether the agent follows the optimal decision procedure.

than they are.

The final objection worth considering can be called the “extensive spy state” objection. This argument is a familiar one to those who have studied utilitarianism. Indeed, if one swaps the phrase “an extensive spy state” with “slavery,” one produces the very famous slavery objection to utilitarianism. Here is the objection in Hare’s (1979, 104) words, “It is often said that utilitarianism must be an objectionable creed because it could in certain circumstances condone or even commend slavery, given that circumstances can be envisaged in which utility would be maximized by preserving a slave-owning society and not abolishing slavery.” The slavery objection, as Hare (Ibid.) shows, fails, and I want to argue that the extensive spy state objection fails for more or less the same reasons.

To envision the circumstances in which a spy state would be utility maximizing, we need to conjure up a rather precarious society. It needs to be precarious enough that it is obvious that abandoning the spy state would be suboptimal. Otherwise, the two-level utilitarian could claim that we ought to fall back on the well-tried liberal principles that respect liberty, privacy, and autonomy. As Hare says in the case of slavery, “If it [the utility of the slave society] were not clearly greater, utilitarians could argue that, since all judgements of this sort are only probable, caution would require them to stick to a well-tried principle favouring liberty, the principle itself being justified on utilitarian grounds...and thus the example would cease to divide them from their opponents, and would become inapposite.” But we must also be careful not to imagine the society that is too dangerous, thereby making the utility of the extensive spy state too great, because as the dangers become more grave, many alternative theories to TLU might also endorse the extensive spy state, assuming such a state turns out to be the best means for staving off these grave dangers.

The following example, I believe, meets these requirements. Imagine that af-

ter 9/11, terrorist attacks by Islamic extremists continued relatively frequently and that these attacks were aimed not only at cities in the United States but also at cities in Canada. Los Angeles, Toronto, Chicago, Montreal, Dallas, Seattle, and Vancouver were all struck. In response to these attacks, America sticks strictly to its tradition of civil liberties, spying on citizens almost exclusively when warrants demonstrating probable cause can be obtained. Following these rules, America has some successes capturing and/or killing terrorists; it works toward addressing those things often thought to lead to terrorism – poverty, military installations abroad, historic grievances, etc. – but attacks continue at more or less the same frequency. Widespread fear cripples the economy. Anti-Islamic sentiments explode, leading to scores of hate crimes against innocents. Meanwhile, Canada implements an extensive program of government spying. The government monitors most personal communications. Government agents secretly observe public places for suspicious activities. Agencies develop dense networks of informants. For the most part, spying is performed without restrictions. The effect of Canada’s policy of unrestricted government spying is a considerable reduction in attacks in Canada. Canadian intelligence agencies catch and kill terrorists at a much higher rate than their American counterparts. Although the private lives of Canadians shrink appreciably in response to the policy; and politics within Canada becomes more conformist, these costs are outweighed by lives saved and fears allayed. Indeed, the success of the Canadian program of unrestricted spying is so apparent, Americans begin pressuring their government to abandon outmoded traditions and to emulate the Canadian policy.

Two-level utilitarians can follow two strategies for denying that the extensive spy state objection is sound. They can claim, by questioning the stipulated facts, that the Canadian policy is not utility maximizing. Or, they can quibble with the idea that an extensive spy state is always wrong. Both moves, I believe, bear fruit.

First, assume that the imagined facts of Extensive Spy State are beyond question. Do we have a strong conviction that an extensive spy state is wrong in the stipulated circumstances? I do not think we do. Many people accept that in extraordinary circumstances, governments should sometimes be granted wider and less restricted powers. In ancient Rome this belief was codified in the institution of the dictatorship. Today, most countries have procedures for declaring martial law. Because spying is typically thought of as one of the crucial powers for governments in dangerous times, it does not seem problematic to infer that most people accept that in extraordinary circumstances (like those in Extensive Spy State) governments should be provided wider and less restricted powers to spy.

But perhaps I am wrong and many ordinary people disagree with this intuition. Should this disagreement count as evidence that Canada's actions in Extensive Spy State were wrong? It should not. The reason is that our substantive moral intuitions lose their reliability in exceedingly unlikely cases. Many people probably reject the Canadian solution in Extensive Spy State *because* they have developed their convictions about state spying in a normal range of circumstances, circumstances in which there are no good reasons to cede nearly unlimited powers to the government to spy.¹⁴

From the perspective of the two-level utilitarian, that people's intuitions are adapted to respond to normal cases is a good thing. In most contemporary democracies, circumstances like those in Extensive Spy State do not obtain, so when citizens in those democracies contemplate the right policies for government spying, it is for the best that they have the conviction that unrestricted state spying should remain off the table, for if they didn't, they "might be tempted, whether through ignorance

¹⁴One could also argue that people's conviction in Extensive Spy State are driven more by haunting examples from literature and the movies that immediately come to mind than a careful reckoning of the facts in the imagined case. It is difficult to extricate cases like Extensive Spy State from more haunting cases, such as those in *1984* or in *Minority Report*, which people rightfully condemn.

or by self-interest, to condone” spying “in cases in which, though actually harmful, it could be colourably represented as being beneficial.” (Hare 1979, 115)

TLU is therefore committed to saying what seems to be a paradox: that Canada’s actions in Extensive Spy State are right but that it is for the best that most people have the conviction that Canada’s actions are wrong. But this is no real paradox. The principles that we use to guide our actions are adapted to the range of cases that we are likely to face, leading us to the right actions most of the time and particularly when our actions matter most. Creating principles to respond to extremely unlikely cases like Extensive Spy State would be disastrous for the production of utility, since it would lead us to reject perfectly good principles for the strong preponderance of cases that we are likely to face, or it would complicate our principles to the point which they are no longer practically useful.¹⁵ Hard cases make bad law.

But suppose I am wrong. Can the extensive spy state stand up as utility maximizing? I do not think it can. I made every attempt to make the case plausible, but I remain deeply skeptical that such a case could actually turn out to be utility maximizing. An assumption in the case, for example, is that the Canadian government, after it is granted unrestricted powers to spy, remains more or less the same. But this seems unlikely. Officials with unrestricted powers to spy would be strongly incentivized to turn their powers against their political opponents. Canada’s thriving democracy, normally faithfully carrying out the interests of its citizens, could quickly turn into a one party state, advancing the interests of only a tiny minority.

Second, it is dubious that the gains in policing effectiveness would jump considerably because governments are permitted to spy more widely, as the Extensive Spy

¹⁵As Hare (1981, 47) says, what we shouldn’t do is “call to mind the improbable cases that novelists, or philosophers with axes to grind, can dream up, and ask whether in those cases the outcome of inculcating the principle would be for the best.”

State case assumes. Existing decision procedures (as well as the principles I defend above) provide law enforcement officials with considerable leeway, especially when the stakes are high and serious harms can be prevented. It is thus reasonable to think either that as decision procedures become more permissive than those in the United States, the costs of spying would increase faster than the benefits, or alternatively, that there would be gains to marginally more permissive decision procedures for spying, but *unrestricted* spying still could not be justified.

Finally, it is difficult to accept that granting the state nearly unlimited powers to spy is the best way to thwart deadly terrorists because there are so many other less costly ways that the problem could be addressed. Tightening security at likely targets, pouring more resources into tracking and tracing terrorist suspects, sincerely addressing the grievances of terrorists – all of these tactics have promise, and none of them seem to come with the risks of giving the government nearly unlimited powers to spy. Hence the assumption that Canada’s actions in Extensive Spy State are optimal is not a plausible one.

6.5 Conclusion

In this chapter I derived the same five intuitive principles from two-level utilitarianism that I derived from widespread intuitions in chapter 3. I also argued, on utilitarian grounds, that these principles should be institutionalized, and I introduced two “ancillary principles” necessary for the five principles to be efficacious. In the remainder of the chapter I returned to some objections raised to classical utilitarianism in Chapter 4 and showed either how moving to the two-level view avoids the objections or why the objections are not persuasive in the first place. By now, I hope to have persuaded my reader of the plausibility of the two-level view, even if she

is not convinced that two-level utilitarianism is the best moral theory. Even more importantly, I hope to have persuaded her that the principles I argued for in chapter 3 and again in this chapter are good ones for government agents to follow.

Chapter 7

Principles for Foreign Spying

In previous chapters, I argued for principles to regulate domestic government spying. Many people may wonder to what extent, if at all, these principles extend to foreign spying. The principles that guide clandestine foreign agents, they may think, should be different from and likely more permissive than those that guide local cops. James Bond should be given more leeway to spy than his counterparts in the Scotland Yard. This intuition is certainly consistent with the way institutions are configured in the United States. In former CIA assistant general council John Radsan's words,

The hope for the U.S. intelligence community is that the beast can be contained. That is, we hope that the lawlessness will remain outside the U.S. jurisdiction. The CIA's black bag jobs, which it conducts in capitals around the world, are supposed to be off limits in our own capital. Outside the United States, the CIA prowls the alleys without a leash.

Inside the United States, the CIA is supposed to behave as a domesticated animal.(Radsan, 2007, 618)¹²

In this chapter, I want to suggest that this (arguably widespread) intuition is mistaken. The principles for domestic spying should not be altered drastically for the foreign context. If my thesis is correct, it has important implications for the conduct of spy agencies and the design of institutions in the United States and in many other countries. It suggests that spy agencies should not be permitted to “prowl without a leash” abroad, but that they should follow similar procedures for foreign spying to those they follow for domestic spying.

To make this argument I first need to clarify what I mean by “foreign” spying. Foreign spying could mean one of five things (See Table 7.1). It could mean spying on citizens abroad, spying on foreign individuals abroad, spying on foreign individuals at home, spying on foreign states abroad, or spying on foreign states at home. Plausibly each of these types of spying requires slightly different treatment. Each may even require a separate set of normative principles.

Rather than explore each of these categories in depth, my plan is to scrutinize the two categories that make up the lion’s share of foreign spying: spying on foreign individuals abroad, and spying on foreign states abroad. States have aggressively spied on one another since the beginning of recorded history, and continue to expend enormous resources to covertly gather intelligence about one another. Further since

¹A black bag job is a covert and illegal operation to secure information. They usually require breaking and entering, hacking, safe cracking, etc.

²Sepper (2010) also notes, “Where legal restraints [on spying] do exist, they impose limits almost exclusively on domestic activity. These include requirements that foreign and domestic intelligence be separated, residents’ and citizens’ information not be intercepted, and, as applies to intelligence networks, residents and citizens data be shared only in accordance with domestic data protections. Even in democracies, however, there are generally no statutory permissions or limitations on intelligence work outside national borders or intelligence relations to foreign counterparts.”

Table 7.1: The Targets of Government Spying

	Citizens	Foreign In- dividuals	Foreign States
Located Do- mestically	Chs. 3-6		
Located Abroad		Ch. 7	Ch. 7

the end of the Cold War, the non-state threats facing most developed states have become more significant and therefore more prominent. A consequence of this shift is that states are doing more spying than ever on non-state actors, individuals and groups inside foreign states.³

The position that I reach by the conclusion of the chapter is *not* that principles for domestic spying should simply be extended to foreign spying. Some important alterations to domestic principles are in order both in the case of spying on foreign individuals abroad and in the case of spying on foreign states. But in both cases, principles for foreign spying should be institutionalized, and government agents should follow a set of principles that look and feel markedly similar to those they should follow in the domestic case.

The method that I use to develop the principles in this chapter is primarily consequentialist, since most of the chapter's space is spent comparing the likely harms and benefits of foreign spying to those in the domestic case. I do not ignore arguments that we should be partial to our compatriots, however. On the contrary, I consider both instrumental and intrinsic reasons for partiality toward compatriots. I conclude that while there are instrumental reasons to be more hesitant to spy on our compatriots, intrinsic arguments for partiality, even if they are successful, are unlikely to weaken our duties not to spy on foreigners.

³In the notes, I try to indicate how the considerations I examine might affect the remaining three kinds of foreign spying. But more work is required to fully develop principles for these categories.

The argument in the chapter proceeds as follows. The next section takes up the case of government spying on foreign individuals abroad. I first argue that, just as in the domestic case, concerns about abuse support institutionalizing principles of government spying. Next, I argue that the domestic principles for spying should be altered in three respects for spying on foreign individuals. Since states have a comparative advantage spying on their own citizens, there should be a presumption against spying on foreign individuals abroad; since *ceteris paribus* the potential harms of spying on foreign individuals are less severe, the bar for just cause should be lowered; and since spying is likely to have weaker behavior conditioning effects, the principle of discrimination should be less restrictive.

Section two examines spying on foreign states. I argue, again, that the case for institutionalizing principles is sound. Further, I argue, that domestic principles should be altered in three ways for spying on foreign states. The bar for just cause should be lowered (even further than in the case of spying on foreign individuals) and it should be altered to take into account situations where other states do not uphold their obligations not to spy; the principles should be adapted to take into account the real risk of retaliation; and the principle of discrimination should be jettisoned, since the conditioning effects of spying are minimal when the target of spying is a state. In the final substantive section, I say a few words about how principles for foreign spying can be institutionalized. In chapters that follow, my focus is almost exclusively on domestic institutions. The reason, I explain in this section, stems from the very difficult practical problem with creating international institutions to regulate spying. The final section is a short conclusion.

7.1 Spying on Foreigners

The first task is to determine whether the argument that I made in chapter four – that principles for domestic government spying should be institutionalized – can be extended to the case of government spying on non-citizens residing abroad (“foreigners” for short). The case for institutionalization in the domestic context, recall, was that since law enforcement officials and intelligence agents have much to gain professionally from spying, and their handlers have much to gain politically from spying, and since citizens have much to lose from government agents abusing their powers to spy, government agents should not be the only arbiters of how principles of just government spying apply.

In order to determine whether an analogous case can be made for institutionalizing principles of government spying on foreigners, three questions must be answered. Do law enforcement officials and intelligence agents have professional incentives favoring spying on foreigners? Similarly, do politicians have incentives favoring spying on foreigners? Finally, do foreigners have much to lose when foreign states abuse their powers to spy? If it turned out that government spies and their handlers had few incentives to abuse their powers to spy on foreigners, or that foreigners have little to fear from foreign governments abusing their powers to spy, then the case for institutionalization would be on shaky ground.

The incentives of intelligence agents often strongly favor spying on foreigners, even when this spying cannot be justified. The principal professional aim of intelligence agents is to identify and prevent threats to the state. From the perspective of these agents, more information is almost always thought helpful toward achieving their aims. If intelligence agents have a hunch that a foreigner plots or schemes against their state’s interests, they typically wish to follow this hunch without legal encum-

branches.

Since intelligence agents have strong incentives to spy, they tend to make errors in the direction of over- rather than under-collection. Consider that one of the most common problems cited in foreign intelligence gathering is the glut of information that intelligence agencies have – so much, in fact, that mountains of data routinely go unanalyzed. (Lowenthal, 2009) This surfeit of information might be caused by the fact that under current institutions intelligence agencies are typically free to follow their hunches when they spy on foreigners. If judgments about whether to covertly collect information were made by what Justice Jackson called a “neutral and detached magistrate,” then the problem of over-collection might be mitigated.

The same argument applies to law enforcement officials. If officers reach an impasse on a case but they have a hunch that a suspect’s relatives abroad might have information about her whereabouts or about the crime, then the officers have strong professional incentives to snoop abroad.

Law enforcement officials and intelligence agents thus have strong incentives to spy on foreigners. Government officials who “handle,” “run,” or “manage” spies also sometimes have incentives to collect more information than is justified. Their responsibilities and perhaps more importantly their electoral incentives are to promote and protect the welfare *of their own citizens*. So it would be no surprise if the executive or her administration spied overzealously on foreigners to secure or promote the interests of her constituents. This asymmetry is an important and under-appreciated aspect of foreign intelligence. State officials can often generate benefits for their constituents, for example heightened security, while ensuring that the costs of generating these benefits are disproportionately borne by non-constituents. They can benefit themselves and their constituents, that is, while externalizing the costs of producing these benefits to others.

So far the argument for institutionalization in the domestic case seems to extend to the case of spying on foreigners. In both cases spies and their handlers have professional, prudential, or political incentives to spy that often diverge from their moral responsibilities. Hence they cannot be trusted to be impartial arbiters of how the principles of government spying apply in particular cases. However, it is possible that abusive spying on foreigners leads only to trivial harms, thereby undercutting the case for institutionalization.

Since the potential harms of spying on foreigners are similar to those in the domestic case, rather than developing the harms from scratch, my strategy below is to begin with the domestic harms that I developed in chapter 5 and then point out important ways the potential harms to foreigners differ.⁴

I want to argue that on one hand there are a variety of considerations that suggest that the potential harms from domestic spying are greater compared to spying on foreigners. For example, states have greater capacities to control and otherwise harm their own citizens than do foreign states. States tend to have special relationships with their own citizens not possessed with foreigners. And the trust between citizens and their own states seems more valuable than the trust between citizens and foreign states. But, on the other hand, the potential harms from spying on foreigners are still significant, and some potential harms, for instance the harms of retaliation, may be *more* significant compared to the domestic case. Let me attempt to unpack and defend these claims by comparing some of the potential harms when the target of spying is a citizen to when she is a foreigner.

Spying, as I suggested in Chapter 5, can be used to gather sensitive information that is then utilized to embarrass, exploit, or shame the target or others, and in po-

⁴I summarize the consequences to spying in Table 5.1 and Figures 5.1-5.2.

litical contexts, sensitive information can be used to fix a jury, to destroy or discredit political opponents, etc. How do these harms differ when the target of the spying is a citizen or a foreigner?

After states have collected sensitive information about their own citizens, typically they can exploit this information better than could foreign governments. They have more levers to pull to produce downstream costs and benefits, since they tend to have (among other things) access to more information about their citizens than do foreign states. Whatever information is collected from states spying on their own citizens will tend to be added to a more complete picture, and other things being equal, the more complete a dossier a state has compiled on an individual, the more effectively and efficiently it can harm that individual. Beyond informational advantages, states also, because of proximity and jurisdiction, tend to have a greater capacity to harm or control their own citizens than foreigners. States, for example, typically cannot arrest, prosecute, and imprison foreigners in the same ways they can their own citizens.

Still, it is important not to exaggerate these differences. States sometimes secretly gather information abroad that is useful for prosecuting domestic cases. In addition, states – and powerful states in particular – sometimes arrest, prosecute, and imprison foreigners. The United States, for example, has apprehended people all over the world, imprisoned them, labeled them as “enemy combatants,” and tried some of them in military tribunals, while detaining others indefinitely. Finally, although states often do not have the capacity to prosecute foreigners, they often *do* have the capacity to harass, coerce, manipulate, and otherwise harm them.

A second potential harm from spying is the chilling effect that spying can have on people’s conduct. People condition their behavior when they suspect they are under observation, especially when they suspect that knowledge of their behavior will lead

to sanctions or rewards. There are reasons to think that spying will have a more significant chilling effect when citizens suspect that it is their own government doing the spying. As we have seen, states typically have, and are usually believed to have, more power over their citizens, and state officials have, and are usually believed to have, more incentives to interfere in domestic matters. The chilling effect is magnified when the observer is believed to be capable of and interested in intruding into more aspects of the target's life. Many citizens also likely believe that their own states will protect them, to some extent, against interference by foreign governments. When their own government interferes in their lives, in contrast, they tend to think they have little recourse.

The differences, again, should not be exaggerated. Government spying can still have material effects on the conduct of foreigners. Imagine a group protesting an American military base in Germany, Japan, or Iraq. Given America's capacity to project its power around the world and its interests in having bases strategically located around the world, it would be no surprise if some dissenters abandoned their protests when they learned they were under covert surveillance by the American government.

A third potential harm of spying is that it can evoke strong negative emotions, such as anxiety, paranoia, or even humiliation. In extreme cases, when spying is done exclusively or disproportionately on disadvantaged groups, it can go beyond emotional harms: it can undermine people's dignity or self-respect.

In the domestic context, there is a reason to think harmful emotional responses are likely to be greater. The reason relates to the special relationship that often obtains between states and their citizens. States are meant (many people think) to protect the rights of their citizens, and they are meant to create and sustain the conditions of mutual recognition and respect. Hence, when a state violates one of its citizen's

rights or it undermines her dignity or respect, it may occasion a deeper harm than if the same violations were perpetrated by a foreign state. As Mill (1979) says, few “hurts which human beings can sustain are greater, and none wound more, than when that on which they habitually and with full assurance relied fails them in the hour of need.” Just as to be harmed by one’s own brother or sister is often worse than to be harmed by a stranger, it is often worse to be harmed by one’s own state than a foreign state.

Even though, other things being equal, the emotional responses to domestic spying may be stronger than the emotional responses to foreign spying, foreign spying can still result in profound feelings of paranoia, humiliation, etc., especially when the spy is seen as a brash hegemon, an occupier, or a colonizer.

Yet another potential harm of spying when it is suspected or discovered is the spy’s loss of trustworthiness. Spying often breaks entrenched norms of appropriate observation, prompting some to discount the degree to which they trust the spy (and those by whom she is employed). The potential harms of lost trust may be worse in the domestic context, since it seems plausible that citizens must trust their governments to some degree for institutions to function in a relatively orderly way. Fragile are the institutions maintained by coercion alone. (Hart, 1997, Ch. 4) But less rides on the trust that citizens have in foreign governments. Maybe some trust is necessary to maintain amicable relations among states, but the fabric of society is not undermined by citizens losing their trust in foreign states.

Finally, spying when it is suspected or discovered can provoke retaliation against the spy, or against the agent for whom she spies. Domestically, government spying is unlikely to be met with significant retaliation. The state is more powerful than its citizens and their voluntary associations by an order of magnitude. Retaliation is thus irrational, even stupid. But spying on foreigners, when it is suspected or discovered,

could trigger retaliation by foreign states, which could end in extensive harms.

To summarize, there are good reasons to think that the potential harms of spying on citizens are more severe than the potential harms of spying on foreigners (with the exception of the potential harms of retaliation). However, just because spying on foreigners is *relatively* harmless does not mean it is absolutely harmless. On the contrary, the discussion above suggests the potential harms of foreign spying are sizable.

Hence the institutionalization argument extends from the domestic context to spying on foreigners. Spies and their handlers cannot be trusted to impartially apply principles of government spying on foreigners, and a good deal is at stake with the right application of these principles, since the potential harms from abusive spying on foreigners can be considerable. Thus the principles of government spying on foreigners ought to be institutionalized.⁵

An objector might suggest that I am underestimating the differences between foreign and domestic harms. The potential harms of domestic spying, she might claim, are of an entirely different magnitude, since there are certain goods that are only attainable within political communities. Governments thus have *special* obligations not to spy on their own citizens.

More generally, it is sometimes argued that:

⁵We can attempt to draw out the implications of the above considerations for the other categories of foreign spying mentioned in the introduction by making a few observations. The harms of state spying on foreigners within the state's borders may be similar to the harms of spying on foreigners abroad, if the foreigner's stay is temporary. But the harms of spying on foreigners who aspire to join the state's political community may be more like the harms to citizens at home. Someone who aspires to membership in a political community may, for example, experience harms of recognition in roughly the same ways as do existing members. Similar inferences can be drawn for citizens abroad. The harms to those who are only temporarily residing abroad are probably similar to the harms to citizens at home. But the harms to those who are seeking citizenship abroad – and particularly to those who take steps to renounce their citizenship – are probably more like the harms to foreigners abroad.

1. Certain profoundly (or even intrinsically) valuable goods, such as social justice, deliberative democracy, or equal recognition are attainable only within political communities.
2. These goods are constituted (or promoted) by special obligations, in this case meaning obligations we have only to members of our political communities.
3. Hence, given our lives will go better if we attain these goods, we have special obligations to those in our political communities.⁶

What is unique about this kind of argument for special obligations is that it is consistent with both impartiality during moral deliberation, and consequentialist ethical theories. (Driver, 2005) Theorists who make this kind of argument are not claiming that we ought to give more consideration to the concerns of, or more weight to the interests of, our fellow citizens. The argument is rather about attaining or promoting

⁶For instance, Mason (1997) argues that our special obligations to our compatriots derive from the good of recognition. Citizenship, on Mason's view, "has intrinsic value because in virtue of being a citizen a person is a member of a collective body in which they enjoy equal status with other members and are thereby provided with recognition." (442) Further, "[p]art of what it is to be a citizen is to incur special obligations...In particular, citizens have an obligation to each other to participate fully in public life and an obligation to give priority to the needs of fellow citizens." (Ibid) Finally, the benefits of citizenship cannot be obtained without the members of states taking on special obligations to one another. Citizens cannot enjoy equal status and recognition without their fellow citizens fulfilling their special obligations. Similarly, David Miller (1988; 1997; 2004) argues we have special duties to fellow nationals on account of the ethical significance of nationality. He claims that groups generate special duties when membership in them is of intrinsic value, when they're not "inherently" unjust, and when the duties are essential to the group's value. He further argues that nations can (although they do not always) fit these three conditions. Relationships between fellow nationals are intrinsically valuable, says Miller, because people's lives go better when they are members of nations, and although some of the benefits of national membership are instrumental, fellow nationals must first "believe that their association is valuable for its own sake, and be committed to preserving it over time, in order to be able to reap the other benefits that national solidarity brings with it." (2004, 67) There is no reason to think that nations are inherently unjust (although some nations no doubt are patently unjust). Finally, the special duties to fellow nationals are integral to nationhood, Miller thinks, because nations function to "underpin political values like social justice or deliberative democracy" and they could not do this without presupposing "that nations are ethical communities whose members have special responsibilities both to support one another and to preserve their community." (Ibid, 69).

certain valuable goods. Its structure is teleological.

Could an argument of this sort generate special obligations not to spy on our compatriots? Let us consider, first, the argument using deliberative democracy. Assume that a deliberative democracy is not possible, or that it functions suboptimally when it crosses political boundaries - perhaps because people struggle to identify or to empathize with people who are too much unlike them. (Miller, 2009) Is it plausible that deliberative democracy requires certain prohibitions against government spying on members of the demos? It is. Deliberative democratic procedures demand a bundle of rights, including “freedom of (political) speech, freedom of the press, the right to form and join political parties, freedom of assembly and the right to present grievances, and perhaps even certain anti-discrimination measures that protect the social and political status of ‘entrenched minorities.’” (Freeman, 2000, 381) Many, if not all of these rights depend on protections against government spying.

Take the right to speak freely. As I argue in chapter 5, if people suspect they are being spied on by their government when they speak out publicly against it, they may censor their speech. Spying can thus have a chilling effect on public speech. But it is not only public speech that could be affected by spying. People need protected spaces outside of the gaze of others to try out political or ethical positions.

Protections against government spying are thus crucial for the right to speak freely. A similar case could be made for most of the other rights supporting deliberative democratic procedures. But it is a mistake to think that these arguments establish *special* obligations, that is obligations we owe *only* to compatriots. A well functioning deliberative democratic procedure relies on citizens having protections not only against their own government’s spying, but also against spying by foreign governments. Citizens in the forum or in other deliberative spaces can not be expected to deliberate freely when they suspect foreign governments are monitoring

them, especially when they believe there may be repercussions for the claims they make. Nor can deliberation be called free when foreign governments use spying to manipulate citizens or citizen groups. Hence securing deliberative democracy demands protecting all democratic citizens against government spying, regardless of whether the government is their own. If we have a special duty not to spy, it is a duty owed to all democratic citizens, not just to the citizens of our own state.⁷

Similar arguments can be made for social justice and recognition. Social justice depends on citizens not being bullied or silenced by their government, but it also relies on citizens having these protections from foreign governments. Recognition entails that one group of citizens is not disproportionately or unfairly monitored, but it also means that citizens are not unfairly monitored by foreign governments. In large, powerful states it is easy to think that one's own state is always the chief threat to liberty, justice, or self-respect. But in smaller, less powerful states the principal threat to these goods is often foreign meddling.

This type of instrumental argument for special obligations not to spy on citizens does not succeed, then. Of course there are other non-instrumental arguments for special obligations, and I will put off considering these arguments until later. Here my purpose for taking up instrumental arguments for special obligations was to establish that my comparison of the harms of foreign and domestic spying above is not far off the mark.

Now that the case has been made for institutionalizing spying on foreigners, let us turn to determining the principles to regulate government spying on foreigners.

⁷In fact the duty probably cannot even be limited to democratic citizens. Because free speech is so important for establishing democracies, it seems the duty may be owed more broadly.

Since states often have a greater capacity to monitor their own citizens than do foreign governments, one might think that states should be assigned the responsibility for monitoring their own citizens. At the extreme, one might follow this argument to an outright prohibition on foreign spying. States all monitor their own citizens, hence there is no role for foreign spying. Such a strong argument, I want to suggest, cannot be supported, but this “assignment” argument nevertheless can yield some restrictions on foreign spying.

The assignment argument has been developed in a series of works by Robert Goodin (Goodin, 1985; Pettit and Goodin, 1986; Goodin, 1988). It begins from the premise that we all have general obligations to protect and promote the welfare of one another. In order to best fulfill these general obligations, it is often the case that we assign particular agents special responsibilities. We assign parents the responsibility of ensuring their children attend school, for example, and we assign police officers the responsibilities of pursuing and arresting suspected criminals. If we all tried to do these tasks simultaneously, fewer children would make it to school, and fewer suspected criminals would be apprehended. Assignment is thus an efficient way to fulfill our general obligations.

Goodin (1988, 681) points to a number of reasons for assigning responsibilities to particular agents: specialization, informational limitations, psychological predispositions, and institutional configurations. But he stresses that special responsibilities derive “wholly from the fact that they were appointed, and not at all from any facts about why they were appointed.” (680) If it turns out that those appointed are incapable, then “of course it is perfectly proper for us to retract their commissions and appoint others in their places.” (Ibid.) But it is not a good practice to continually reconsider and remake appointments. Often suboptimal appointments are vastly superior to no appointments at all.

A variant of the assignment model already operates for spying within the nation state. We have obligations to protect others from serious harm, and spying is often one way to prevent these harms. But we do not all share equally the responsibility of monitoring potential threats. Instead we assign this responsibility to the local police, the FBI and other law enforcement and intelligence agencies. Assignment permits these agencies to develop expertise and centralize information, making them vastly more efficient spies than myriad individuals sharing the responsibility.

At the international level, the assignment argument begins from an observation I have already made: that states are better situated than foreign states or non-state organizations to spy effectively and efficiently on their own citizens. They have a comparative advantage spying on their own citizens. As I mentioned above, proximity permits states to more easily peer into the lives of their own citizens. So too does their better access to enormous quantities of information collected and stored about citizens. State officials also typically have knowledge of and connections to their fellow citizens they don't have vis-à-vis foreigners. Further, compared to non-state agents, states tend to have more resources for spying and considerably more power to act on the information that they collect from spying.

A second consideration is the bond that is often shared between fellow nationals. People are less likely to abuse their power when the abuses affect someone they consider to be one of their own. Abuses affecting only members of out groups are more likely to be discounted or dismissed. Hence, assigning states the responsibility of monitoring their own citizens may lead to fewer abuses than assigning this responsibility to other states or supra-national organizations. Although there may be sub-state groups that are even more fundamental to people's identities than are states, such groups rarely have the organizational, informational, or monetary capabilities to spy effectively. If some do, then perhaps assignment should be done at the

sub-state level. For now, however, I assume that the assignment argument points strongly toward states, which is likely accurate in the majority of cases.

The assignment argument has merit, then. But it does not support an outright prohibition against foreign spying, since not all states have the capacity to monitor their citizens or can be trusted to do so. Some states, for example, lack sophisticated intelligence services or cutting-edge spy technology. Other states assist, tolerate, or collaborate with the very agents they should be invigilating. Hence assignment does not imply prohibitions against spying on “troubled” states. There also may be cases in which a state has the requisite capabilities to monitor its citizens, but unjustifiably prioritizes dangers to foreign countries below domestic threats. The resources it directs toward a threat, in other words, are not commensurate with the threat’s magnitude.

A critic might object that the assignment argument supports no prohibition at all: the exceptions overwhelm the rule. Surely, she might worry, an agency such as the CIA could sincerely claim that its capacity to spy around the globe is greater than that of other foreign intelligence services; that many states cannot be trusted to give threats to American interests priority; and that many of the severe threats they monitor have too few resources committed to them, even with multiple intelligence agencies on their trail. Hence, she might argue, the assignment argument does not prohibit the CIA from engaging in *any* kind of foreign spying.

The objection helpfully highlights that the assignment argument likely entails that foreign spying will often be permitted, especially by well-resourced and well-intentioned intelligence agencies. But the objection is, ultimately, unsuccessful. Monitoring most citizens probably does not require the manpower or the technology possessed by the CIA. Many state intelligence agencies, particularly those in developed countries, have the resources to monitor their own citizens. Moreover, they likely do

more with less, given their connections on the ground. Finally, most of these governments are strongly motivated to extinguish threats by citizens and citizen groups affecting foreign states, not because they are altruistic, but because the threats that face foreign states (terrorism, organized crime, drug trafficking) often profoundly affect domestic interests as well.⁸

So the assignment argument favors a *prima facie* prohibition on spying on foreigners and this prohibition can be overridden when states lack the capacity to monitor their own citizens, they cannot be trusted to monitor their own citizens, or they cannot be expected to appropriately prioritize foreign threats. In rare cases, there may be threats so considerable that countering them requires the resources of many states.

The assignment argument thus suggests that certain kinds of permissible domestic spying may not be permissible on foreigners. My earlier consideration of the potential harms of spying on foreigners led mostly in the opposite direction. Other things being equal, the potential harms of domestic spying are more severe than the harms of spying on foreigners.

How do these cross-cutting conclusions affect the principles for domestic spying developed in previous chapters? I want to argue that domestic principles should be modified in three ways, one major and two minor. The major change involves the addition of a sixth principle: a principle of comparative advantage. This principle is implied directly by the assignment argument. It stipulates that it is not permissible

⁸We can work out the implications of the assignment argument to other categories of foreign spying with a few observations. Proximity favors assigning states the responsibility for spying on people within their own borders, whether they are citizens or foreigners. But, because of the information states tend to have access to about their own citizens, a case could be made for assigning the state the responsibility for spying on its citizens, regardless of where they are located. A reasonable conclusion for foreigners located domestically and citizens located abroad, then, may be dual assignment. For instance, perhaps both the United Kingdom and the United States should have the responsibility for spying on British citizens residing in the United States and for American Citizens residing in the U.K.

for governments to spy on foreigners unless one of the exceptional cases enumerated above (e.g. that a foreign state does not have the capacity to spy on its own citizens) obtains. The minor changes are to the principles of just cause and discrimination.

Since the harms of domestic spying will likely be more severe than the harms of foreign spying, it is reasonable to think that more just causes should be included on the list I developed in Chapter 6. The justificatory bar for each of the three types of spying identified in that chapter, in other words, should be shifted down. How far they should be shifted down is a difficult empirical question, and answering it likely requires data far more comprehensive than what is currently available. But my intuition is that the shifting should be relatively minimal, since the potential harms to foreign individuals from government spying remain considerable.

The second minor change is to the principle of discrimination. The principle of discrimination (I argued in Chapter 6) helps set agents' expectations about where and when they may be permissibly spied on. Citizens not connected to serious harms, under the principle of discrimination, know they are in very little danger of being spied on, while those breaking or scheming to break the law know they risk being secretly monitored by the government. A discrimination principle, thus, deters harms, while protecting legal and beneficial activities.

The principle could have the same effects at the international level, although one might expect that the magnitude of the deterrent decreases, given that states tend to have less capacity to interfere with foreigners. More importantly, however, since foreign governments are especially unlikely to interfere with foreigners when they are engaged in harmless or beneficial conduct, innocents probably require less extensive protection from spying by foreign governments.

Precisely where these considerations leads is unclear, but I want to suggest that it is plausible that government agents should be permitted to engage in foreign spying

for the broader purpose of threat identification, instead of just threat prevention, especially when the threats are of the most extreme nature, such as terrorism and nuclear deterrence.⁹ What separates threat identification from threat prevention is the epistemic or evidentiary hurdle that law enforcement officials or intelligence agents must jump over in order to permissibly spy. In the case of threat prevention, government agents must show that “the targets of their spying can reasonably be believed to be engaged in or assisting the harm that the government agents aim to prevent in securing the just cause.” In the case of threat identification, in contrast, the epistemic hurdle is lower. Government agents need only show that “it is reasonable to believe that the targets of their spying *could* be engaged or assisting the harm that the government agents aim to prevent in securing the just cause.”

The benefits following from such a change could obviously be large, to the extent that the greater leeway for intelligence agents leads to preventing catastrophic harms. But expanding the principle of discrimination in this way risks miring government agents in costly or potentially abusive fishing expeditions, signaling to foreigners that a country’s spying is done more or less indiscriminately, and even worse, encouraging government agents to illegitimately discriminate (consciously or not) against certain groups of foreigners in their threat identification procedures. For reasons discussed above, these risks could be less severe in the foreign context, but the risks, even in the foreign context, are nontrivial.

My intuition is that these risks can be mitigated by a discrimination principle that permits government agents to engage in threat identification *only* for a limited set of grave threats, and that compels agents to justify the procedures by which they

⁹A case could be made that for the most severe threats, discrimination should be modified even for the domestic case. See e.g. Posner (2006, 2008) and Solove (2008). But the cost/benefit analysis above suggests the case for altering discrimination in the foreign case is on stronger ground.

identify threats. But my intuition, admittedly, is based on a small and biased sample. Very little is known about threat identification programs, such as the NSA's colossal data mining operations, let alone how successful they have been. Many of these programs are not public knowledge, and there is no way of knowing whether the public programs are indicative of the set of all threat identification programs. (Bamford, 2002; Mayer, 2011)

I now want to address two worries that one might have about my analysis thus far. The first worry is that although I have carefully contrasted the harms of foreign and domestic spying, I have not done an equally thorough comparison of benefits. Perhaps the potential benefits of spying vary significantly when the target is a citizen or a foreigner. I think this is unlikely. Any threat by an individual or non-state group to a state and its citizens that can be imagined abroad can also be imagined domestically, and vice versa. For every Osama bin Laden, we can imagine a Timothy McVeigh. In the long-term, most states will face existential threats from inside and outside their borders. Sometimes domestic threats will be more significant and more exigent than foreign threats and other times it will be the reverse.

The second worry concerns special obligations to compatriots. My arguments thus far have relied exclusively on the harms and benefits of spying on foreigners. I have not ignored the possibility that we may have reasons to give our compatriots special consideration, but the only argument that I have considered for partiality toward compatriots was an instrumental one. Yet some theorists think there is a non-instrumental or "intrinsic" case to be made for partiality toward compatriots. Whereas instrumental reasons for partiality toward compatriots point to how rela-

tionships between compatriots or relationships between compatriots and their states or nations promote or partially constitute well-being, intrinsic reasons for partiality toward compatriots point only to these relationships themselves as sources of moral reasons.

If these theorists are right and there are intrinsic reasons for partiality toward compatriots then the principles I developed above are potentially mistaken, since they do not reflect these reasons. I want to suggest, however, that this worry is misplaced: even if we accept intrinsic reasons for partiality toward compatriots, we should not alter the principles for spying on foreigners that I argued for above, since reasonable partiality cannot justify weakening our negative duties not to harm others.

What is the intrinsic case for partiality toward compatriots? Jeff McMahan (1997, 129) points out that compatriots tend to cooperate together in a number of political and non-political shared projects, including “sustaining and continuously re-creating their culture and way of life as well as transmitting the cultural heritage to their descendants.” When a person benefits from the contributions that her compatriots make to these cooperative endeavors, McMahan argues, she “acquires duties of fair play to reciprocate.” (Ibid) Hence one source of special duties among compatriots “overlaps with the theory of political obligation.” (Ibid) Similarly, McMahan argues that since “one is deeply indebted to one’s nation and its culture,” given that it provides language, the stories by which we understand ourselves and our relations to others, and the basic “social infrastructure” that makes a good life possible, “the nation itself, as a transhistorical entity, is one’s benefactor, and there are duties one owes to it in consequence.” (130) So, according to McMahan, the intrinsic case for partiality toward compatriots rests on duties of reciprocity toward one’s compatriots and duties of gratitude toward one’s nation.

Thomas Hurka (1997) makes a different intrinsic case for partiality toward com-

patriots. He notes that nationalists often try to assimilate for their own purposes the intrinsic case for partiality toward family members, which many theorists take to be on strong intuitive ground. The problem many theorists point out, however, is that nations are not like families. Whereas we interact constantly with members of our family, many of our compatriots we have never even met. Although Hurka concedes that families and nations are dissimilar with respect to frequency of interaction, he argues that they are strongly analogous in a more important way.

For Hurka the intrinsic basis of partiality is shared history. Hence the crucial fact about the relationships within families and nations are the histories that members share. In his words,

Some activities and states of people, most notably their doing good or suffering evil, call for a positive, caring, or associative response. Others, such as doing evil, call for a negative or dissociative response. Partiality between people is appropriate when they have shared in the past in the first kind of activity or state. (152)

Hence when compatriots do good or suffer evil together, these shared experiences serve as the building blocks for a justified partiality.

Although I am skeptical that either McMahan's or Hurka's argument succeeds, let us assume that one of them is right and thus that there is a persuasive intrinsic case to be made for partiality toward compatriots. Intuitively, partiality toward compatriots would still be limited, and both Hurka and McMahan are careful to acknowledge this point. In Hurka's words, "It may be that any morally acceptable national partiality must be constrained by respect for the basic rights of all persons, both within one's nation and outside it." (155) Similarly, McMahan says "... there are obviously limits to the degree of priority that one is permitted to give even to one's closest family

members.” (132)

Partiality toward compatriots can be limited in one of two ways: the *extent* of partiality could be limited, or the *kind* of partiality could be limited. I want to focus on the latter sort of limitation. McMahan argues that the most intuitively plausible kind of partiality occurs in cases in which benefits are being distributed. It may be permissible, for example, for a community to use its tax dollars to build public schools for its children, even when this money could potentially do more good if it were directed to famine relief abroad. Less intuitively plausible but still sometimes reasonable is partiality in cases involving the prevention of harms. If I am in the position to save the life of only one soldier, I may be permitted to save my fellow compatriot, rather than a soldier from an allied country. It is not clear, however, that I would be permitted to save only one of my fellow soldiers if doing so meant letting two soldiers from an allied country die. Finally, McMahan argues that partiality in “cases involving the causation of harm and, in particular, cases in which an act causes a harm as a means of preventing a different harm or providing a benefit” is the least intuitively plausible. It does not seem reasonable, for example, for an American to kill two Germans or two Saudis in order to save just one American life. Nor does it seem, more generally, like we are any more justified in stealing from, lying to, or spying on someone merely because she is not our compatriot.

McMahan’s logic suggests that partiality toward compatriots should play a minimal if non-existent role in altering our negative duties toward foreigners. Our reasons not to harm others or impose risks on them do not require adjustment given the relationship in which we stand to those others. Thomas Pogge (2002) argues that this universal duty to avoid wrongfully harming others is the most defensible form of cosmopolitanism. He says, “. . . the stringency of our most important negative duties does not vary with the presence or absence of compatriotism. You do not have more

moral reason not to murder a compatriot than you have not to murder a foreigner. And you do not moderate your condemnation of a rapist when you learn that his victim was not his compatriot.” (87)

I have argued that the potential harms that follow from spying may be less considerable when the target of the spying is a foreigner compared to a citizen, and this difference leads us to endorse slightly different principles for spying on citizens and foreigners. I did not consider the possibility that we have less reason to worry about causing an identical harm to a foreigner than to a citizen. If McMahan and Pogge are right as I think they are, then my neglect of this consideration is not a problem. Our duty not to spy on others derives from the potential harm we may cause these others. These harms may vary in strength or frequency between citizens and foreigners leading us to justifiably treat citizens and foreigners differently. But it is not plausible to think that we should count identical harms differently, simply because in one case the harmed is our compatriot and in another case she is not.

7.2 Spying on Foreign States

Although spying on foreigners has grown considerably recently with the rise of non-state threats, spying on foreign states still makes up the lion’s share of foreign spying. So let us turn our attention to spying on foreign states.

Similar to the previous section, in this section I first consider whether the principles of government spying on foreign states should be institutionalized, and then consider how domestic principles should be altered for spying on foreign states. My focus is again on harms. I defend this focus at the end of the section.

There are a couple of respects in which the argument for institutionalization seems weaker in the case of spying on foreign states than in the case of spying on domestic

citizens, but since the argument in the latter case is strongly overdetermined, this relative weakness is not a decisive concern. I will not retrace the arguments regarding the incentives of intelligence agents and law enforcement officials, since I think they hold just as well when the target is a foreign state as they do for foreign individuals. I also think the asymmetry argument that I made above – that political officials have incentives to offload costs on foreigners to generate benefits for their constituents – apply when the target of spying is a foreign state. But let me say a few more words about the incentives of the political officials who handle spies.

In most modern democracies the executive typically directs spying on foreign states. If we assume that the executive's primary aim is to secure reelection, then she has incentives to abuse her power to spy. First, the executive may have constituents, especially large corporations, who would like to know the secrets of foreign states, and who would be willing to reward the executive handsomely with campaign contributions if she provided this information. State spying for domestic corporations is hardly an unheard of practice, so it would not be surprising if an executive used it for electoral advantage. (Nasheri, 2005)

The executive may also have prudential reasons to spy – or spy disproportionately – to secure her foreign policy goals, even when these aims are not sufficiently weighty to justify the spying, or even worse when the aims are unjust. For example, the George W. Bush administration bugged the United Nations (UN) Secretary General Kofi Annan's office to assess which UN members were likely to vote for the Iraq War resolution in the UN's Security Council. (Gendron, 2005, 411) This spying was unnecessary, disproportionate, and lacked a just cause. More extreme cases are easily imagined, such as an executive spying to bolster her case for an unjust war. Spying, she might reason, could uncover information that could be twisted and spun in such a way to make her case to the public more plausible.

This last example is particularly important, given executives have many times throughout history spied to prevent phantom threats. They have spied, that is, to prevent inflated, remote, unsubstantiated, or even implausible threats. Purported threats, such as weapons of mass destruction in Iraq (or Iran), dominos falling in South East Asia, the “missile gap” during the Cold War, and Islamic terrorists acquiring nuclear weapons have led to spying on a massive scale. In some of these cases, spying may have been justified, but the concern is that in other cases the threats were blown out of proportion, and thus much of the spying ordered in response to the threats was not justified. Precisely why the executive is prone to chase ghosts is not entirely clear. Sometimes combating even imagined threats may be electorally advantageous, assuming these threats can be sold to the general public. But the empirical tendency of the executive to err in the direction of conjuring up or inflating threats rather than ignoring or downplaying them is an important consideration for institutional design. It suggests not only that the executive should be institutionally checked, but also that she should be checked by a cooler, more cautious agent.

Since both intelligence agents and their handlers have professional, prudential, and/or political incentives to spy on foreign states that diverge frequently from their moral responsibilities, the first step toward extending the institutionalization argument is successful. But perhaps the dangers of abusive spying on foreign states are insignificant, leaving no reason to institutionalize the principles of government spying on foreign states.

Since I have already mentioned numerous harms that could follow from the misuse of government spying on foreign states, it should be clear that I do not think the dangers of abuse are trivial. But let me say more about the potential harms that could follow from state spying.

In so doing, I shall again reference the potential harms from domestic spying that

I developed in chapter 5. But since I developed these harms under the assumption that the target of spying was an individual, I need to consider how, if at all, my analysis changes when the target of spying is a state. As a preliminary to answering this question, I should say that while I think it makes perfect sense to talk about a state being harmed, I *do not* think that it makes sense to believe that these harms are morally important unless they lead to harms to individuals.¹⁰ Hence I will attempt to show how spying on foreign states leads to nontrivial harms to individuals.

Modern states keep a breathtaking number of secrets. Even the most open societies conceal information about their military capabilities and strategies, their positions in international negotiations, what they know about other states, and how they collect intelligence around the globe. In many cases these secrets are kept for good reasons, since their revelation could lead to considerable risks or harms to the state and its citizens. For example, in order to plan for its defense, a state needs to know its greatest military vulnerabilities, but if this knowledge were acquired by the state's enemies, the state and therefore its citizens could be rendered more vulnerable to violent attack.

So spying on foreign states can lead to the collection of secrets, which can then be used to do all kinds of harms to the state and its citizens. Some of these harms are severe, like those I mentioned, but others are more prosaic. Spying could, for example, harm states in bilateral or multilateral negotiations, or it could disadvantage states competing for foreign investment. It is important to reiterate, however, that harms to the state and harms to its members need not always coincide, and it is only the latter with which we are concerned ethically. Some states are manifestly unjust, perpetrating genocide or other egregious acts on their own citizens. Spying, when it

¹⁰For a defense of this view see Pogge (1992) and Beitz (2005).

harms *these* states may be justified, since it may lead to stopping or reducing the severity of unjust acts.

Spying on foreign states can also lead to corruption, both of individuals and institutions. Indeed, the Cold War industrial spy complexes in the United States and the Soviet Union were arguably manifestations of such corruption. Both countries created many overlapping institutions and agencies to spy, and thereby created (and reinforced) a bureaucratic interest in continued spying. The case for spying became as much about keeping agencies busy as it was about countering serious threats. (Cf. Moynihan 1998, especially Chapters 6 and 7.)

Third, whereas in the cases of spying on citizens and foreigners the risks of spying conditioning behavior were significant, it is not clear that spying on foreign states leads to similar conditioning. Citizens are relatively powerless compared to their states: states have seemingly unlimited resources to affect their citizens' conduct, and individual citizens have little capacity to retaliate against state attempts at control. In contrast, the balance of power among states is more equal. No doubt there are powerful hegemonic states and weak relatively undeveloped states, and the former can, with carrots and sticks, sometimes control the conduct of the latter. But most state-to-state dyads feature relative equality, and control is more difficult to exercise.

More importantly, privacy norms play a different role between states than they do between states and individuals. In chapter five, I argued that privacy norms were crucial for the protection of individuals against the powerful force of majority opinion. Without strong privacy norms, individuals cannot engage in a range of harmless or beneficial behaviors. Those who have counternormative beliefs or inclinations are often forced to live in fear or self-denial. The force of majority opinion in international affairs is, in contrast, less prodigious. The norm of noninterference that prevails between states in the international community is much stronger than the norm of non-

interference typically protecting citizens against their states. States generally have little reason to fear interference from other states, so long as their conduct is harmless and non-threatening. Further, interference with another state is costly, difficult and risky, not just because of the relative power equality between states, but also because the tools of interference are comparably blunt.

I do not wish to claim here that spying on foreign states has *no* conditioning effect. This claim would clearly be false. To cite one example of such an effect, both Russian and American nuclear proliferation policies have been conditioned by the other side's capacity to covertly determine the type and extent of nuclear stockpiles. (Chesterman, 2011, 33) My claim is rather that the conditioning effect of spying on foreign states is relatively minimal, and that it is practically nonexistent when states are not engaged in conduct that is harmful or threatening to other states.

Spying can, fourth, lead to a loss of enjoyment, a host of negative emotions, such as paranoia, and a loss in status. Since strictly speaking, states do not enjoy things, have emotional responses, or experience the harms of a loss in status, if these harms are important in the context of spying on foreign states it is because they are experienced by the target state's citizens. Citizens are unlikely to have the enjoyment of their daily activities altered in any serious way by suspecting that their state is being spied upon. But such suspicions, in some contexts, could lead to emotional and status harms. Those who identify with their states, for example, can be humiliated or made to feel second class when they learn their state is under covert surveillance by a hegemonic or colonial power.

Fifth, when its spying on a foreign state is suspected or revealed a state can lose trustworthiness, but the effects of this loss in trustworthiness is different than it is in the domestic context. In the domestic context, the chief worry is that governmental institutions cannot function smoothly in the absence of sufficient trust. Citizens have

to trust their governments to some extent for democracy, markets, etc. to work. In the international context, in contrast, states that lack trustworthiness may not be able to enter into mutually beneficial agreements or they may struggle to develop and keep allies. This trust between states is likely connected to the general welfare of citizens in those states, but it is unlikely as crucial for citizens to live a flourishing life as is the trust between citizens and their own government. As Hobbes (1994 [1651], Book I, Chapter xiii) suggests, citizens can live decent lives, even while their states are in a “posture of war,” “with guns upon the frontiers of their kingdoms, and continual spies upon their neighbors.”

Finally, in extreme cases a breach of trust from spying can lead to retaliation. This risk, I argued, is minimal in the case of domestic spying and more serious in the case of spying on foreigners. It is *most* severe, I want to suggest, in the case of spying on foreign states. Retaliation may follow the logic of tit-for-tat: one state’s spying triggers spying by the target state in response. But retaliation can also be more severe. Captured spies, for example, are often imprisoned and sometimes killed. Spying can even in rare cases lead to war. Kant was perhaps the first to articulate this worry. He says, “spies (*uti exploratoribus*)” exploit “only the dishonorableness of others (which can never be entirely eliminated).” (Kant, 1991, 109-110, his emphasis) Such practices “destroy the trust requisite to establishing a lasting peace in the future” (Kant, 2012, 117); they “cannot long be confined to war alone[they] will also carry over to peacetime and will thus undermine it.” (Kant, 1991, 110) Hence, whereas the risk of retaliation is minor in the domestic context, it can be considerable when governments spy on foreign states.

We can now conclude that the argument for institutionalization remains on solid ground in the case of spying on foreign states, just as it did in the case of spying on foreign individuals. Intelligence agents and their political handlers have professional,

political and/or prudential incentives favoring spying on foreign states, and the potential harms that can follow from government spying on foreign states are appreciable. Perhaps the case for institutionalization is weakest in the context of spying on foreign states, but it remains sufficiently strong for the argument to go through.

That government spying on foreign states should be institutionalized does not, however, suggest that the principles by which government agents should spy when they spy on foreign states should be identical to those that they should use when they spy on their own citizens. The consequences of spying on foreign states may be sufficiently different that they merit distinct principles.

I want to suggest that the principles developed for domestic spying ought to be altered in four ways for spying on foreign states. Two of these modifications are similar to the modifications made in the case of spying on foreigners above. First, as was the case with spying on foreigners, more just causes should be admitted than in the domestic case, since other things equal, spying (unless it significantly risks retaliation) is likely to be less harmful when it is on foreign states than when it is on citizens. The justificatory bars for each of the three types of spying (identified in chapter 6) should again be shifted down, and arguably these shifts should be larger than they were in the case of spying on foreigners, since some of the harms in the cases of spying on citizens or foreigners either do not apply or do not apply to nearly the same extent in the case of spying on foreign states.

The principle of just cause also might require slight alteration in cases when other states do not respect their obligation not to spy. Because of the anarchic nature of the international system of states, states cannot always expect that other states will uphold their obligations, which raises the question: Should a state uphold its obliga-

tions not to spy on another state if it has reason to believe that the other state is not fulfilling its obligations not to spy? France, for instance, is often accused of spying on its allies to collect industrial secrets, which it then passes on to French corporations, thereby providing those corporations with a competitive advantage. Suppose the U.S. government knows that the French are engaging in this sort of industrial espionage on American soil. How, if at all, does America's knowledge that France is not fulfilling its obligations not to spy affect America's obligations not to spy?

It is implausible, I want to suggest, to think that America's obligations suddenly disappear and thus that America has a general permission to spy on France. The fact that French intelligence agencies steal the blueprints of consumer electronic devices made in the United States would not, for example, permit American intelligence agencies to steal French nuclear secrets or bug the residence of the French president.

But it is also implausible to think that America's obligations not to spy remain entirely unchanged. Leaving the obligation not to spy unchanged when others do not respect it would provide incentives to countries to cheat and it would impose undue penalties on virtuous countries. At the very least, then, America should be permitted to spy on France in order to foil France's spying, assuming of course that America's spying meets conditions of proportionality, minimization, and necessity. Arguably, however, this is no significant amendment to the principles that I developed above, since countries already have a just cause to spy on foreign states in order to prevent fraud, deception, or theft.

Beyond the permission to engage in counterespionage, America may also be permitted to spy on France on grounds of fairness. A variety of international practices that promote the general good, such as trade or treaty making, have built-in standards of fairness. These standards of fairness often include norms of appropriate information collection, norms which usually prohibit spying. If, as is no doubt the

case in the French example, a country violates these norms of appropriate information collection, there may be instances when other countries engaged in the practice can justifiably follow suit and violate the same norms in order to put themselves back on to a level playing field. So, we can add to the list of just causes promoting the fairness of a justified practice.

It also makes sense, third, to alter the principles of just cause and minimization given the risks of retaliation when one state spies on another. Since the risks of retaliation can be so severe, only preventing the gravest harms would be a just cause for any spying that materially risks retaliation.

One might point out that these risks are accounted for by applying the principle of proportionality, and thus no change to the principle of just cause is necessary. As a theoretical matter, this point is correct, but given the possible severity of retaliation, it seems judicious to formally account for the threat of retaliation in order to ensure the risks of retaliation are fully accounted for in deliberation. Further, in the domestic case, I suggested a number of practical considerations for the principle of minimization, for example ensuring that all reasonable precautions are taken to secure the information collected from spying. To these considerations should be appended a requirement to take all reasonable precautions to prevent retaliation.

The fourth and most significant change to domestic principles I want to suggest is dropping the principle of discrimination. As I discussed above, government spying is much less likely to condition the behavior of states than it is to condition the behavior of individuals. The consequence is that the importance, on utilitarian grounds, of the principle of discrimination is significantly diminished.

Dropping the principle of discrimination permits governments to cast wider nets when they spy on foreign states than when they spy on citizens or foreigners, and other things being equal wider nets can produce more benefits. The chief worry about

letting governments spy on innocents in the domestic case was the conditioning effect spying might have on harmless or even beneficial behavior. But, as I argued above, the conditioning effect of spying on foreign states is likely to be minimal. Hence states are more likely to produce net benefits by rigorously pursuing just causes, even when this sometimes means spying on “innocent” states.

One might object that dropping the principle of discrimination is counterintuitive, since it gives no special assurance to a state’s allies. In response it should be noted that if states followed the four principles (without discrimination) it is still likely that most of their spying would be on their enemies. Spying on allies would thus remain comparatively rare.

Further, sometimes states’ allies attempt to cross, betray, or take unfair advantage of them. Allies, after all, are not equivalent to innocents: it is perfectly possible that a state’s allies engage in all manner of wrongs, wrongs which in some cases justify spying. Without the principle of discrimination if a state’s allies do engage in harmful activities, they open themselves up to be spied upon, and plausibly this is a good thing, since their machinations are more likely to be discovered and prevented.

I have assumed so far that the benefits of spying on foreign states vary mostly in degree and not in kind from the benefits of spying on (domestic or foreign) individuals. This assumption permitted me to focus primarily on the how the potential harms of spying vary from case to case while fashioning principles. Let me now defend this assumption.

In his widely read textbook on intelligence, Mark Lowenthal (2009) argues that states have intelligence agencies for four main reasons: to prevent strategic surprise,

to provide long-term foreign policy expertise, to support the policy making process, and to maintain national secrets. Let me consider each of these benefits in turn.

The term “strategic surprise” is used in the intelligence literature to refer to events such as Pearl Harbor or 9/11, events that are unforeseen and significantly undermine a nation’s interests. Historically the threats of strategic surprise have been from states, but as the case of 9/11 demonstrates, this tendency is rapidly changing. The most dangerous threats to states are now often non-state actors and they even sometimes reside within states’ borders. Hence the benefit of preventing strategic surprise is not exclusively produced by spying on foreign states. Increasingly the prevention of the most dangerous threats to states involves spying both on citizens and foreign nationals.

Lowenthal, second, argues that intelligence agencies provide long-term foreign policy expertise. The analytical and executive positions in intelligence agencies, he says, tend to have greater stability than in agencies focusing on defense or diplomacy. Hence “much knowledge and expertise on national security issues resides in the intelligence community.” (3) Lowenthal is no doubt right that intelligence agencies possess a great degree of knowledge and expertise, and he is probably right that a good deal of this knowledge and expertise is regarding foreign policy. But this knowledge and expertise is increasingly about foreign individuals and non-state actors, not just about foreign states. Further, domestic intelligence agencies, such as the FBI or MI-5, possess deep knowledge and expertise regarding domestic policy.

Third, intelligence agencies support the policy making process. As I just suggested, however, intelligence agencies support both the domestic and the foreign policy making process. Hence this is not a benefit unique to spying on foreign states.

Finally, Lowenthal points to the importance of intelligence agencies for securing national secrets. Intelligence agencies are constantly on the lookout for hackers and

spies trying to steal the government's secrets. Indeed, counterintelligence is a central function of today's intelligence agencies. But in keeping with the trend of the most dangerous threats coming increasingly from individuals and non-state actors, the threat of a state's secrets being stolen is no longer just from foreign states. State secrets are vulnerable to hackers and terrorist groups, who may wish to use the information to harm the state or just to sell the information on the open market.

So, while I think Lowenthal is mostly right about the main functions of intelligence agencies, none of the benefits that accrue from having intelligence agencies accrue exclusively from spying on foreign states.

But perhaps even Lowenthal has failed to see one of the major benefits of spying on foreign states. Glen Sulmasy and John Yoo (2007) argue that spying on foreign states reduces the risks of interstate wars. "Any international agreement or norm that makes it more costly for states to gather better information, and hence reduce uncertainty," they argue, "would only increase the possibility of war."¹¹ (636) In other words, spying on foreign states tends to equip states with more and perhaps better information, which reduces uncertainty, and thereby reduces the likelihood of war. If Sulmasy and Yoo are correct, then there is an important class of benefits to spying on foreign states that I have not yet considered, and plausibly by incorporating these benefits one would reach a different set of principles for spying on foreign states than those I reached above.

¹¹Although they do not say so explicitly, presumably their argument does not extend to the domestic case because of a set of observations often made by international relations theorists. Those theorists (e.g. Waltz (1959)) have often pointed out that the system of states is best characterized as anarchical. When states have conflicting interests, there are (typically) no higher authorities to appeal to. Conflicts must be sorted out between and among states. In contrast, domestically states are best characterized as hierarchical. When conflicts arise between individuals or corporations, these individuals or corporations can resolve their conflicts in the courts. Whereas conflicts between states can lead to violence and war, states settle conflicts between individuals or corporations within their borders.

Their conclusions follow, they think, from a bargaining model of war. States have preferences over issues, which sometimes conflict. Most of the time, these conflicts can be resolved with bargaining because there exists a set of outcomes on the issues that both parties would accept (a “win set”). But there is no guarantee that the result of bargaining will be a point in the win set. Parties in negotiations rarely show their hands. If an outcome is not in the win set, both parties do an expected value calculation to decide whether to go to war. If this calculation is positive, war ensues.¹²

With the bargaining model of war taken as a given, Sulmasy and Yoo reason that “[b]etter information from intelligence gathering, whether covert or overt, can actually promote the potential for peace and reduce international tension.” (634) Why would this be the case? They argue that “a significant obstacle to reaching a negotiated settlement is the problem of imperfect information” (635) Neither side knows the other’s true bargaining position. Either side could be bluffing. So in some cases there are acceptable outcomes to both parties, but the parties, because of imperfect information, cannot converge on any of these outcomes. One way to solve or alleviate the problem of imperfect information, Sulmasy and Yoo think, is to permit both parties to spy. Good intelligence reduces uncertainty. “Better intelligence allows a state to determine more accurately the military strength of the other side and the value of the disputed asset.” (Ibid) Thus, better intelligence gained by spying reduces uncertainty and thereby reduces the likelihood of war, and rules that restrict permissible spying abroad “would only increase the possibility of war.” (636)

I want to suggest some reasons to be skeptical about this argument. First, spying

¹²If the outcome of bargaining $< p(\text{winning war}) * [\text{value of preferred outcome} - \text{likely costs of war}]$, for one of the parties, then war ensues. The likely costs of war are a function of the relative military strengths of both sides. (634-635)

may produce information, but more information does not obviously lead to greater certainty, especially in the intelligence world. Intelligence agencies already collect so much information that it is often difficult to know which information is important. This problem is referred to by intelligence analysts as either the “wheat versus chaff” problem or the “noise versus signal” problem. (Lowenthal, 2009, 72). What drives uncertainty is usually not a lack of information, then, but instead insufficient expertise for placing the information in context, or the skill to see what is crucial and what is irrelevant.

There is also the problem that the information collected from spying has a non-trivial chance of being false. Almost every country has a counterintelligence operation the purpose of which is not just to catch spies but also to feed them with disinformation. Spies themselves also sometimes have incentives to feed false information to their handlers. Recall the “Curveball” case during the run up to the Iraq war.¹³ Thus, although it may be safe to assume that spying will produce new information, there is no reason to jump to the conclusion that this new information will reduce uncertainty.

A final concern is that spying could lead to conflict rather than prevent it. This was Kant’s argument, canvassed at the end of chapter 5. High profile espionage cases have in the past disrupted international negotiations. Consider the 1960 U2 incident when American pilot Francis Gary Powers’ spy plane was shot down over Soviet airspace. Since the incident, which came weeks before the East-West summit in Paris, turned out to be a significant embarrassment for the United States and led to a precipitous worsening in U.S.-Soviet relations, it does not seem like a stretch to

¹³See Jervis (2006, 29). The agent known as “Curveball” was an Iraqi defector, who falsely claimed to have worked in a mobile biological weapons laboratory. The Bush administration used his claims as evidence that Iraq had reconstituted its weapons of mass destruction program.

think that the U2 incident made war between the United States and the Soviet Union more rather than less likely.

There are strong reasons to question the plausibility of Sulmasy and Yoo's argument, then. But for the sake of argument, let us suppose that their argument is sound. How, if at all, would their conclusion affect the arguments I made for principles for spying on foreign states above? First off, it would *not* affect the principles of proportionality, minimization, or necessity. The calculations of proportionality would change with the incorporation of the new potential benefit, but the *need* for the principle would not. Further, agents should still be morally required to show that less costly alternatives to spying are not available and that all reasonable precautions are taken to minimize the harms of spying.

Sulmasy and Yoo's argument may, however, affect the principle of just cause. If preventing interstate conflict can be done indirectly by covertly collecting information about foreign states' military capabilities, then perhaps there is a just cause for spying whenever it is likely to produce information about a foreign state's military capabilities. I do not think this thought is defensible, however. In most state-to-state dyads the probability of war is vanishingly close to zero. Hence permitting states in these dyads to spy on one another is unlikely to reduce the likelihood of war in any significant way. Even in dyads where the probability of war is nontrivial, permitting all spying to collect information about foreign states' military capabilities seems mistaken. Some states in these dyads face nontrivial likelihoods of war *because* of their own unjust actions, and it does not seem plausible to give greater permissions to states to spy because they behave unjustly. Those states which have not behaved unjustly likely already have a just cause to spy, since spying to protect themselves against violent attack obviously counts as a just cause. Hence, even if Sulmasy and

Yoo's conclusion that spying on foreign states reduces the likelihood of war is correct, it does not materially alter the principles I argued for above.

7.3 Institutionalizing Principles for Foreign Spying

So foreign spying – whether it is spying on foreigners or spying on foreign states – should be institutionalized and it should follow principles not drastically dissimilar to those regulating domestic spying. In this final section I want to say a few words about the institutional implications of this argument. My remarks, I hope, will serve as a bridge to the institutional analyses that follow in the final chapters.

There are two broad options for institutionalizing constraints on foreign spying: codifying constraints in international law and extending constraints on domestic spying to foreign spying. The latter option is considerably more promising practically, and hence it is the option that I explore in the remaining chapters, but the former option is less problematic than has been suggested by many international legal scholars.

There are roughly three kinds of institutions that comprise international law: customary international law, treaties between or among states, and international organizations.¹⁴ Presently there are no international organizations regulating spying. Nor are there significant treaties to speak of, which restrict spying between or among

¹⁴These are sometimes reduced to two because international organizations are created by treaty, but as Buchanan (2010, 80) notes international organizations are increasingly making international laws.

states in significant ways.¹⁵ International customs on espionage appear to be in tension with one another; hence they provide scanty guidance for foreign spying. Many international legal scholars believe that intelligence gathering in peacetime is either legal or at least not obviously illegal. Indeed, as Sulmasy and Yoo (2007, 628-629) point out, “Nowhere in international law is peaceful espionage prohibited” and “no serious proposals have ever been made to prohibit intelligence collection as a violation of international law.” Yet, it is also generally accepted in international practice that states have the authority to punish captured spies, and these punishments can be severe. Hence states are free to severely punish people for engaging in what is not technically considered an offense.

If states wish to codify constraints on spying into international law, they will have to do so by making treaties, by setting up new international organizations and empowering them to regulate state spying, or by expanding the authority of existing international organizations. Many legal scholars think such codification is little more than a pipe dream. Radsan, for example, thinks unchecked foreign spying is an unalterable feature of the nation state system. So long as there are nation states, there will be aggressive spying among states.¹⁶ Sulmasy and Yoo (2007, 636) think international regulatory efforts could hinder important intelligence aims. “Any international regulatory scheme could hamper efforts to frustrate al Qaeda,” they argue. They also point out that any international agreement regulating spying would unlikely be universally reciprocated.

¹⁵There are, however, agreements for intelligence sharing. For example, in 1947 the United States and Britain signed the United Kingdom-USA Intelligence Agreement (UKUSA), which was later joined by Australia, Canada, and New Zealand. See Chesterman (2006, 1093 fn. 97).

¹⁶“Until the system of nation states is replaced, until regional and international integration really take hold, intelligence services will be around to do their states’ bidding. National intelligence services are far from being integrated into regional commands. . . or into an international peacekeeping function. . . Intelligence. . . remains just one more arena for national competition.” (613)

But these critiques seem to be aimed at attempts to prohibit spying entirely, a proposal I obviously do not support, given the principles I argue for above. If the principles I have endorsed were institutionalized, constraints on foreign spying would not eliminate spying abroad, instead they would rein in unjustified or abusive spying. Of course, there is some danger that the institutions would thwart crucial instances of justified spying. But this danger can often be minimized with institutional innovations.¹⁷

If the proposal is for institutions that codify the principles of foreign spying I argued for above rather than for outright prohibitions on spying, the problem of reciprocity also seems to be a dead letter. Sulmasy and Yoo suggest that if most states reached an agreement on intelligence gathering, non-state groups like al Qaeda would not honor such an agreement. But it is not clear why this matters. Any plausible set of institutions is very unlikely to restrict spying against groups like al Qaeda. If anything, institutions will ensure that more of the intelligence resources of those party to the agreement will go toward spying on serious threats to international order, rather than toward political vendettas or bureaucratic hobby horses.

A more interesting problem would be if countries agreed not to engage in industrial espionage, and some failed to honor the agreement. Spying in these cases could lead to competitive advantages for defectors. But this is really only a problem if the agreement has no teeth.¹⁸ One could envision a set of institutions that function like the WTO's anti-dumping agreements. If companies export products at prices lower than they charge domestically, then they are dumping these products into interna-

¹⁷I discuss some of these innovations in the final two chapters.

¹⁸It could also be a problem if the spying was almost never discovered. Presumably if it is only discovered some of the time, penalties could be set high enough to make it in countries interests not to spy.

tional markets. The WTO provides a framework for countries who are harmed by dumping to legally retaliate against those who dump.

Hence there are reasons to think that some international lawyers overstate the case against institutionalizing ethical principles for spying into international law. Yet these international lawyers are probably right to be skeptical about international institutions, at least in the short run. There is little if any political will, even among liberal democracies, to build international institutions in areas connected to national security like spying.

On February 18, 1976 President Gerald Ford signed Executive Order 11905, which among other things banned political assassination.¹⁹ Ford's order was in response to the Church Committee's investigation of the intelligence community's many attempts to assassinate foreign leaders, including Fidel Castro of Cuba, Patrice Lumumba of Congo, Rene Schneider of Chile, and Rafael Trujillo of the Dominican Republic. The Committee concluded that the procedures for using assassination were unclear, in fact, so unclear that the committee could not verify whether many of the assassination attempts it reviewed were authorized by the proper authorities. Furthermore, it concluded that, short of war, assassination was contrary to American values.

Executive Order 11905 is an instance of a state constraining its behavior abroad with domestic policy. One can imagine similar policies with respect to spying, and these policies have the advantage of not requiring international collective action. Rad-

¹⁹The order read, "No employee of the United States Government shall engage in, or conspire to engage in, political assassination." The order was later strengthened by Carter's EO 12036 and Reagan's EO 12333 but then weakened by interpretations under the Clinton and George W. Bush administrations that excluded individuals connected to terrorism.

san (2007, 619) considers and rejects a similar policy. It is worth quoting his remarks at some length here because I fear he rejects a straw argument. He argues:

It is possible for the U.S. espionage statutes to be amended to have full extraterritorial effect. . . . But the reward for such self-righteousness would be mockery and disbelief. Other states would not then preclude their intelligence services from stealing secrets from foreigners and foreign governments. Even if they took that step, they would not enforce the preclusion.

It is odd that Radsan considers expanding America's espionage statutes, which make it illegal to steal diplomatic and military secrets. If the goal is preventing executive and bureaucratic overstep it would seem more relevant to extend to people residing outside America's borders the protections given to American citizens in the Fourth Amendment, Title III and FISA. Most spying is not done by adventuring American citizens stealing secrets abroad, but rather by America's richly financed intelligence agencies. Second, although "mockery and disbelief" may follow from a policy that entirely prohibits foreign spying, most countries abroad may think a policy of extending constraints on foreign spying not only virtuous but prudent, since constraints on foreign spying are necessary to ensure that foreign spying is done in a focused and efficient way and with a modicum of respect toward non-citizens, to minimize backlash.

7.4 Conclusion

I argued in this chapter that government spying on foreign individuals and on foreign states should be institutionalized and that it should follow principles similar but not identical to those that governments should follow in the domestic context.

In the final section, I argued that codifying principles of ethical spying into international law faces thorny (although not intractable) collective action problems that governments are not particularly enthusiastic to address. Since extending domestic political institutions has more promise practically my focus in the final two chapters will be almost exclusively on these domestic institutions.

Part III

The Institutions of Government

Spying

Chapter 8

Controlling Government Spies: The American Model

In the last five chapters I developed and defended principles to regulate foreign and domestic government spying. I provided no reason to think that government agents will faithfully follow these principles even if they were made aware of them, however. On the contrary, I argued that government agents often have prudential and political reasons to spy that separate from their moral obligations, and that without political institutions to compel compliance, even the most virtuous government agents are likely to fall prey to the temptation to spy for personal or political gain.

Yet my call for institutions to compel compliance with ethical principles is not enough, since it is far from clear *which* institutions are most likely to secure compliance. Achieving ethical conduct in politics requires careful institutional analysis and design; it requires taking the real-world circumstances of political actors seriously; and it requires marrying the best normative theory with the best social science.

Hence my aim in this chapter and the next is to begin this project of understanding which institutions are best suited to ensure that government agents comply with

ethical principles for spying. In this chapter I examine and critique the two primary institutions designed to control America's spy agencies. In the next chapter I step back from America's institutions, characterize the universe of possible institutions to constrain spy agencies, and propose a set of institutions better suited to secure compliance with ethical principles.

The institutions regulating spying in the United States are complex. Generally, government spying outside of America's borders is controlled primarily by legislative oversight, whereas government spying within America's borders is regulated mainly by judicial review. But this characterization admits many exceptions. Most notably, for domestic spying, judicial review is typically only required for cases involving wire-tapping. Many other kinds of spying, such as tailing a suspect or collecting a suspect's telephone or bank records do not require judicial warrants. But the characterization is precise enough for my purpose in this chapter, which is not to give a detailed description of American institutions, but rather to pick out key components of those institutions and to explore their strengths and limitations.

I defend two main claims in the chapter. The first is that the legislative oversight of spy agencies as a mechanism of control is neither effective nor normatively appealing. The second is that while judicial review is a more effective mechanism of control than legislative oversight, it is also on shaky normative ground, and even more importantly it does not inform or shape the strategic decisions made by spy agencies. Hence as a standalone institution, neither legislative oversight nor judicial review is an attractive mechanism of control. Either or both of the mechanisms of control may play a role in the optimal set of institutions (a consideration I take up in the next chapter), but by itself neither institution is sufficient to secure compliance with the ethical principles that I developed in previous chapters.

My argument proceeds as follows. In the first section I focus on legislative over-

sight, the primary mechanism of control employed for spying outside of America's borders. I examine two waves of theorizing in political science about bureaucratic agencies, the first maintaining that bureaucratic agencies operate independently of political control and the second claiming that agencies are controlled by the Congress. I then outline the intelligence literature's engagement with and criticism of the second wave. Intelligence scholars persuasively argue that legislative oversight is considerably less effective in intelligence than in other areas of public policy. Finally, I add to the problems with legislative oversight by arguing that oversight, at least in the American case, is normatively problematic. Unrepresentative and inexpert subcommittees cannot guarantee that most people's interests are accounted for.

In the second section I turn to judicial review, the main mechanism of control for spying within America's borders. I argue that when it is performed *ex ante*, judicial review can be a more effective mechanism of control than legislative oversight, but it nevertheless runs into the same structural and normative problems that legislative oversight faces. Two further problems confront *ex ante* judicial review: it is costly compared to other mechanisms of control, and more importantly, it does almost nothing to structure the most important strategic decisions made by spy agencies. It does not, for example, ensure that spy agencies target the most dangerous threats or that they employ their scarce resources effectively.

8.1 Legislative Oversight

Until the beginning of the 1980s, the dominant view in political science was that bureaucratic agencies operate independently, carrying out their own agendas, with-

out considering the policy preferences of other political actors.¹ According to this view, Congress mostly fails in its responsibility to oversee and shape the conduct of the bureaucracy. Although Congress has the tools to oversee agencies' conduct, it rarely uses them. It calls few hearings and fewer investigations, and the hearings and investigations it does call are often conducted in a superficial way; it lacks detailed knowledge of agency operations and the effects of agency choices; and it rarely passes legislation to alter the structure or conduct of agencies. (Weingast and Moran, 1983, 767) In Pearson's (1975, 281, 288) words, "...congressional oversight remains basically weak and ineffective...It is a vital yet neglected congressional function." The result is that the bureaucracy is "in many respects a prodigal child. Although born of congressional intent, it has taken a life of its own and has matured to the point where its muscle and brawn can be turned against its creator." (Dodd and Schott, 1979, 2).

Proponents of this view offer a variety of explanations for why Congress fails to control agencies. Chief among them is the information asymmetry that exists between bureaucratic agencies and the Congress. Bureaucrats tend to know their policy domain better than legislators. They have more expertise. But, more importantly, bureaucrats control the flow of information. To perform its oversight responsibilities, Congress depends on the agencies it is supposedly overseeing to provide timely, accurate, and complete information. As McCubbins, Noll and Weingast (1987, 251) point out, "[i]n a sense, the agency both keeps the books and performs the audit."

Beginning in the 1980s, scholars challenged and eventually replaced the traditional view.² These scholars did not dispute the empirical evidence brought to bear to

¹See e.g. Anderson (1975), Dodd and Schott (1979), Rourke (1976) and Wilson (1975, 1980).

²Weingast and Moran (1982, 1983); Calvert, McCubbins and Weingast (1989); Weingast (1981, 1984); Barke and Riker (1982); McCubbins and Schwartz (1984); Fiorina (1981, 1982); McCubbins (1985)

support the traditional view – that Congress calls very few hearings, for example – rather they sought to explain the facts differently. They drew heavily on principal-agent theory from economics. (Alchian and Demsetz, 1972; Holmstrom, 1979; Jensen and Meckling, 1976) And they argued that sometimes in order to shift blame or take credit or to utilize outside experts the Congress often has incentives to delegate certain activities to bureaucratic agencies. (Fiorina, 1977, 1982; Fiorina and Noll, 1978; Epstein and Sharyn, 1994; Huber, Shipan and Pfahler, 2001) The Congress does not give agencies unlimited discretion to perform these activities as they please, however. Instead it plays the role of principal, and the bureaucratic agencies to which it delegates stand in as its agents. Principals devise incentive mechanisms, that is rewards and sanctions, to shape their agents’ conduct. The better the principal designs incentive mechanisms, the less she has to invest in monitoring the moment-to-moment behavior of her agents. Her agents do as she wishes because their incentives are aligned with her aims.

The new view, often referred to as the “congressional dominance” theory, holds that Congress is a remarkably sophisticated principal. Congress has designed complex incentive mechanisms, which ensure that agencies do its bidding, without burdening it with costly monitoring. (Weingast and Moran, 1983; Moe, 1987) Congress does not neglect its oversight responsibilities, this new view holds, rather it rationally constructs – using various incentive mechanisms – a form of oversight that maximizes its net benefits.

To understand the incentive mechanisms that Congress purportedly creates for agencies, it is helpful to distinguish between two forms of oversight, proposed by McCubbins and Schwartz (1984). *Police-patrol* oversight encompasses most of what is typically thought of as oversight. By holding hearings, conducting investigations, etc., Congress periodically examines agency activities to detect and discourage aberrant

conduct. As mentioned, oversight of this sort is costly in time and resources, and it comes with the risk of signaling congressional mismanagement to constituents.

By employing *fire-alarm oversight*, on the other hand, “Congress establishes a system of rules, procedures, and informal practices that enable individual citizens and organized interest groups to examine administrative decisions (sometimes in prospect), to charge executive agencies with violating congressional goals, and to seek remedies from agencies, courts, and Congress itself.” (166) Members of Congress outsource monitoring to citizens and interest groups, who often have more time, resources, and motivation to track agencies’ activities, and they step in only when they are likely to receive credit for assisting their constituents. Fire-alarm oversight, McCubbins and Schwartz claim, “serves congressmen’s interests at little cost.” (168) Congress has neither the time nor the inclination to monitor all of the buildings in their neighborhood for fire, so it places “fire-alarm boxes on street corners, builds neighborhood fire houses, and sometimes dispatches its own hook-and-ladder in response to an alarm.” (166)

Examples of fire-alarm rules, procedures, and practices include requiring transparency or access to administrative decision making, giving standing to citizens and interest groups in the administrative decision making process, and facilitating collective action among relatively unorganized groups. Although members of congress predictably support fire-alarm rules, procedures, and practices, thereby empowering their constituents, it is important to note that these rules, procedures and practices tend to be “particularistic,” in the sense that they “emphasize[] the interests of individuals and interest groups more than those of the public at large.” (172) Members of Congress tend to “stack the deck” in favor of those constituents most likely to secure them reelection.

Empirical support of the congressional dominance theory has been mixed. Some

find evidence that Congress strongly influences agencies. (Moe, 1985; Scholz, 1991; Weingast and Moran, 1983; Wood and Anderson, 1993) Others find that Congress's influence is minimal. (Eisner and Meier, 1990; Moe, 1987; Wood, 1988).

The chief theoretical objection to the congressional dominance theory has been that the principal-agent model it is based on is too simple. Agencies have not one but multiple principals. At the very least, the objection goes, if one wants to understand bureaucratic behavior in the American context, then one must factor in not only the influence of the Congress, but also the influences of the President and the courts. In Moe's (1987, 477) words, "if it [the current theory of congressional dominance] is to live up to its promise, it must transcend its fixation on Congress and get on with the task of building a larger theoretical perspective."

Theorists have attempted to meet this challenge by developing a variety of spatial models of agency policy making (Ferejohn and Shipan, 1990; Hammond and Knott, 1996; Steunenberg, 1992; Shipan, 2004) These models situate bureaucratic agencies in a particular institutional context. They theorize that agency policy making is a sequential game. An agency first makes a proposal, and then various political actors have a chance to respond to the agency's proposal. In most models, the first responder is the congressional committee that has oversight responsibility for the agency, which chooses whether to accept the proposal (and end the game) or introduce a new bill. The bill is then, depending on the model, given an up or down vote on the floor, or it proceeds to the floor under an open rule. In the simplest models, the agency, the committee, and the floor are the only players. More complex models build in judicial review and Presidential veto power.

These models offer the most promise for understanding agency behavior. But they likely require considerable refinements for particular agencies and particular institutional contexts. Since existing models focus almost exclusively on regulatory agencies,

and since they have typically been tested on a single agency or agencies of the same type, how well they generalize to other kinds of agencies is not yet clear. (Meier and O'Toole, 2006, 179) It is therefore an open question whether models designed for the Securities and Exchange Commission or the Food and Drug Administration are also applicable for the FBI or the CIA.

The intelligence literature is less theoretical than the literature in political science. (Born and Caparini, 2007; Johnson, 2004, 2007; Smist, 1994; Tsang, 2007; Zegart, 2012) Its theoretical engagement has been limited primarily to pointing out problems with applying the congressional dominance theory to intelligence agencies. Like the literature on political science, the intelligence literature focuses chiefly on legislative oversight. It chronicles the history of intelligence oversight, with a particular focus on the American case; it collects the successes and failures of the American oversight model; and it provides a deep well of recommendations for reform.

In order to develop the intelligence literature's critique of the congressional dominance theory, it is helpful to first delve into its telling of the history of American intelligence oversight, for it is this history that motivates intelligence scholars' pessimism about intelligence oversight.

The history of American intelligence oversight begins with The National Security Act of 1947. It created America's first intelligence agencies, but it failed to create a detailed oversight framework. From 1947 to 1975 subcommittees within the Appropriations and Armed Services Committees in both the House and Senate divided oversight responsibilities. Oversight was performed informally, irregularly, and infrequently, usually by the committee chairs. No records were kept and the subcommittees sought almost no input from citizens, the press, or interest groups. According

to Zegart (2012) the subcommittees met only a handful of times annually, and in some years only once. The attitude in these meetings was not one of agency intransigence, but rather congressional reluctance. In the words of a former CIA legislative council, “We allowed Congress to set the pace. We briefed in whatever detail they wanted. But one of the problems was you couldn’t get Congress to get interested.”³

Congressional reluctance stemmed, in part, from congressional trust in intelligence agencies. As Congressman Robert Ellsworth noted, “The Political zeitgeist of the time was that the CIA was wonderful. In politics, anybody who wanted to make trouble for the CIA was seen to be a screwball and not to be countenanced.”⁴ It also stemmed from a worry that by inquiring into the agencies’ operations, members of Congress were only likely to discover information that they would rather not know. Leverett Saltonstall, a member of both subcommittees overseeing the intelligence community in the Senate during the 1950s, noted, “It is not a question of reluctance on the part of CIA officials to speak to us. Instead it is a question of our reluctance...to seek information and knowledge on subjects which I personally, as a Member of Congress and as a citizen, would rather not have.”⁵

Between 1947 and 1974 there were many opportunities to enhance congressional oversight of intelligence, but none were taken.⁶ For example, in 1966 Senator J.W. Fulbright introduced Senate Resolution 283. It promised to establish a Senate Committee on Intelligence operations, made up of members from the Armed Services, Foreign Relations, and Appropriations subcommittees. The Senate voted 61-28 against

³Quoted in Smist (1994, 5).

⁴Ibid.

⁵Quoted in Zegart (2012)

⁶According to Zegart (2012), the two chambers voted against over two hundred resolutions between 1947 and 1974 to consolidate or strengthen intelligence oversight.

the resolution. The belief among many was that the agencies would do their best work free from congressional meddling. In the words of Richard Russell, who for some time chaired both subcommittees on intelligence in the Senate, “If there is one agency of the government in which we must take some matters on faith, without a constant examination of its methods and sources, I believe this agency is the CIA.”⁷

As suspicion of the executive branch escalated in response to mismanagement in Vietnam and to the Watergate scandal (in which many believed the CIA played a role), Congress’s faith in American intelligence agencies began to crumble. Then came late December 1974. On December 22, 1974 the *New York Times* published Seymour Hersh’s now (in)famous article on its front page, detailing a laundry list of the CIA’s transgressions. Hersh had obtained a copy, or fragments of a copy, of what is now referred to as the “Family Jewels,” a document commissioned in 1973 by CIA director James Schlesinger to collect ongoing and past CIA initiatives that fell outside the CIA’s charter. The Family Jewels, which were only released publicly in June 2007, detailed extensive agency misconduct. For nearly the entire history of the agency, it had conducted activities outside of its charter, activities including assassination plots, domestic surveillance, illicit wiretapping, and experimentation on human subjects. Hersh’s article focused on domestic surveillance. He summarized the activities as follows: “The Central Intelligence Agency, directly violating its charter, conducted a massive, illegal domestic intelligence operation during the Nixon Administration against the antiwar movement and other dissident groups in the United States.”

Hersh’s article set fire to the trust that the American people and the Congress placed in America’s intelligence agencies. The CIA and its sibling agencies could no longer be trusted to carry out the people’s will or to protect the people’s interests.

⁷Quoted in Smist (1994, 6).

Although few doubted the necessity of having spy agencies, a consensus emerged that spy agencies could no longer go on operating unconstrained. Well-functioning intelligence agencies, most agreed, are vital to the safety of democracies, but *unconstrained* intelligence agencies can undermine the very values that they are meant to protect. The flourishing of liberal democracies thus depend on institutions that permit their intelligence agencies to identify and prevent threats from terrorist organizations and hostile foreign states, while at the same time preventing these agencies from undermining cherished liberal rights and democratic values.

The reaction to Hersh's article and the many that followed it in early 1975 was a series of investigations. First came a commission in the executive branch, appointed by President Ford and headed by Vice President Nelson Rockefeller tasked with investigating the accusations of intelligence abuses in the *New York Times* articles. But the Congress was no longer willing to trust the executive branch to rein in its intelligence agencies: both chambers launched their own investigations.

The first, and ultimately the most important investigation, took place in the Senate, beginning in late January 1975. Led by Senator Frank Church, a Republican from Idaho, the investigation lasted over a year. The committee produced 14 volumes of reports and transcripts. It focused on enumerating the abuses by America's intelligence agencies, and its list turned out to be startlingly long. Abuses by America's intelligence agencies were not occasional missteps carried out by a few overzealous agency employees. They were widespread, and often backed by the highest links in the chain of command. America's agencies plotted to assassinate foreign heads of state; they carried out extensive domestic spying and covert operations against non-violent political groups; and they experimented with various drugs on unknowing subjects. The committee concluded, "...that intelligence activities have undermined the constitutional rights of citizens and that they have done so primarily because checks and

balances designed by the framers of the Constitution to assure accountability have not been applied.”⁸

The Church committee recommended sweeping changes to enhance congressional surveillance and control of intelligence agencies. Of the 96 it put forward, none received more attention and a more immediate response than its final recommendation: “The Committee reendorses the concept of vigorous Senate oversight to review the conduct of domestic security activities through a new permanent intelligence oversight committee.”⁹

Less than a month after the Senate launched its investigation, the House launched its own investigation, chaired by Lucien Nedzi. The Nedzi committee, however, never began its investigation in earnest. It was plagued by internal dissension. In July, the House created a second committee, chaired by Otis Pike. But the Pike committee also suffered from internal discord. Although it completed almost a month of hearings, the House eventually voted to suppress its report. The report was later leaked, however, triggering investigations of the investigators, and undermining the committee’s credibility.

In contrast to the Church committee, the focus of the Pike committee’s investigations was not intelligence agency abuses. Instead it focused on the effectiveness of America’s intelligence agencies. It wanted to understand whether the agencies made the best use of tax-payer dollars.¹⁰ Like the Church committee, the Pike committee’s principal recommendation was the creation of a standing committee in the House for intelligence oversight.

⁸Book IV, Section A.

⁹Book IV, Section C, Part xii.

¹⁰See Smist (1994, 154).

The chief legislative response to the congressional investigations was the creation of select intelligence committees in both the House and the Senate. The select committees finally established a framework for legislative oversight of intelligence agencies. Impromptu oversight had been replaced by regularized meetings and formal reporting. The expectation was now that the congressional committees would be informed of the various intelligence agencies' activities. Congress' mantra concerning intelligence agencies had shifted from "trust" to "trust but verify."

According to the congressional dominance theory, the establishment of formal mechanisms of oversight should have ushered in a new era of congressional control of intelligence agencies. But intelligence scholars tell a starkly different history. In their telling, although the creation of the select committees was an important turning point for the accountability of intelligence agencies, Congress hardly has had the intelligence community on a short leash since the late 1970s. On the contrary, the period following the creation of the standing committees is riddled with intelligence abuses. In Johnson's (2004, 3) words, legislative oversight failed "to hold the intelligence community in check during the Cold War, leading to a significant erosion of civil liberties in the United States."

The conclusion that intelligence scholars have reached is that the congressional dominance theory is not true, at least not for intelligence. Maybe the congressional dominance theory is on less problematic grounds for regulatory agencies, like the SEC or the FDA, but the theory simply does not fit the available evidence for intelligence oversight, they argue. By any measure, Congress does not have a great degree of sway over intelligence agencies.

This conclusion – that the congressional dominance theory is wanting in the realm of intelligence – has led intelligence scholars to think carefully about what is special about intelligence. Loch Johnson and Amy Zegart have been at the forefront of sys-

tematizing the reasons why the congressional dominance theory does not apply to intelligence agencies. I summarize and expand on their reasons below.

First, the assumptions of the fire-alarm model on which the congressional dominance theory is built, although perfectly legitimate for certain regulatory agencies, are problematic for intelligence agencies. The congressional dominance theory assumes that the steps that legislators must take to avoid blame or take credit for policies are more or less clear and that these steps will pay off with a high degree of certainty. In intelligence, however, the link between policies and outcomes is often exceedingly difficult to forecast. In Zegart's (2012, 1296) words, "if a legislator argues that more money should be spent on Predator drones or tacitly approves a covert CIA action, he may not be regarded as a champion of American security; he may be accused of being a warmonger, a CIA lackey, or both." If, on the other hand, a legislator argues for increasing farm subsidies or for lowering taxes it is fairly clear how she will be regarded by her constituents.

The congressional dominance theory also assumes that agencies are agents of Congress. Yet the intelligence community sees itself primarily as an agent of the executive. Nearly all of the "products" produced by the intelligence community, such as the President's Daily Brief or National Intelligence Estimates are aimed primarily at members of the executive branch, for example. Further, in matters of national security, the President nearly always makes the call, and the intelligence community envisions itself supporting these decisions.

Finally, and most crucially, the congressional dominance theory assumes that there are numerous, powerful, informed citizens and/or interest groups that will sound the alarm about bureaucratic overreach. But interest groups focused on intelligence are comparatively scarce. Despite spending on defense taking up the lion share of the federal discretionary budget, among America's registered interest groups, less than

five percent focus on foreign affairs. (Zegart 2012, 1442-1451) Lobbyists in the defense industry make up only five percent of all lobbyists, and earmarks to intelligence projects make up less than one percent of the total budget spent on intelligence. (Ibid.)

The relative lack of interest groups might not be problematic if concerned voters could be expected to collectively pressure their legislators to reign in the bureaucracy, but they cannot. Intelligence is a national issue, which means concerned voters are geographically dispersed throughout the country. Hence compared to issue areas that are predominantly regional or local, such as farm subsidies or fishing regulations, voters concerned with intelligence face comparatively large barriers to collective action. (Olson, 1971)

Further, even if there were significantly more interest groups focused on intelligence, and voters could overcome their collective action issues, the problem remains that these interest groups and concerned voters will almost certainly not be well informed. No issue area is more opaque than national security, and within the national security bureaucracy, few agencies are more secretive than intelligence agencies. The mechanisms providing voters and interest groups transparency often do not apply to intelligence agencies. The Freedom of Information Act (FOIA), which mandates that government agencies disclose requested records, for example, has nine “exemptions,” limiting citizens’ right to know. First among these exemptions is any information classified as secret by an executive order. Since nearly everything that America’s intelligence agencies do is classified as secret, this exemption effectively guarantees that voters and interest groups will be ill-informed relative to their legislators. Yet the whole point of fire-alarm oversight is that legislators strategically rely on their more informed constituents to alert them to bureaucratic overreach.

Hence the assumptions of the fire-alarm model do not seem plausible in the world

of intelligence. To make matters worse, there are reasons to think police patrol oversight also is an insufficient means of controlling intelligence agencies. First and foremost, not just citizens and interest groups struggle to inform themselves about the activities of intelligence agencies, even congressional overseers are often kept in the dark. Committees do not observe the conduct of intelligence agencies first hand. They rely, rather, on agencies self-reporting their conduct. Such reliance creates what Rahul Sagar (2007, 408) calls a “structural dilemma.” Those who select the information to share with accountability holders are those supposedly being held accountable, that is, those who have an interest in withholding information to hide mistakes, misconduct, etc.

The problem with police patrol oversight is not just informational. Even if legislators were capable of collecting complete information on intelligence agencies’ activities, they probably would not do so. Members of the legislature tend to devote their time and attention to issues that will maximize their electoral success, and electoral success is driven primarily by delivering services and favors to influential constituents. But, as I just noted, very few voters and interest groups concern themselves with intelligence, and those that do tend to be dispersed throughout the country. Further, they have a limited capacity to monitor intelligence policy and outcomes, which makes it difficult to know whether policymakers have delivered on their promises. In the rare cases when legislators are motivated to oversee intelligence agencies, it is often for the wrong reasons. They use hearings and investigations to score political points, rather than to ensure that intelligence gathering is performed legally and effectively.¹¹

¹¹Johnson (2007, 61), for example, documents a recent exchange between two Senators after the Republican leadership in both the House and the Senate decided to dismiss allegations of FISA violations against the Bush administration for warrantless wiretapping. Senator Rockefeller, a democrat, accused the Senate committee of being “basically under the control of the White House, through its chairman.” Senator Roberts shot back and claimed that the Democrats were engaging in “gotcha oversight.”

Beyond the informational and motivation problems, lack of expertise further hinders police patrol oversight in intelligence. When overseers lack expertise, the risk is that they will not ask the right questions, request the right information, or even have the requisite knowledge to evaluate the performance of intelligence agencies. But for reasons connected to the already discussed problems of information and motivation, legislators rarely develop the required expertise for intelligence oversight. It is difficult to become an expert in an area in which you are often denied a complete picture of events, and few legislators invest in developing expertise if that expertise is unlikely to have electoral payoffs. Further, in the United States, the problem of expertise is exacerbated by term limits on the House Permanent Select Committee on Intelligence.¹²

The final problem with police patrol oversight is the risk that overseers will become coopted. The risk of cooption is connected to the problem of motivation. In both cases, the worry is that overseers will not conscientiously perform their responsibilities. But the worry about cooption is not that legislators will devote nearly all of their and their staffs' time to issue areas offering more electoral promise, it is that overseers will begin taking their orders from those whom they are supposed to be overseeing. They will, in other words, do the intelligence agencies' bidding. Cooption apparently has happened often throughout the history of American intelligence. Johnson (2007, 59), for example, suggests that during the period of 1992-2001 "Republican lawmakers on the [Senate Select Committee on Intelligence] or the [House Permanent Select Committee on Intelligence] became less gimlet-eyed reviewers of intelligence programs

¹²The Senate Select Committee on Intelligence had term limits for nearly thirty years, but abolished them in 2005.

than uncritical advocates of whatever the secret agencies wanted.”¹³

What the intelligence literature shows, then, is that (1) Congress does not control intelligence agencies as the congressional dominance theory suggests, and this is not surprising since (2) the assumptions of the fire-alarm model (on which the congressional dominance literature depends) are highly problematic for intelligence oversight, and since (3) problems plague more traditional (police-patrol) oversight mechanisms.

The problems with legislative oversight are in part structural, then. Legislators will, for instance, always struggle to gather sufficient factual information in order to paint a detailed picture of the activities of the intelligence agencies they oversee. But the problems are also due to legislators, and particularly legislative subcommittees, being comparatively poor principals. Let me develop this latter point more systematically. Doing so will simplify the comparison between legislators and alternative principals below.

The chief purposes of legislative oversight, I have suggested, are broadly to prevent abuse and to ensure that intelligence agencies effectively and efficiently use their resources. When legislative oversight is the primary mechanism of control, what constitutes abuse, as well as what constitutes the effective and efficient use of resources is determined by the beliefs and values of the intelligence agencies’ principals. As I write this sentence, in America that means seventeen men and three women in the House and twelve men and three women in the Senate. Thirty five individuals who hardly constitute a representative sample of Americans, and none of whom could claim to be intelligence experts before they were selected to serve on one of the committees.

Normally the lack of representativeness and expertise might not be problematic for an overseer. If the activities of the agency are open for all to see and judge, for

¹³He also notes that the dominant Republican on the House committee was labeled an “unrelenting cheerleader” for the intelligence agencies.

example, then the standards of abuse as well as good agency conduct become a part of the public conversation. The beliefs and values of the few could not so easily be imposed on the many. But the intelligence environment, as I have pointed out, is far from normal. Very little of what intelligence agencies do is visible to the public. The consequence is that the principal's beliefs and values are comparatively unmediated by public pressure. Hence *if* oversight were effective, the sub-committees would more or less get what they want from intelligence agencies.

Many might think this is a morally dubious state of affairs. Why should a relatively small, unrepresentative, and inexpert group that self-selects serve as the moral compass of intelligence agencies? Such a group, one might think, is dangerously likely to put intelligence agencies to the service of its own personal and political interests rather than to the service any common interests. Certainly there are considerations of secrecy, politicization, and specialization. For example, debating the intricacies of intelligence policy on the floors of the House and the Senate, even if those debates were sealed from public eyes and ears, may lead to worse outcomes than debates in subcommittees. But is it plausible to think that the best possible intelligence overseer is the one Americans currently have? I do not think it is.

A good overseer, that is, one who will induce spy agencies to comply with ethical principles, I want to suggest, will grade highly on four criteria. It will have a high level of expertise, it will be trustworthy with secrets, it will be motivated to conscientiously perform its oversight responsibilities, and it will be unbiased. I have already spoken of the importance of expertise and motivation. Let me briefly touch on trustworthiness with secrets and unbiasedness.

When overseers cannot be trusted with secrets, both the effectiveness of the oversight relationship and the success of intelligence operations diminishes. The effectiveness of the oversight relationship decreases because when intelligence agents fear that

the information they share with overseers could be leaked to hostile parties, they will be less likely to share information. Leaks can undermine the success of intelligence operations by tipping off potential targets of observation, allowing them to avoid being spied on, or permitting them to feed the spies with misleading information. In some cases leaks can even be deadly, for example, when the identities of secret agents gets into the hands of a hostile party.

If overseers are biased, then oversight will tend to advance narrow special interests instead of national or global interests. When bias systematically favors intelligence agencies, as is the case when cooption takes place, then oversight ceases to function in any meaningful way. Rather than provide a check against abuse or the inefficient use of resources, overseers provide a false stamp of legitimacy to the conduct of intelligence agencies. When bias is systematically against intelligence agencies, then legitimate spying will often be rejected. Intelligence agencies will struggle to successfully fulfill their obligations. Often, however, bias will not be systematically in favor or against intelligence agencies. Overseers' biases will be more subtle. They will count some people's interests more than others, or they will discount other people's interests entirely.

Since nearly every individual, to the extent that she has moral and political opinions, is biased, it is virtually impossible to select unbiased overseers. Unbiased (or less biased) overseers are therefore typically created institutionally. Subcommittees in the Congress, for example, are often made up of roughly equal numbers of Republicans and Democrats. But, since these committees are selected from a pool of mostly rich, elderly, white men, even these committees likely have strong biases.

On the criterion of unbiasedness alone, the ideal overseer would be constituted by a representative sample of the population. All interests would be represented in the same proportion as they occur in the population. No interests would receive special

treatment. Of course the representative sample is unlikely to fare well on the other three criteria of a good overseer enumerated above. Nevertheless, representativeness is a good proxy for unbiasedness. More representative groups will tend to be less biased.

Based on the four criteria of the ideal overseer, congressional subcommittees receive a low grade. As discussed above, they typically lack expertise and motivation. Although relatively trustworthy with secrets, occasionally they have political motivations to leak sensitive information. Finally, they have strong biases in favor of protecting privilege and power, and they lack even a moderate level of representativeness: they are significantly less racially, religiously, and ideologically diverse than the population.

Given the gravity of the problems with the legislative oversight of intelligence, one might expect intelligence scholars to give up on legislative oversight altogether, or at least to search for alternative mechanisms to constrain intelligence agencies. But reform recommendations from intelligence scholars have tended to focus not on alternative mechanisms of control but on marginal improvements to intelligence oversight. Johnson (2006, 10), for example, recommends incentives and perks to motivate legislators to conscientiously perform their oversight responsibilities. “Incentives,” he says, “could include prestigious awards presented by the congressional leadership and civic groups to dedicated and accomplished overseers, Capitol Hill perks dispensed by the leadership based on the devotion of lawmakers to accountability, and publicity in national and hometown newspapers underscoring admirable oversight achievements by individual members.” He also recommends clearer jurisdictional lines for oversight responsibilities. Zegart (2012) focuses more on improving the expertise of legislative overseers and bolstering their budgetary power. She recommends abolishing term limits on the House Permanent Select Committee on Intelligence, increasing the sub-

committees' abilities to use congressional resources such as the General Accounting Office, strengthening congressional staffs, and consolidating budgetary power in the House and Senate intelligence committees.

This focus on reforming legislative oversight is strange in two respects. First, reforming legislative oversight may be like making minor repairs to a sinking ship. Even if the reforms turn out to be improvements, they may be improvements to a deeply flawed mechanism of control, and thus may do more to perpetuate ineffective institutions than repair them. Second, existing mechanisms of control are hardly limited to legislative oversight, especially in the domestic context. The courts, as I will discuss in the next section, for example, have a strong hand in domestic intelligence policy.¹⁴ If, at bottom, the concern for liberal democratic societies is preventing intelligence abuses and ensuring that intelligence agencies make efficient use of resources, then there is no reason to think that oversight is the only available mechanism of control, nor is there any reason to think that that legislative subcommittees are the most suitable overseers.

8.2 Judicial Review

So let us turn our attention away from legislative oversight and examine the main mechanism of control for spying within America's borders, judicial review. It is helpful to begin by distinguishing two kinds of judicial review: *ex post* judicial review and *ex ante* judicial review. In both cases, judges review the conduct of intelligence agencies and evaluate whether they have behaved reasonably. The standards indicating what counts as reasonable often shift with circumstances, for example the bar for

¹⁴There are other institutions that play less central roles. For example, the executive branch has inspectors general overseeing many of the intelligence agencies.

reasonable spying in the United States is lower if the target is a non-resident or an agent of a foreign government.¹⁵ But the standards are usually formal. Judges apply them, they do not create them *de novo* for each case.

Ex post judicial review is simply judicial oversight, legislative oversight with a different principal. Judges evaluate the actions of intelligence agencies after the fact. *Ex ante* judicial review, however, is a different mechanism of control. The courts actually share the power to spy with spy agencies, since government agents cannot legally spy without first securing approval from the courts.¹⁶ Typically *ex ante* judicial review of spying is done with warrants. Government agents who wish to spy, in other words, must secure written certification that their spying is reasonable.¹⁷

In the United States, most spying is subject to judicial review. Whether a particular instance of spying requires judicial review depends on whether it counts as a search, since the U.S. Constitution provides individuals protections against unreason-

¹⁵The Foreign Intelligence Surveillance Act sets the standards for spying on foreigners (or agents of foreign governments) within the United States. FISA accelerates the process of getting warrants and it provides less judicial scrutiny over the details of the surveillance. It created a new court – the Foreign Intelligence Surveillance Court (FISC) – which meets secretly to consider applications for warrants by federal law enforcement and intelligence agencies. Importantly, FISA does not require probable cause for all warrants. In the case of U.S. persons, located in the United States, a showing of probable cause of the standard kind is necessary. Indeed, to get a FISA warrant on a U.S. person located in the United States, it must be demonstrated to the FISC that the person is committing a crime related to secret intelligence gathering, terrorism, or identity fraud. But in the case of non-U.S. persons located in the United States, government agencies need not show that a crime is under way; rather they need to show only that the target of surveillance is an agent of a foreign power or a member of a terrorist group. FISA originally required that government agencies certify that the primary purpose of their surveillance is “to obtain foreign intelligence information.” (Foreign Intelligence Surveillance Act, Pub. L. No. 95-511, 104(a)(7)(B), 92 Stat. 1789 (1978)). But later, the Patriot Act weakened this requirement: now it is sufficient that a purpose of the eavesdropping is to obtain foreign intelligence. (50 U.S.C. 1805(a)(3) (2006).)

¹⁶I develop this distinction in more detail in the next chapter.

¹⁷In what follows, I will sometimes call *ex post* judicial review “judicial oversight” and *ex ante* judicial review “judicial warrants.” When I refer to judicial review, I mean to include both kinds.

able search or seizure, but does not mention spying or surveillance directly. The text says:

The right of the people to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures, shall not be violated, and no Warrants shall issue, but upon probable cause, supported by Oath or affirmation, and particularly describing the place to be searched, and the persons or things to be seized.

Some spying in the eyes of the Supreme Court uncontroversially counts as a search, for example when a government official breaks into your house and secretly copies your files; and some spying clearly does not count as a search, for instance when a government official covertly watches you in a public place.¹⁸ But in many cases it is not clear whether spying should count as a search. In *Olmstead v. United States*, for example, the Supreme Court argued that wiretapping did not constitute a search: “There was no searching. There was no seizure. The evidence was secured by the sense of hearing and that only. There was no entry of the houses or offices by the defendants.”¹⁹ But this decision was later reversed in *Katz v. United States*. Now there is a detailed regulatory framework in the United States for wiretapping consistent with the Fourth Amendment.²⁰

So the American regulatory framework does not require judicial review for all government spying within America’s border, it only requires judicial review for gov-

¹⁸Many legal scholars have complained that the court has wrongfully narrowed the scope of the term search. Amar (1994) somewhat comically summarizes fourth amendment jurisprudence as follows, “Warrants are not required – unless they are. All searches and seizures must be grounded in probable cause – but not on Tuesdays.”

¹⁹*Olmstead v. United States* 277 U.S. 438 (1928) at 464.

²⁰See Title III of the Omnibus Crime Control and Safe Streets Act of 1968.

ernment spying that the Supreme Court counts as a search. Further, the text of the Fourth Amendment does not require warrants for government searches, it merely stipulates that searches and seizures be reasonable and that warrants shall not be issued without probable cause. In practice the Fourth amendment has been interpreted to mean that the most intrusive searches, such as searches in a person's home or wiretapping a person's phone require judicial warrants. Less intrusive searches, in contrast, do not require judicial warrants but may be subject to *ex post* scrutiny.²¹

For now, I want to abstract from the intricacies of American law and focus on the general strengths and limitations of judicial review for compelling intelligence agencies to spy ethically. To aid this abstraction, it is helpful to assume that all spying, or at least all spying within America's borders that meets a threshold for potential harm, is subject to judicial review, and not to worry about whether a particular kind of spying counts as a search.

Let us first examine judicial oversight. Judicial oversight, as I mentioned, is similar to legislative oversight, the key difference being that judges do the overseeing, rather than legislators. Judges, of course, have different institutional interests and incentives from legislators, and they usually have a different set of rewards and sanctions available to them. Although judges are far from apolitical, they are typically more insulated politically than legislators, since, for example, they are often not subject to elections. Unlike legislators, judges typically lack the power of the purse, nor do they have the power to create, to destroy, to expand or to contract bureaucratic agencies. Typically, the strongest stick available to judges is the ability to exclude evidence from judicial proceedings. Usually what this means is that if evidence is not gathered according to standards of reasonableness (with e.g. unjustified spying),

²¹See, e.g., Amar (1994).

judges can suppress the evidence thereby barring its use. Judges can also, during courtroom proceedings and opinions, praise and shame intelligence agencies. Finally, the courts can award damages to those who were unlawfully or unreasonably spied on.

While evaluating legislative oversight above I employed five criteria. Four criteria – motivation, expertise, trustworthiness (with secrets), and unbiasedness – I used to evaluate legislators as principals. The fifth criteria – the extent to which the principal has unmediated access to information – I used to evaluate oversight as a mechanism of control. It is helpful to begin with the same criteria to evaluate judicial oversight.

On the criteria of trustworthiness with secrets, motivation, and expertise, judicial overseers grade marginally better than legislative overseers. Judges have fewer political reasons to leak secrets since they are more isolated from the political process than legislators. They also have fewer disincentives to spend time on intelligence issues and thus to develop intelligence expertise, since they have little if any pressure to please constituents, who demand expertise on more local issues. But unless a judge's tenure is for life, she still has to worry about those who make and influence her appointment. Secrets could be used as currency to buy favor from these parties. Judges also typically see a wide variety of cases, and intelligence issues are likely to figure in relatively few of them. Hence judges typically do not have strong incentives to focus on intelligence issues or to develop intelligence expertise.

On the criterion of unbiasedness, however, judges grade slightly worse than legislators. Motions to suppress evidence and cases seeking damages are often heard by just one judge. Further, it is often pointed out that judges are biased by the fact that evidence demonstrating or supporting the guilt of the defendant has been found, otherwise there would be no hearing to suppress. It is more difficult for a judge to conclude that a search was unreasonable, when the officer's hunch paid off. (Stuntz,

1991, 911-913)

Finally, judges have no better access to information from law enforcement officials and intelligence agents than legislators do. Law enforcement officials and intelligence agents can withhold information and provide false information, and judges are rarely in the position to gauge whether they are being deceived or kept in the dark.

Reaching beyond the five criteria employed above, there are a few more problematic features of judicial oversight. First and most importantly, when the primary stick at the disposal of the courts is the suppression of evidence, law enforcement and intelligence agents are almost never sanctioned for unreasonable spying when their hunches do not pay off. If police officers unreasonably bug a suspect's house, for example, but never find evidence of her guilt, the courts have very little power to reprimand the officers' conduct. Similarly, if law enforcement or intelligence officers wish to control or coerce a person in ways other than arresting her, the threat of evidence suppression is not a deterrent. Recall the case, discussed in Chapter 5, of FBI officials threatening to expose Martin Luther King Jr.'s marital infidelities. Those officials wished to bring King to ruin, and they were willing to use nearly all available means to achieve their ends.

The threat of damages is also a meager deterrent for abusive spying. Since most spying is no doubt never discovered, the probability of being sued for damages is exceedingly low. Further, the harms from wrongful covert surveillance are notoriously difficult to quantify. (Stuntz, 1991, 901) Recall some of the potential harms of spying canvassed in chapter 5. An illegal search can, for example, signal broadly to members of a community that their personal information is less secure than they thought. Or it can be demeaning if it is performed arbitrarily on a disadvantaged group. These harms, although they cannot be ignored in a liberal society, are almost impossible to quantify in dollar terms, and further since class action lawsuits are rare (even nonex-

istent) for spying, the tendency is for the legal system to considerably underestimate these harms.

As a mechanism of control, judicial oversight fares no better than legislative oversight, then. However, it still may play an important role in a complex set of institutions designed to control intelligence agencies. Certainly, in some circumstances, the threat of evidence suppression or censure by the courts is enough to prevent government agents from abusing their power to spy.

In contrast to judicial oversight, judicial warrants, I want to argue, are a much more effective mechanism of control than legislative oversight – at least for the day-to-day decisions that intelligence agencies make to spy. To make this point, I will again draw on the five criteria I used above. It may seem repetitive to evaluate judges a second time, but the evaluations for four of the criteria actually change in interesting ways when the judicial review is done *ex ante*.

With legislative oversight, the crux of the motivational problem is that legislators have strong political incentives to spend their time and energy on issue areas that are more likely to generate electoral benefits than intelligence, and there is nothing about the mechanism of oversight that compels them to give intelligence their undivided attention. Both of these problems are less pronounced with judicial warrants. With judicial warrants, each request to spy receives scrutiny, making motivation less of an issue. Certainly the courts can do a half-hearted job considering requests, but they cannot choose *not* to consider requests. As I mentioned above, judges also tend to have fewer competing commitments. Unlike legislators, they are not tempted to spend their time on issues more likely to produce electoral advantages.

Hence on the criterion of motivation, judicial warrants fare better than legislative oversight. Judicial warrants also fare better on the criterion of expertise for the simple reason that judges are compelled by the institution of judicial review to tend to

each case. Further, for panels of judges selected exclusively to review spying requests, such as the Foreign Intelligence Surveillance Court, the expertise advantage can multiply, since appointees can be selected for their experience in intelligence or national security.

On the criterion of unbiasedness, judges also perform better when their review is done *ex ante* than *ex post*. During *ex post* review, as I explained above, it is uncommon for judges to find spying unreasonable when the spying produces damning evidence against the defendant. Such bias obviously cannot play a role during deliberation over whether to grant a warrant, however, since the judge will not know whether the spying produces damning evidence or not.

Still, judges are far from unbiased principles. Even when it is a panel of judges reviewing requests rather than a single justice, these panels are not representative of the population. For one, panels tend to be small (usually eleven or fewer), so it is difficult to make them truly representative of the population. More importantly, however, those who appoint judicial panels have incentives to ensure that judges are consistent with their political values. Hence panels of judges will tend to be systematically biased toward the political views of those who appoint them.

So judicial review runs into the same problem as legislative oversight, even when review is done *ex ante*: there is no guarantee that the court's bias will not sometimes steer agencies away from ethical principles rather than toward them. Further, judicial warrants suffer from the informational problems that plague legislative oversight. In one respect the information problem is less severe for judicial warrants than it is for legislative oversight. Legislative overseers, who typically evaluate countless intelligence decisions *ex post*, struggle to know what information is relevant. Which decisions should be scrutinized? What information is required to evaluate these decisions? Judges, in contrast, evaluate particular decisions *ex ante* when it is much

clearer what kind of information is required. But information asymmetries remain a deep problem for judicial warrants. Judges rely on law enforcement and intelligence agents to provide them with the evidence to make good judgments. Classification exacerbates these asymmetries, since when law enforcement or intelligence agents fail to pass relevant information to judges, the judges cannot rely on concerned citizens or interest groups to alert them to this fact.

Hence as a mechanism of control, judicial warrants fare slightly better on the five criteria than legislative or judicial oversight. But like legislative and judicial oversight it suffers from information asymmetries, and when it *is* a successful mechanism of control there is no guarantee that it will compel intelligence agencies to behave ethically.

I want to suggest that judicial warrants have two further weaknesses. First, compared to legislative or judicial oversight, a system of judicial warrants is expensive; and second the request-by-request nature of a system of judicial warrants makes it a poor tool for addressing strategic problems facing intelligence agencies.

The claim that judicial warrants are relatively expensive should not be surprising, since with *ex ante* review every single instance of spying receives scrutiny, whereas with *ex post* review typically only a small sample of instances receive scrutiny. As Stuntz (1991, 887) argues, “Ex post enforcement of conduct rules generates vast economies, since it allows the system to achieve a given level of deterrence while litigating only a small fraction of the possible violation.” These economies, he explains, are possible for two reasons: most actions tend to be compliant with the law, and many violations tend to do only trivial harm.

The primary cost of judicial warrants is labor. Since *ex ante* review scrutinizes many more instances of spying, it requires many more judge/hours for scrutiny. A consequence of these high resource requirements, Stuntz points out, is that warrants

tend to be granted with very little deliberation: most search warrants are granted *ex parte* in just a few minutes. The process is typically much less thorough than motions for evidence suppression or hearings for awarding damages.

So, *ex ante* judicial review may be a stronger mechanism of control than legislative or judicial oversight, but it also considerably more costly. Perhaps the most important weakness of judicial warrants, however, is the reactive position in which they place the courts. Judges are tasked with reviewing the requests to spy that come before them. They are not positioned to make broader, more strategic evaluations of the performance of intelligence agencies.

One of the virtues of oversight is that it permits judges or legislators to ask broad questions. Are intelligence agencies identifying and responding to the most dangerous and most exigent threats? Do intelligence agencies use their resources efficiently? Should more resources be placed in human intelligence and less in signals intelligence (or vice versa)? Do intelligence agencies exhibit a culture of compliance, or is there a tendency for them to behave like “rogue elephants”? Could the money spent on intelligence do more good if it were spent on education, health care, or other programs? *Ex ante* review procedures typically do not permit principals to engage in this kind of broad strategic evaluation.

8.3 Conclusion

My aim in this chapter was to examine and critique the two primary mechanisms employed in the United States to control the conduct of intelligence agencies. Both legislative oversight and judicial review, I argued, are by themselves insufficient mechanisms of control. Further, both mechanisms exhibit systematic biases, suggesting that they will sometimes compel intelligence agencies to follow narrow or partisan

aims, rather than to secure common interests.

These conclusions do not suggest that the aim of compelling intelligence agencies to behave ethically is hopeless, however. There may be creative possibilities to combine institutions to achieve better outcomes than legislative oversight or judicial review alone. Exploring these possibilities is the purpose of the next chapter.

Chapter 9

Reforming the Control of Government Spies

In this final chapter my primary aim is to go beyond legislative oversight and judicial review and consider the full range of mechanisms of control that could be brought to bear to constrain intelligence agencies. I also have a secondary aim to begin the conversation about how to reform the institutions that control America's intelligence agencies.

I argue that there are five principal mechanisms of control that can be employed to constrain the conduct of intelligence agencies. Three of these mechanisms – accountability mechanisms, mechanisms of direct interference, and information alteration mechanisms – are agentive, that is they rely on a principal deliberately shaping an agent's options or her perceptions of those options. The other two – selection mechanisms and partitioning mechanisms – are non-agentive.

All five mechanisms of control can be employed with some success, but it is partitioning mechanisms, I suggest, that offer the most unexplored potential for constraining the conduct of intelligence agencies.

The potential of partitioning mechanisms is borne out in my reform proposals. I propose that all spying above some threshold for potential harmfulness should be subjected to *ex ante* review. Although it may seem natural for the courts to play the role of reviewer in such a system, I argue that the courts are not the principal best suited for this role. Elected panels are likely to be less biased than courts and they are just as likely to have the necessary expertise. I further propose that the procedures for review should be reformed. In particular, I argue that they should incorporate a devil's advocate.

But, as I argued in the previous chapter, *ex ante* review is by itself an insufficient mechanism of control, since those who review day-to-day requests are often poorly positioned to perform strategic oversight. Hence I recommend that the size of the elected panel be expanded, so that panelists can rotate between day-to-day and strategic oversight responsibilities. I also recommend that the panel be given a host of further powers and administrative capabilities to ensure that they have full access to the information required for them to perform their role as strategic overseer.

The chapter begins by distinguishing the five abovementioned mechanisms of control. I first unpack oversight into its two component mechanisms and then introduce the remaining three mechanisms of control. As I describe the five mechanisms, I also indicate when they are likely to be effective, and to what extent they are likely to be effective in the intelligence environment. The chapter's second section develops my proposal for reform. I begin by arguing that pairing judicial review and legislative oversight for all spying over some threshold of potential harm would be an improvement on existing American institutions. But such a proposal leaves a number of problems with these mechanisms of control identified in the previous chapter unresolved. By addressing these problems, I work toward my final proposal of the elected panel.

9.1 Mechanisms of Control

In order to illuminate the universe of possible mechanisms to control intelligence agencies, it is helpful to begin with oversight. Oversight, I want to suggest, is primarily an accountability mechanism. It is not, however, exclusively an accountability mechanism. It is, rather, a bundle of control mechanisms, the most important of which is an accountability mechanism.

“Accountability” is a term with many senses, and it is easy to conflate these senses and fall into confusion. Two senses are particularly important in political theory. The first sense is accountability as *responsiveness*. Agent A is accountable to agent B for X to the extent that A’s decisions in X are responsive to B’s interests, preferences, etc.¹ Accountability, in this sense, is often thought to be a virtue in democratic theory, at least when the accountable agent is responsive to the interests of all those affected by her decisions. According to this sense, any institution that fosters A’s responsiveness to B’s interests is an accountability institution. All mechanisms of control could therefore potentially be accountability institutions.

But this usage is confusing, since there is a second *institutional* sense of accountability. Scholars disagree about the precise definition of this institutional sense.² The most common formulation holds that A is accountable to B for X when B has the power to reward or sanction A for her performance of X.³ On this formulation accountability institutions are concerned principally with *ex post* rewards and sanctions. When accountability relationships “work” B compels A to perform X according to

¹See e.g. Dovi (2007, 66).

²It may be that some of the mechanisms of control I discuss below are sometimes ignored or overlooked by democratic theorists because calls for accountability (as responsiveness) are mistakenly interpreted as calls for institutional accountability.

³See e.g. Fearon (1999); Grant and Keohane (2005).

her (B's) standards with the threat of punishment or the promise of reward. Accountability institutions thus rely mostly on anticipatory power, what political scientists sometimes refer to as power's second face.⁴

A competing formulation, defended most notably by Philp (2009), maintains that A is accountable to B for X when B has the power to compel A to explain or justify her performance of X. On this version, institutions of accountability are primarily about B "giving an account," that is explaining or justifying why she did X the way she did. Although giving an account may itself be a kind of punishment for some agents, anticipatory power, at first glance, plays a less central role in this conception, especially when B does not have any further power to reward or sanction A for her behavior in X. On closer examination, however, in many cases there are other agents who have the power to reward or punish A for her performance of X, but who lack the capacity to compel A to explain or justify her behavior. A, in these cases, may anticipate those agents' responses when she expects that B will compel her to explain or justify her performance of X.⁵

A third formulation holds that accountability is a combination of the first two formulations. Mansbridge (2009) and Rehfeld (2005, 189), for example, argue that accountability relationships have both deliberative and sanctioning "elements" or "dimensions."⁶ In actual accountability relationships, one or the other elements may be prominent, but accountability relationships typically have both elements.

For the rest of this chapter when I use the term "accountability," I mean in-

⁴See Lukes (2004); Hayward (2000).

⁵And she expects that the account she gives to B is shared with others. When the powers to compel A to explain or justify her performance of X and to reward or sanction A for her performance of X are possessed by different agents, we have what Rubenstein (2007) calls surrogate accountability.

⁶Rehfeld suggests that it is a conceptual fact that accountability relations have both of these elements. Mansbridge, on the other hand, suggests that the fact is merely an empirical one.

stitutional accountability. My own view about institutional accountability is that Mansbridge and Rehfeld are probably correct: most accountability relationships exhibit both deliberative and sanctioning elements. But I am ambivalent about whether either or both of the sanctioning and deliberative elements are necessary or sufficient components of institutional accountability.

Oversight as it is practiced in the United States and in many other countries includes both deliberative and sanctioning elements of accountability. In the case of deliberative accountability, legislative overseers have the power to call hearings in which they can compel intelligence officials to explain and to justify their conduct. Similarly, they also have the power to initiate and carry out investigations, and with the power to call investigations typically comes the power of subpoena, the power to compel testimony or the release of records. The sanctions available to legislators include publicly shaming agency personnel, cutting or even eliminating agency budgets, enacting costly procedural requirements, and denying future appointments.

Yet oversight is not just an accountability mechanism. Accountability institutions, I have argued, function primarily by bringing to bear anticipatory power: agencies condition their behavior because of the threat of sanction or the promise of rewards. But some of the typical oversight powers available to legislators can be employed not just *ex post* but also *ex ante*. For example, legislators can use their power of the purse not just to punish (or reward) agencies for their bad (good) conduct, but also to rule out an agency's options or to favor some options over others. Similarly, legislators can use their powers to make and change statutes to remove an agency's authority to pursue certain options, or to eliminate the agency altogether. (MacDonald, 2010) Hence oversight is not just an accountability mechanism, it is also a mechanism of direct interference.

Oversight is likely to be successful when the conditions of effectiveness are met

for both accountability mechanisms and mechanisms of direct control. Accountability mechanisms are more effective when the accountability holder has rewards and sanctions weighty enough to motivate the accountable agent, when the accountability holder has more complete and unmediated access to information about the accountable agent's conduct, when the accountability holder is motivated to conscientiously perform her responsibilities, and when it is clear to the accountable agent how she is supposed to behave to avoid sanctions and earn rewards.

Mechanisms of direct control are more effective when the controlling agent can remove or alter any of the options available to the agent she seeks to control, when the options the controlling agent can remove or alter are ends rather than means, when the controlling agent can accurately predict which options the agent she seeks to control will select, and when the controlling agent is motivated to conscientiously perform her responsibilities.

As I argued in the previous chapter, many of these conditions of effectiveness for oversight are not met in the intelligence environment.

Accountability mechanisms and mechanisms of direct interference are both agentive control mechanisms, that is they rely on a principal deliberately shaping an agent's options or her perceptions of those options. A third agentive control mechanism, not often associated with oversight, is the power to alter an agent's information environment. Since agents do not have direct access to reality, but rather beliefs based on the information available to them, and since agents' conduct is determined, in part, based on their beliefs about the set of options facing them and their beliefs about the likelihoods of the various costs and benefits following from those options, one way of altering an agent's conduct is by altering her information environment

thereby altering her beliefs.⁷

A wide range of tactics could potentially alter an agent's information environment, such as persuasion, manipulation, and brainwashing. I can't go into these tactics here, but I want to give an example of how altering agents' information environments can shape their conduct, and then speculate about how the example could apply to institutions in intelligence. In his (2012) book *The Honest Truth about Dishonesty*, psychologist Dan Ariely recounts an experiment he ran with Nina Mazar and On Amir.⁸ In the experiment two groups of students were given twenty matrices and tasked with finding two numbers in each matrix that added up to ten. Students had five minutes and were asked to "solve" as many matrices as possible. Further they were told they would receive fifty cents for every correct answer.

In previous similar experiments the researchers had the first group verify every correct answer with a proctor, while they had the second group count up their correct answers, shred their worksheets, and then report their number of correct answers. In so doing, they learned whether and how much students would cheat if they thought they could get away with it. On average, participants in the shredder condition claimed to have solved two more matrices than those who had to check their work.

In the new experiment what Ariely *et al.* wanted to know was whether priming participants with moral or ethical reminders would make them less likely to cheat. It did. Of the students placed in the shredder condition, half were asked to sign their University's honor code. Those who did not sign the honor code cheated to a degree similar with past results. But those who did sign the honor code appeared not to

⁷It is possible to alter a person's beliefs not by altering her information environment but by altering her physical states, for example by putting a powerful magnet to her head. I shall ignore these cases here for simplicity.

⁸Amir, Ariely and Mazar (2008).

have cheated at all. Ariely concluded, “it seems when we are reminded of ethical standards, we behave more honorably.” (43)

There are numerous ways that intelligence agents could be primed like the students in Ariely’s experiments in order to prevent them from engaging in illicit or ill-advised spying. Warrants could be fashioned to call to mind moral and ethical beliefs. They could, for example, compel agents to swear that their proposal meets constitutional standards, and that all information is true and complete. Agency signs and slogans could be designed to remind employees of their duties to respect individual rights, to protect innocents, etc. Steps could also be taken to give intelligence agents the sense that their activities are under observation by supervisors, regulators, or overseers. Although each of these interventions would likely alter the behavior of intelligence agents in only minor ways, cumulatively, they could make concrete strides toward producing a culture of compliance.

Generally, mechanisms that change agents’ information environments are likely to be effective under the following set of circumstances. First, the principal is more likely to be able to control the agent’s conduct by changing her information environment when the agent’s views are not hardened. When people’s views are strongly held, most recent social science suggests that evidence confirming their view will tend to be accepted, while disconfirming evidence will tend to be simply tossed away and ignored.⁹ Second, the principal is more likely to be able to control the agent’s conduct when she knows the agent’s views. By knowing the agent’s views, the principal can avoid costly appeals unlikely to sway the agent. Finally, the principal is more likely

⁹A classic study is Lord, Ross and Lepper (1979), which found that “People who hold strong opinions on complex social issues are likely to examine relevant empirical evidence in a biased manner. They are apt to accept “confirming” evidence at face value while subjecting “disconfirming” evidence to critical evaluation, and as a result to draw undue support for their initial positions from mixed or random empirical findings” (2008)

to be able to control the agent's conduct when the number of others who can change the agent's information environment is small. When the agent is bombarded by information from multiple principals, the power of any of these particular principals to affect her choices is attenuated.

Given that none of these conditions are met to a high degree in intelligence, it may not be possible simply by changing the information environment to counteract particular risks of abuse. But it still may be possible, over time, to instill or reinforce values that will favor sound deliberation and tamp down the tendency to engage in reckless or self-regarding conduct. Some may be skeptical that those who specialize in manipulation and deception could be so easily primed, and certainly intelligence agents make for a limit case, but there is a growing literature in psychology that suggests priming matters more than people think.

So far we have seen that there are three mechanisms of agentive control: mechanisms of direct interference, accountability mechanisms, and information alteration mechanisms. Not all mechanisms of control rely on a principal interfering or threatening to interfere with another agent's options or her perceptions of her options, however. Certain outcomes can be made more likely by *selecting* agents who are more likely to behave in particular ways or by *partitioning power* to create a structure of decision making that generates predictable results.

In a recent article, Mansbridge (2009) helpfully clarifies the distinction between accountability mechanisms and selection mechanisms. She says.

This [selection] model works only when a potential agent already has self-motivated, exogenous reasons for doing what the principal wants. The principal and agent thus have similar objectives even in the absence of the principals sanctions. As a general rule, the higher the *ex ante* probability

that the objectives of principal and agent will be aligned, the more efficient it is for the principal to invest resources *ex ante* in selecting the required type rather than investing *ex post* in monitoring and sanctioning. (369)

So, as a mechanism of control, selection does not work by threatening sanctions or promising rewards, it works by choosing an agent who is likely (because of her own internal motivations) to behave in a particular way.

Selection is likely to be an effective mechanism of control when there is a diversity of candidates on relevant characteristics. Since it is possible but unlikely that in a small pool of candidates the selector will find her ideal candidate, the selector is more likely to find a desirable candidate when the pool of potential candidates is diverse. Diversity is particularly important when the candidate selected is empowered across a wide range of issues. If she faces hundreds or thousands of possible decisions, then selection will tend to be less effective, unless the pool of candidates is extraordinarily diverse.

Selection is also more likely to be effective when some of the more desirable candidates display a good deal of constancy. If all of the candidates display extreme fluctuations in their principles, policy positions, etc. it complicates and diminishes the quality of prediction for the selector and thereby tends to reduce the effectiveness of selection. Finally, the effectiveness of selection depends on the ability of the selector to make reliable predictions about how well candidates will conform to her standards of good conduct. Good prediction requires not only knowledge of probability and statistics, patience, and sometimes a good deal of costly work, but it also requires trustworthy information about the beliefs and behavior of candidates.

Selection already plays a role in constraining the conduct of intelligence agencies. Anyone hired into the American intelligence community, for example, has to survive

rather thorough background checks. These background checks presumably weed out a host of unsavory characters – especially criminals and foreign agents – who would be more likely than average to abuse their powers to spy. Further, the directors of many of America’s intelligence agencies are selected by the President.

There seems to be some promise to selection as a constraint on the power of spies, then. But the extent to which selection can be an effective constraint on power is rather limited. Principals rarely have complete information for the selection of their agents, and since many of the agents who operate in secretive environments specialize in concealment, it is naive to think that principals could obtain a complete file on potential candidates. Principals also routinely face selection dilemmas. The agents who are the best executives, managers, scientists, and spies, will often not be those most likely to respond to the principal’s interests. Principals are therefore forced to balance competence and responsiveness. Finally, even if the principal had perfect information and she could identify an agent who, based on historical experience, is both competent and responsive, it is dubious whether historical experience will lead to an accurate prediction, since secrecy can have a treatment effect. Conduct in an open environment is not necessarily indicative of conduct in a secretive environment. Evidence that an agent behaves well in a strongly constrained environment, no matter how extensive, is thus not sufficient to infer that the agent will continue to behave well in a less constrained environment.

Less well understood than the distinction between accountability mechanisms and selection mechanisms is the distinction between accountability mechanisms and mechanisms that partition power. To elucidate this distinction it is helpful to first illustrate the range of ways that power can be partitioned. Imagine some agent, Rex, who has the power to govern the territory Elysium. Rex has absolute power: only he can make, interpret and enforce Elysium’s laws. Now imagine that, perhaps for the good

of his kingdom, Rex decides to cede some of his power. How could Rex's power be divided?

It could be divided *functionally*. Rex could cede, for example, his power to make the laws, but keep his powers to enforce and interpret them. It could be divided *spatially*. For instance, Rex could keep all of his governing powers in one corner of Elysium, but cede his powers in the rest of the kingdom. It could be divided *temporally*: Rex could cede his power during even years, but keep it during odd ones. Finally, it could be divided *topically*. Rex could keep his power to determine Elysium's security policy, but cede other powers, such as the power to tax.

These methods of dividing power should be distinguished from sharing power. Rather than dividing up the range of his power, Rex could, by sharing his power, make authoritative decisions depend not just on his own will but also on the will of another agent. Rex would thus have to cooperate with another agent to govern Elysium.

Of course, these are not the only five options for Rex to divide or share his power because any of these methods could be combined. Rex could, for example, cede his legislative power on odd years in a tiny corner of Elysium, or cede his judicial power and share his legislative and executive powers. Since topics and functions can be defined in many different ways and time, and space can be sliced up into infinitesimally small parts, the possibilities for dividing and sharing power are endless.

Historically, many of these methods for dividing or sharing political power have gone by different names. For example, dividing power functionally is usually referred to as the separation of powers.¹⁰ Carving off a set of issues and ceding them to con-

¹⁰See e.g. Cooter (2002, Ch. 9), Gwyn (1965), and Vile (1998). The separation of powers is sometimes confused with mixed government. While the two do share features, the former focuses primarily on functional division, while the latter concerns itself chiefly with involving different parts or classes (e.g. the aristocracy, the people, etc.) of society in government. Notice that mixed

stituent (usually territorial) political units is often called federalism.¹¹ Term limits are one kind of temporal division of power, so too are sunshine provisions. When power is shared within a governmental function (or perhaps within a topic) it is sometimes said that one agent “checks” or “balances” the other. But usage of these labels has not always been consistent. For example, the separation of powers is sometimes said to be distinct from checks and balances; other times the two ideas are thought to be indistinguishable.¹²

Partitioning power, that is dividing and/or sharing power among agents, I have suggested, is analytically distinct from other mechanisms of control, and especially mechanisms of accountability. But in one sense, it seems the two cannot be distinct: agentive mechanisms of control, such as accountability, presuppose the partitioning of power. One agent has the power to act and another agent has the power to reward, sanction, or compel the explanation or justification of these actions. But the partitioning that I am concerned with is not political power generally, but rather partitioning a *particular* power, in this case the power to spy.

Partitioning, in this more specific sense, is clearly distinct from accountability (and other mechanisms of agentive control). Power can be partitioned and accountable, as in the case where two intelligence agencies are accountable to a congressional sub-committee. It can be partitioned but not accountable, as was more or less the case with intelligence agencies prior to the creation of the standing intelligence committees in the House and Senate. It can be accountable but not partitioned, if, for

government could be achieved with the separation of powers, but it could also be achieved topically, temporally, or spatially.

¹¹See e.g. Karmis and Norman (2005).

¹²Finer (1949, 84) treats the two as identical. But Madison argues in Federalist 48 and 51 that checks and balances are necessary for the maintenance of a separation of powers.

example, all intelligence functions were consolidated in one agency that was accountable to a congressional subcommittee. Finally, it could be neither accountable nor partitioned, as was more or less the case with the KGB before the fall of the Soviet Union.

Partitioning power is likely to be an effective mechanism of control when power is partitioned among agents powerful enough to deter one another from (re)consolidating power or when there is another more powerful agent who will enforce the separation. All agents need not have roughly equal power though: the usual balance of power logic applies.¹³ When one agent overreaches or attempts to overreach, multiple agents can band together to check the overreach. But because banding together is not costless, the presence of a relatively powerful agent will tend to mean that some overreach will go unchecked. When no combination of agents is powerful enough to check a powerful agent from overreaching, dividing and/or sharing power is unlikely to be an effective constraint on power.

The effectiveness of partitioning power also depends on the motives of the agents among whom the power is partitioned. Agents with nearly identical motives are unlikely to check one another. They will tend to behave monolithically. At the other extreme, when agents have no overlapping motives, power will tend not to be exercised at all. Gridlock will ensue. Hence if the aim is to prohibit certain exercises of power, then the best institutional solution may be to simply divide this power among agents never likely to agree on how it should be exercised. But if the aim is that power be exercised in particular ways, then, designing effective institutions becomes a sophisticated game theoretical exercise of placing agents with the particular set of motivations in a particular set of institutions.

¹³See, e.g., Waltz (2001).

In the United States, the power to spy is divided in a number of ways. Sixteen agencies alone conduct intelligence activities. Compared to the many divisions of the power to spy, however, there is relatively little power sharing. Agencies, for the most part, have independent powers to spy, although in the domestic context these powers are often shared with courts, as warrants must typically be sought to spy on citizens. Inter-agency bargaining is typically not necessary to authorize spying.

The potential of partitioning power for constraining the power of spies is promising, but the full potential remains unclear. It is probably too optimistic to think that institutional arrangements dividing and sharing power could ever be so finely tuned that they could guarantee that power is wielded more or less in accordance with a fairly complex system of rules. But it is also unlikely that conformity with fairly complex rules could ever be achieved in severely asymmetric information environments without partitioning power.

A set of institutions with fewer divisions of power but more power sharing would likely better constrain spy agencies than the existing American model. Fewer divisions would likely be better because there is little reason to think that the gains to constraining power increase to any considerable degree after power is divided four or five times, and there are strong reasons to think redundancies increase the likelihood of abuse. Economists who study industrial organization, for example, distinguish monopolies, duopolies, and oligopolies, but after the number of firms in an industry reaches four, something resembling a perfectly competitive market tends to obtain.

To say there is little to be gained by further dividing power among America's intelligence agencies is not, however, to say that power is divided optimally. Most of America's intelligence agencies are "full source" intelligence agencies. Each has diverse collection capabilities, for example from human sources, from signals etc., and each has its own analytic capabilities. Further, most intelligence agencies work across

a range of threats. Plausibly power could be divided among America's agencies differently to capture more specialization benefits, to enhance data sharing, or to curb inefficiency or abuse.

More power sharing is likely to lead to better outcomes because under the current institutional configuration, large swathes of spying, especially foreign spying, can still be carried out at the discretion of the intelligence agencies. No cooperation is required from other branches of the government.

I have now highlighted five mechanisms of control: accountability mechanisms, mechanisms of direct interference, information alteration mechanisms, selection mechanisms, and mechanisms that divide or share power. One may be tempted to tick through the conditions of effectiveness for each mechanism of control and conclude that one mechanism in particular is the best mechanism of control for the intelligence environment, but such a strategy is wrongheaded, since there is no reason that the mechanisms of control must be used in isolation, and since there will often be considerable interaction effects between mechanisms.

Another way to summarize the five mechanisms is with a list of choices for institutional design. Suppose we begin with a political power, such as the power to spy, and we want to design a set of institutions to ensure that this power is exercised legitimately. We can think of ourselves as Rousseau's (2002, Book II, Ch. 7) legislator, designing institutions for posterity. Our first choice is whether to leave the power to spy whole and give it to one agency, or to partition it among multiple agencies. We could, for example, give the power to spy at home to one agency and the power to spy abroad to another. We also have to choose whether the power to spy, once it is divided, should be shared. Should the courts, for example, have a veto on intelligence

agencies' decisions to spy? Should other arms of the bureaucracy have a say? For example, perhaps the Department of State should sign off on requests to spy.

Second, we must choose how to select the personnel to populate the agencies to which we have delegated power. The options here are limitless. We could choose one leader, for example, and let her appoint her subordinates, successors, etc. Or, we could mandate periodic appointments.

Designing selection institutions includes not just a choice about *how* to select agency personnel, it also includes a choice about *who* selects agency personnel. Agency personnel could, for example, be selected by existing branches of the government, by the people at large (or sorted into territories, etc.), panels of experts, etc.

Even after we partition the power to spy and select good agents to exercise these powers, we still might not be satisfied that the power to spy will be exercised justly. Consequently, we might choose an overseer or a set of overseers. Choosing an overseer involves not only determining who should oversee the agencies, but also what powers they have to control the agencies' conduct. By granting the overseer powers, we can incorporate any of the agentive mechanisms of control discussed above. By granting the overseer the power to reward and sanction the agencies, for example, we could establish an accountability mechanism.

Finally, we may want to add ancillary institutions (what are sometimes called "administrative procedures" in the bureaucracy literature) to ensure the proper functioning of the mechanisms of control. For example, since accountability mechanisms do not function well when the overseer does not have visibility into the conduct of the agency, we may set up reporting requirements, sunshine laws, or subpoena procedures.

To summarize, we have six choices when designing institutions to constrain government spies. We have to choose whether (and how) to:

1. Partition the power to spy among agencies (courts, legislative subcommittees, etc.) {A1, A2, . . . ,An};
2. Appoint a selector of agency personnel;
3. Set procedures for selecting agency personnel;
4. Appoint an overseer;
5. Grant powers to overseer(s) (to control agencies), and set procedures for agency control; and
6. Set ancillaries/administrative procedures.

9.2 Reforming American Mechanisms of Control

With the institutional choices available to policymakers to control intelligence agencies laid out more or less clearly let us turn our attention to how policymakers ought to make these choices in the American context.

I proceed by first examining the idea of marrying *ex ante* judicial review with legislative oversight for all spying – domestic *and* foreign – that meets some threshold of potential harm, and explain why such an institutional configuration would be an improvement on the American model. I then argue that this pairing of judicial warrants with legislative oversight leaves unsolved a number of problems identified in the previous chapter with these institutions. For the rest of the section, I propose modifications that I believe address these problems without raising problems of comparable significance.

This method is useful because it leads relatively quickly to concrete reform proposals that improve on existing institutional forms. It is important to stress, however,

that it is *not* a comprehensive method: I do not systematically answer the six questions I posed above. I do not consider, for example, whether there should be fewer American intelligence agencies, or whether the divisions among America's intelligence agencies should be made differently. Hence, my proposal should be taken as a starting point for a conversation about intelligence reform, a conversation that I hope over time as additional proposals join the fray will become more comprehensive.

Let me begin by defending the claim that marrying judicial warrants with legislative oversight for all spying – domestic and foreign – that meets some threshold of potential harm would be an improvement on existing American mechanisms of control. First, under current American institutions, large categories of potentially harmful spying are subject only to *ex post* legislative scrutiny, and as we saw in the previous chapter, oversight by itself is an insufficient mechanism of control in the intelligence environment. Spying conducted outside of America's borders requires no judicial review, for example. Similarly, many kinds of intrusive domestic spying are subject only to *ex post* judicial or legislative review. For example, a large amount of spying in the United States is done with national security letters (NSLs). NSLs permit intelligence agencies to collect the financial transactions and the phone, internet, or email records of individuals, so long as the information is relevant to an investigation either of terrorism or clandestine intelligence activity. Since 2007, nearly 15,000 NSLs requests have been issued each year.¹⁴ NSLs do not require *ex ante* judicial review, though the letters can be challenged in the courts by the recipients. Requiring judicial warrants for all spying that meets some threshold for potential harm would thus reduce the likelihood of unreasonable or abusive spying.

A second reason pairing legislative oversight and judicial warrants for all spying

¹⁴epic.org/privacy/wiretap/stats/

under a threshold of potential harm would be an improvement on America's current institutions is that the extension of judicial review would inform the legislative oversight process. Judicial oversight places volumes of information on the record for use by those on the congressional subcommittees with oversight responsibilities. Information produced by judicial review would permit legislators and their staffs to examine more systematically what kinds of operations intelligence agencies are involved in, how well resources are employed, and how well the agencies predict the costs and benefits of their operations. In short, the panels would permit legislative overseers to better evaluate the quality of their intelligence agencies.

One worry many might have about pairing judicial warrants and legislative oversight is that the costs of administering such a regime would be considerably larger than current institutions. America engages in thousands of spy operations abroad and in thousands of spy operations domestically that currently do not require judicial warrants. Mandating judicial review for these operations would require thousands of hours from judges to review warrant requests and thousands of hours from intelligence and law enforcement agents to fill out these warrant requests. A further worry is that these administrative costs will sometimes deter justified spying, or they will unnecessarily delay spy operations until they are either less effective or no longer possible.

These worries about costs are important, but I want to suggest that the costs will be outweighed by the benefits of better agency control. When agencies spy according to the ethical principles I argued for in previous chapters, they focus efficiently on preventing the gravest harms. In contrast, when they abuse their powers to spy, they tend to follow personal and political interests doing less for public security or welfare. The opportunity cost, in other words, of an unconstrained spy agency is considerable, since some of the gravest threats are not prevented. Beyond the opportunity costs, there are the many potential harms that I enumerated in chapter five, all of which

balloon when spying is not properly constrained.

The costs of judicial warrants can also be minimized with institutional innovation. For example, when the time required to apply for a warrant is likely to make justified spying impossible or considerably less effective, procedures can be put into place for warrant applications *ex post*. Similarly, warrants need not be required for every instance of spying. Procedures can be designed so that law enforcement and intelligence agents, in certain circumstances, can apply for and receive framework warrants, that is warrants that permit them to use a predetermined number of different surveillance tactics, on a group of individuals, in a set of places, over a defined amount of time. There is no reason, for example, to require a warrant every time an intelligence agency wants to tap the phone, or track the whereabouts of a known al-Qaeda agent. A framework warrant could provide standing permissions to spy (with stipulated means, in stated places, etc.) on such an agent. Finally, the procedures for filing warrants can be greatly simplified with technology. Warrants, for example, can be applied for via phone or email.

Marrying judicial warrants with legislative oversight for all spying above some threshold for potential harm would be an improvement over American institutions, then. Yet a number of the problems with legislative oversight and judicial review identified in the previous chapter are not solved by this marriage. Bias, motivation, expertise, and information asymmetries all remain issues. Legislative subcommittees are typically unmotivated, unrepresentative, and inexpert. Judges, although they are better motivated, and have better expertise – at least when they are selected for their expertise – are also unrepresentative. Further, both legislative subcommittees and judges rely primarily on the very intelligence agencies they oversee for information.

I want to address these problems with two sets of institutional innovations. The first set focuses on the day-to-day scrutiny of requests to spy. For these requests, I

want to argue that America's procedures for judicial warrants should be overhauled, and even more boldly that judges should be replaced by an elected panel of overseers.

Let me begin with my recommendation for changing the procedures of judicial review. The principal problem with the procedures of judicial review today is that they are one-sided. Those who provide the court with the information necessary to determine whether warrants should be granted are the same agents seeking approval from the court.

The procedures of *ex ante* review can become less one-sided, I want to argue, with the incorporation of a devil's advocate. The term "devil's advocate" comes from an institution in the Catholic Church dating from 1587 to 1983. During the canonization process (the process of declaring a deceased person a saint), the devil's advocate's task was to argue against canonization. He was meant to question the character of the proposed saint by looking for faulty arguments or false evidence supporting the individual's case.

The devil's advocates that I envision for warrant procedures would be experienced intelligence analysts, on partial loan from America's various intelligence agencies.¹⁵ They would have full access to operational details, and they would be responsible for making the best case against spying to those considering warrant requests. Devil's advocates would be "embedded" in operational teams within intelligence agencies. But they would receive specialized training and would report not just to those who manage the operational team on which they work but also to a highly placed executive who manages the devil's advocates and reports directly to the Director of National Intelligence. The organizational structure of the devil's advocates would thus be de-

¹⁵One can even imagine a service requirement in the devil's advocate office for intelligence agents to be promoted to certain managerial or executive levels.

signed to ensure that devil's advocates could quickly sound the alarm when they are being denied important information.

For every warrant (or framework warrant) request, the devil's advocate would produce and submit a dissent to those considering the warrant request, offering up the strongest reasons why the operation should not go forward. The devil's advocate could, for example, poke holes in the logic of the arguments in the warrant request, she could offer information not produced in the warrant, or she could offer alternative interpretations of data presented in the warrant.

The idea of the devil's advocate has been employed before in intelligence. For example, since the Agranat Commission in 1973-74, the Military Intelligence (MI) division of the Israeli Defense Force has had a two person devil's advocate team, reporting directly to the Director of MI. The devil's advocate is tasked with making the best case against proposals under consideration within MI. After the Yom Kippur war, in which Israel was attacked by surprise, MI formed the devil's advocate group to ensure that key hypotheses and concepts leading to recommendations did not go unquestioned.

The principal difference between the Israeli case and my proposal is that in the Israeli case, the devil's advocate informs the decisions of the intelligence agency, whereas in my proposal the devil's advocate primarily informs the decision (by judges, etc.) to grant a warrant. In principle these two institutions are compatible. One could have multiple devil's advocates throughout the decision making process. In practice, however, too many skeptics built into the decision making process may slow decision making more than they bolster the quality of decisions.

Since the institution of the devil's advocate is added to a procedure in which someone already makes the case for a warrant, the resulting procedure is an adversarial one. In adversarial procedures the assumption is that neither side possesses the

whole truth or the best reasons. Both sides have full access to relevant information and both sides use this information to make their best case. Ideally, however, the “best” case is presented by neither side. Instead it is constructed deliberately by those who hear both sides of the argument.

The introduction of the devil’s advocate into warrant granting procedures significantly mitigates Sagar’s “structural dilemma,” thereby rendering warrants a significantly stronger mechanism of control. Intelligence agencies are less likely to get away with warrant proposals that withhold or present false information or provide implausible interpretations of the facts when there is someone privy to the same information charged with criticizing their proposal. Hence, even though those who scrutinize warrant requests do not have direct access to the information the requests contain, they have an ally of sorts to certify or call into question the quality of the information in the warrant.

The costs of the devil’s advocate are, like the costs of warrants, mostly administrative. Some might worry, again, that these costs will slow the application process or even deter justified spying. But these costs can be justified by the effectiveness boost they provide to the warrant process, and further they can be mitigated by some of the same institutions used to decrease the costs of warrants. The devil’s advocate would thus be a strong improvement on existing American warrant institutions.

So, the devil’s advocate mitigates the informational problems for warrant procedures. It does not, however, address the problem of bias. Generally, bias in review procedures can be addressed either by making the principal more representative (as I discussed in the previous chapter), or by introducing what Adrian Vermeule (2007, 31) calls a “veil rule.” As a general rule, making a principal more representative of the population of those its decisions affect increases the likelihood that all of the relevant interests will be considered. Boosting representativeness as a strategy for reducing

bias, however, has limitations, since deliberation rarely proceeds as political theorists think it should, especially when groups get large.¹⁶ Pervasive inequalities and stereotypes, for example, lead to privileging some people's contributions to discussion and discounting others'. Since people tend to defer to wealth, power, education, articulate speech, and membership in high status groups, representativeness can trade one kind of bias for another.

Veil rules "suppress self-interested behavior on the part of decisionmakers by subjecting the decisionmakers to uncertainty about the distribution of benefits and burdens that will result from a decision." (Vermeule 2007, 31) If decisionmakers do not know exactly who stands to gain or lose by their decision, it is very difficult for them to successfully bias their decisions. The principal difficulties with attaching veil rules to *ex ante* review are that *ex ante* review is typically done with all the relevant facts, and the decisions of overseers do not usually generalize, that is they do not typically apply to cases other than the one under review. Hence it is almost always apparent to overseers who stands to gain and lose from their decision in the short term, and there is little danger that their immediate decision will affect future decisionmaking. Overseers therefore can stack the deck.

These difficulties with veiling rules can be partially overcome. But like making an overseer more representative, veiling rules have their limitations. Perhaps the most promising way to institute a veiling rule for *ex ante* review is to enact a system of precedent: decisions made today must be respected in future decisions. The effectiveness of such a system would rely on the threat that decisions would be overturned if they are not consistent with past decisions. This threat could be instituted by making the overseers' decisions subject to scrutiny by a higher court.

¹⁶Sanders (1997) highlights many of the limitations of deliberation. The very large empirical literature on juries also highlights these limitations. For an overview, see Devine et al. (2001).

The question is whether by making the overseer more representative or by using veiling rules we can create a more unbiased overseer without sacrificing considerably other important features a panel of judges possesses such as expertise, trustworthiness with secrets, and motivation. For simplicity, we can set aside the issue of motivation, since our overseer will be performing *ex ante* review of all requests to spy.

A quick look at a few institutional forms suggests that if we focus on just one criterion, the other criteria suffer. It is a bit like squeezing one end of a balloon and watching the other end inflate. Suppose, for example, in order to maximize unbiasedness we constitute an overseer with a group of fifty individuals randomly selected from the population. Even if the random sample proved unbiased (which for reasons I mentioned above we should doubt), it would not likely have a high degree of expertise, and there would be some danger that those selected would not be trustworthy with secrets. Hence representativeness seems to be purchased at the price of expertise and trustworthiness with secrets.

Consider a second option. The executive selects a group of experienced and trustworthy professionals. Assuming she does not simply select her cronies (always a danger with appointment), we have relatively trustworthy experts. The overseer is extremely unlikely, however, to represent a wide variety of interests. For example, since expertise in intelligence and national security is almost always acquired within the government, and since appointment by the President likely means selection of a group who view the administration favorably, it is difficult to imagine an appointed overseer that does not have a pro-government bias.

Constituting an overseer that scores highly on all three criteria is no easy task, then. Below I explore three attempts to constitute such an overseer, and I argue that one of these attempts – electing overseers – succeeds marginally better than the other two.

The first attempt begins from the idea of a randomly selected panel of overseers, and seeks to solve the problems of trustworthiness and expertise with ancillary institutions. First, a large number of individuals (approximately 75) would be randomly selected from the population. These individuals would then undergo thorough background checks, to ensure that they can be trusted with sensitive and secret information. While the background checks are being conducted, the individuals would undergo extensive training on the standards by which they will judge cases, on previous decisions made by the panel, and on how to deliberate effectively and to ensure all voices are heard. Those who pass the background checks would then serve on a panel of overseers, presided over by a judge, for a term of approximately three months. The presiding judge's role would be simply to direct the panel to decide by the appropriate standards for the case at hand, just as a judge directs a jury to consider the elements of the law. Including training, each person's service would last roughly four months.

How does this modified random panel grade on the three criteria? On unbiasedness the panel receives an above average grade. It is highly representative of the population. But there is still the worry that, even with deliberation training, the panel will sometimes be "ruled" by stereotypes and background inequalities. On expertise, the panel is slightly better than a purely random sample, since it receives training, but a few weeks of training does not make someone an authority on intelligence or even on the kind of practical reason required to make good judgments. Since the panel's knowledge of prior decisions will be limited, it will also be very difficult to pair with institutions of precedent discussed above. So, on expertise the modified random panel receives a relatively low grade. Finally, on trustworthiness with secrets, the modified random panel receives a good grade, but not as good a grade as a group selected from intelligence professionals with years of experience successfully keeping secrets.

The second attempt to constitute an overseer that performs well on all three criteria also makes use of random selection. This time, however, individuals would not be selected from the population at-large, but rather from the set of people who already have some form of secret clearance. Over 4.8 million Americans have some level of secret clearance. These individuals are spread throughout the national security establishment, including not just intelligence agencies, but also the military, the Department of State, etc. These individuals would not require background checks, but they would still require training on the standards by which they will judge cases, on previous decisions made by the panel, and on how to deliberate effectively and to ensure all voices are heard. Their service would also be presided over by a judge and would last approximately four months.

This panel – call it the secret clearance panel – scores lower on unbiasedness but higher on expertise and trustworthiness with secrets. It scores lower on unbiasedness because it is significantly less representative, since it is not selected from the population at large. There is a further danger that it is systematically biased in favor of intelligence agencies. Although this bias is a possibility, it is important to remember that many of the people with secret clearances do not work for or inside intelligence agencies; they work for other arms of the executive branch or for private contractors. Bringing together people with many different institutional affiliations may increase the scrutiny given to requests to spy, since each agency is likely to view the need for a particular piece of information differently.

The secret clearance panel scores higher on expertise since it mostly selects people who have some experience in national security; and it scores higher on trustworthiness with secrets since it selects people who have not only passed background checks but who also tend to have years of experience keeping secrets. Like the modified random panel, however, since the panel's knowledge of prior decisions will be limited, it will

also be very difficult to pair with institutions of precedent.

The secret clearance panel has an added benefit that it compels those with secret clearance to think carefully about the ethical issues in relationship to their secretive professions. Just as participation on a jury can provide “education in citizenship” by stimulating participants to critically evaluate the conduct of their fellow citizens, participation on the “intelligence panel” could serve as a kind of practical training in professional ethics, by engaging those with secret clearance in the moral decision making of their peers. To the extent that there is an (unhealthy) ethos of loyalty and obedience pervading intelligence agencies, mandatory service on the secret clearance panel may begin to provide a counterweight. As countries like the United States perform more of their activities in secret, the protection of basic liberal and democratic values rely more on reflective professionals willing to scrutinize the decisions of their peers and superiors and courageous enough to raise red flags when they witness wrongful or illegal conduct.

The final attempt employs not random selection or appointment but elections. The idea is to populate the panel of overseers by electing approximately fifty non-partisan individuals from one at-large district to serve for renewable terms of approximately ten years. So that the public would not have to do their homework on hundreds of candidates at a time, the terms of the elected individuals could be staggered. Further, since those elected to the panel would have significant time to learn in detail the standards they are meant to apply to each set of circumstances, there would be no need for judges to preside over their deliberations.

How well would the elected panel grade on the three criteria? To some extent, it depends on how well the public chooses the panel’s members. In principle, since fifty is a large enough number for a diversity of interests to be represented on the panel, and since the panel is chosen in one at-large district, these interests could mirror the

interests of the public at large. If the public does not show up for elections, or it does not do its homework, however, there is no guarantee that its interests will be well represented. Certainly the elected panel will not be as representative as the modified random panel, but it is likely to be more representative than the secret clearance panel. The elected panel is also more likely than the panels selected by lot to deliberate effectively, since the people are likely to choose panelists capable of strongly advocating their own views.

On the criteria of expertise and trustworthiness, the elected panel would also grade highly, assuming the people select mainly on these traits. Since the panelists will not be affiliated with a party, competence, experience, and integrity are likely to play a strong role in the people's selection. In addition, electoral institutions such as debates, the public financing of campaigns, and the information provided to voters before they cast their ballot could be shaped so as to focus the people's deliberation on these traits. Even more importantly, however, the long term of the elected panels would permit elected panelists to acquire expertise and experience protecting state secrets. The long term would, further, make the institution of precedent possible. Hence each decision by the panel particularly in its early stages would carry with it far reaching consequences and thus more uncertainty about whose interest would be effected, thereby building in a veiling mechanism and increasing the unbiasedness of the elected panel.

The table below summarizes my analysis of possible overseers for *ex ante* review. It force ranks the four main options I considered on the three criteria for a good overseer, and it also notes other advantages I mentioned for certain options. As I mentioned above, I think the elected panel fares slightly better than the other options. It has a strong edge on the random panel on both expertise and trustworthiness, and ranks only slightly behind the random panel on unbiasedness. It fares much better than

the appointed panel on unbiasedness, while grading similarly on the other criteria. And it beats the modified random panel on all three criteria, though it does not have the beneficial educative effect that the modified random panel does. If one were to argue that unbiasedness is far and away the most important criteria, I still think it makes sense to pick the elected panel, since it likely fares almost as well on this criterion as the modified random panel. Remember that the elected panel is likely to have a deliberative advantage over the modified random panel and it can incorporate veiling mechanisms that the modified random panel cannot incorporate. Similarly, if one were to argue that expertise is the most important criterion, I would suggest the elected panel is still the best choice, since it is not obvious to me that an executive will be a significantly better selector of expert panelists than the people.

Table 9.1: Overseers for *Ex Ante* Review

	Appointed Panel	Modified Random Panel	Secret Clearance Panel	Elected Panel
Expertise	++	--	-	+
Unbiasedness	--	++	-	+
Trustworthiness	+	--	-	++
Educative			✓	
Veiling	✓			✓

Thus far I have been attending to the day-to-day review of requests to spy. I have argued that warrants should be reviewed not by the courts, but by an elected panel, and that the procedure for granting warrants should incorporate a devil's advocate. I now want to turn away from day-to-day review toward strategic oversight. In the model that pairs judicial warrants with legislative oversight, legislators have the responsibility for this strategic oversight.

The problems facing legislative oversight, I argued in the previous chapter, are similar to those that plague judicial warrants, only more pronounced. Legislators are poor principals: they tend to be unmotivated, inexperienced, and unrepresentative. Further, they rely on those they oversee to provide the necessary information for successful oversight.

What is the best way to address these issues institutionally? Let me first rule out a couple of responses. First, the informational problem with strategic oversight cannot be solved with a devil's advocate. Overseers must have a nearly complete view of intelligence activities in order to subject those activities to strategic review. If overseers wish to know whether an intelligence agency addresses the gravest and most imminent threats, for example, then they need to know the information that intelligence agencies use to assess threats. Similarly, if overseers wish to evaluate whether intelligence agencies spend their resources efficiently, they need to know not only how intelligence agencies spend their resources, but also the myriad alternative ways that intelligence agencies could spend their resources. Hence the informational requirements for strategic oversight stretch well beyond what is required to evaluate particular cases of spying. They include, among many other things, agency strategies, procedures, budgets, head counts, and communications. In practice, strategic overseers would need to be privy to nearly all of the secrets intelligence agencies keep.

Nor can the motivational issue be set aside (as I did above) on account of the review being *ex ante*. Since strategic review is typically done *ex post*, overseers are not compelled institutionally to perform their role faithfully. Even if hearings or investigations were mandatory, it would not compel overseers to engage in a systematic review of intelligence operations. Overseers must have strong internal motivations and very few distractions, then.

I want to suggest that the problems of legislative oversight are best addressed by

moving – again – to an elected panel of overseers. Indeed, I think it makes sense to elect one relatively large panel of intelligence overseers (perhaps as many as 75) that alternate between scrutinizing warrants and performing strategic oversight. Such rotations could deepen the expertise and widen the perspective of overseers.

For strategic oversight, elected panels grade higher than legislative subcommittees on all relevant criteria. They are more likely to be motivated because they are devoted institutionally to intelligence issues, and their bids for reelection (which would be comparatively infrequent) would hinge almost exclusively on their performance as intelligence overseers, which could be evaluated by their peers and by the legislators and bureaucrats who deal with them. They would be less biased because they would be selected by the people from one at-large district, rather than self-selected to a subcommittee. Finally, they would be more likely to be trustworthy experts since they would be selected by the people primarily on these traits, whereas legislators are evaluated on scores of other traits irrelevant to their performance as overseers.

Although it may be possible to *appoint* panelists who are slightly more experienced or more trustworthy, appointment remains an inferior institutional arrangement. One reason is that appointment would not permit rotation between responsibilities (assuming those who perform *ex ante* review are elected). More importantly, however, appointed panelists are more likely to exhibit a systematic pro-government bias than elected panels. Elected panels will likely be more reflective of the diversity of views in the population.

So an elected panel would be the best strategic overseer. Let me now try to fill in how strategic oversight would function. Strategic overseers would have the power to hold hearings, subpoena information, and issue official government reports. No information would be off limits to them. Their principal responsibilities would be to develop an in depth picture of America's intelligence activities, and to evaluate the

effectiveness, efficiency, and legitimacy of these activities. Strategic overseers would prepare and issue classified reports to legislators (most likely the subcommittees overseeing intelligence agencies) and to the executive; they would also issue sanitized reports directly to the public. These reports would focus broadly on the threats facing the country, and on how well intelligence resources are being employed to counter these threats; they would report on whether intelligence activities are consistent with the rule of law; and perhaps most importantly they would report on whether the resources devoted to secret activities can be justified, given other national priorities.

Unlike the panelists who grant warrants, the power of strategic overseers is mostly hortatory. They have no veto on the conduct of intelligence agencies or their congressional overseers. Their power derives, instead, from their expertise, and from their ability to alert the executive, the congress, and most importantly, the public to inefficiency or illegality within the secret branches of government. Strategic overseers do have the powers to shame or embarrass intelligence agencies or overseers, but apart from these powers, they possess little in the way of sanctions and rewards.

An objector might argue that these powers of persuasion are insufficient for the secrecy council to be much of a force. They might point out, for example, that the 9/11 commission along with a handful of other august blue ribbon commissions have called for reform and reorganization in the U.S. intelligence community, but little has been done in response to these calls. (Zegart, 2009) So they might ask why a panel of strategic overseers would succeed where these commissions have failed.

One reason is that commissions typically make their recommendations and then disband. They do not remain around to prod and niggle the Congress and the executive to carry out their recommendations. The strategic oversight panel, in contrast, would have a permanent presence in Washington, and as a branch of the government, it would have its own bully pulpit.

A second, more important, reason is that the strategic oversight panel, as I envision it, would have a unique kind of power that blue ribbon commissions have never had. It would be the final arbiter of what information it makes public. No other branch of government would have the power to redact or censor its reports. It could make public whatever information it determined the people need to know. It would, in other words, have the power to *declassify*.

A final related point is that strategic overseers would *appear* more independent than most commissions, since they have been elected by the people to oversee intelligence activities. Since no branch of government can force strategic overseers to alter or censor their findings, people would be more willing to accept the council's reports as unvarnished.

So, a panel of strategic overseers could spur the Congress and the executive into action on intelligence policy and it could serve as a key line of defense against the abuse of state secrets. But these good effects depend on the secrecy council having access to the executive's most closely held secrets, and one might doubt that strategic overseers could ever achieve this access. The panel of strategic overseers has the power to subpoena information and to hold hearings, but so too do legislative overseers, and these overseers, we have seen, rely on the intelligence agencies they are supposed to be overseeing to provide all relevant information. They are prisoners in Sagar's structural dilemma. Is it realistic to think the panel of strategic overseers can escape this structural dilemma?

Since the structural dilemma is real, it cannot be fully escaped, but it can – as I suggested above in my discussion of warrant procedures – be mitigated. One way to mitigate it is to pair strategic oversight with the warranting process. As I argued above, the panels would produce an informational trove of inestimable value for overseers. Two further institutions can also substantially mitigate the structural dilemma

for the panel of strategic overseers. The first is to provide strategic overseers with all reasonable investigative resources. Hence the panel should have a sizable permanent staff and it should have the power to direct or commission other investigative arms of the federal government, such as the General Accounting Office (GAO).

More intriguingly, I want to argue that the panel of strategic overseers should have its own internal intelligence agency to “spy on the spies.” This agency would develop a small group of informants placed throughout the intelligence community to alert strategic overseers to inefficiency, illegality, etc. and to ensure that the information that strategic overseers receive from the intelligence community is factual and complete.

On its face, this proposal may seem to lead to an infinite regress. If we need spy number two to watch spy number one, who watches spy number two? Perhaps spy number three? But then who watches her? The problem calls to mind Juvenal’s famous question: “*Quis custodiet ipsos custodes?*” But an infinite regress need not follow if the aim is merely to achieve a net reduction in unconstrained power. It is less dangerous to leave a small intelligence agency directed by the panel of strategic overseers “unwatched” than to leave behemoth agencies, such as the CIA, the NSA, and the FBI without supervision. Besides, the panel of strategic oversight’s intelligence agency need not remain entirely unwatched, since it could be a requirement that all of its requests to spy be warranted.

Practical worries no doubt present themselves to this proposal. Spy agencies detect and turn informants – it is a core part of what they do. So there is no guarantee that the information passed from the strategic oversight panel’s informants would be trustworthy. Revelation of the fact that the panel placed informants inside intelligence agencies could also lead to difficulties. Intelligence agencies could, for example, come to distrust the secrecy council, leading it only to guard its secrets more vigi-

lantly. Furthermore, there are the costs of recruiting and managing spies. Running agents costs money and resources that could be spent on hearings, investigations, etc. Finally, as I suggested above, these agents could engage in their own abusive spying.

Yet giving the panel of strategic overseers its own miniature intelligence agency still has merit. Intelligence agencies may be more likely to behave in accordance with legal and ethical precepts when they suspect that their illegal or illegitimate conduct may be discovered, and especially when they suspect that their conduct will be exposed to the public. They also may be more forthcoming with information when they suspect that the secrecy council is likely to discover the information anyway. Hence there are good reasons to think the structural dilemma that faces the panel of strategic overseers can be mitigated.

9.3 Conclusion

In this chapter I stepped back from the American mechanisms of control that I critiqued in the previous chapter and tried to understand the universe of possible mechanisms to control intelligence agencies. I identified five mechanisms – accountability mechanisms, mechanisms of direct interference, information alteration mechanisms, selection mechanisms, and partitioning mechanisms – and I argued that partitioning mechanisms have the most unexplored potential for constraining intelligence agencies.

I then offered my own proposal for intelligence reform in America, which features among a host of other intricacies an elected panel of overseers that would perform both *ex ante* review of proposals to spy and the strategic review of intelligence agencies. I did not claim that my proposal is optimal – on the contrary, I argued that many more possible combinations of mechanisms of control should be considered. But I did claim that my proposed institutions would be better suited for ensuring that

America's intelligence agencies conform to the ethical principles that I argued for in earlier chapters.

Bibliography

- Alchian, Armen and Harold Demsetz. 1972. "Production, Information Costs, and Economic Organization." *American Economic Review* 62:777–795.
- Alexander, Larry. 1985. "Pursuing the Good - Indirectly." *Ethics* 95(2):315–332.
- Alexander, Larry. 1989. "Comment: Personal Projects and Impersonal Rights." *Harvard Journal of Law & Public Policy* 813:813–826.
- Alexander, Larry and Michael Moore. 2012. "Deontological Ethics." *Stanford Encyclopedia of Philosophy* .
- Allen, Anita. 2008. "The Virtuous Spy: Privacy as an Ethical Limit." *The Monist* 91(1):3–22.
- Amar, Akhil Reed. 1994. "Fourth Amendment First Principles." *Harvard Law Review* 107:757–819.
- Amir, On, Dan Ariely and Nina Mazar. 2008. "The Dishonesty of Honest People: A Theory of Self-Concept Maintenance." *Journal of Marketing Research* 45(6):633–644.
- Anderson, James. 1975. *Public Policy Making*. New York: Prager.

- Andreoni, J. and R. Petrie. 2004. "Public goods experiments without confidentiality: A glimpse into fund-raising." *Journal of Public Economics* 88(7-8):1605–1623.
- Appelbaum, Arthur Isak. 1998. "Are Violations of Rights Ever Right?" *Ethics* 108(2):340–366.
- Aquinas, Thomas. 2002. *On Law, Morality, and Politics*. Indianapolis: Hackett Pub.
- Arendt, Hannah. 2006. *On Revolution*. New York: Penguin Books.
- Ariely, Dan. 2012. *The (Honest) Truth About Dishonesty: How We Lie to Everyone - Especially Ourselves*. New York: Harper Collins.
- Aristotle. 1996. *Aristotle: the Politics and the Constitution of Athens*. Cambridge Texts in the History of Political Thought Cambridge: Cambridge University Press.
- Austin, J. L. 1956. "A Plea for Excuses." *Proceedings of the Aristotelian Society* 62:1–30.
- Axelrod, Robert. 2006. *The Evolution of Cooperation*. Revised ed. New York: Basic Books.
- Bachrach, Peter and Morton Baratz. 1962. "Two Faces of Power." *American Political Science Review* 56(4):947–952.
- Baer, Robert. 2002. "See No Evil: The True Story of a Ground Soldier in the CIA's War on Terrorism."
- Bales, R.E. 1971. "Act-Utilitarianism: Account of Right-Making Characteristics or Decisions Procedures?" *American Philosophical Quarterly* 8:257–265.
- Bamford, James. 2002. *Body of Secrets: Anatomy of the Ultra-Secret National Security Agency*. New York, NY: Anchor.

- Banville, John. 1997. *The Untouchable*. New York: Vintage International.
- Barke, Richard and William Riker. 1982. "A Political Theory of Regulation with Some Observations on Railway Abandonments." *Public Choice* 39:73–106.
- Bateson, M., D. Nettle and G. Roberts. 2006. "Cues of being watched enhance cooperation in a real-world setting." *Biology Letters* 2(3):412–414.
- Beitz, Charles R. 2005. Cosmopolitanism and Global Justice. In *Current Debates in Global Justice*, ed. Gillian Brock and Darrel Moellendorf. Vol. 2 Springer-Verlag pp. 11–27.
- Benn, Stanley. 1971. Privacy, Freedom, and Respect for Persons. In *Privacy*, ed. J. Roland Pennock and John Chapman. New York: Atherton Press.
- Bentham, Jeremy. 1838. *The Works of Jeremy Bentham: Published under the Superintendence of His Executor, John Bowring*. London, UK: Adamant Media.
- Bentham, Jeremy. 1996. *An Introduction to the Principles of Morals and Legislation*. Oxford; New York: Clarendon Press ; Oxford University Press.
- Bentwich, Norman. 1910. "Espionage and Scientific Invention." *Journal of the Society of Comparative Legislation* 2(10):243–249.
- Bok, Sissela. 1989. *Secrets: On the Ethics of Concealment and Revelation*. New York: Vintage Books.
- Bok, Sissela. 1999. *Lying: Moral Choice in Public and Private Life*. New York, NY: Vintage Books.
- Bordo, Susan. 2004. *Unbearable Weight: Feminism, Western Culture, and the Body*. University of California Press.

- Born, Hans and Maria Caparini, eds. 2007. *Democratic Control of Intelligence Services*. Burlington: Ashgate.
- Bowers, Rick. 2010. *Spies of Mississippi: The True Story of the Spy Network that Tried to Destroy the Civil Rights Movement*. Washington, D.C.: National Geographic.
- Brandt, R. B. 1990. "The Science of Man and Wide Reflective Equilibrium." *Ethics* 100(2):pp. 259–278.
- Brandt, Richard B. 1998. *A Theory of the Good and the Right*. Amherst, N.Y.: Prometheus Books.
- Bratman, Michael. 1999. *Faces of Intention: Selected Essays on Intentionality*. Cambridge, UK: Cambridge University Press.
- Brink, David. 1986. "Utilitarian Morality and the Personal Point of View." *Journal of Philosophy* 83:417–438.
- British War Office. 1894. *Manual of Military Law*. London, UK: Harrison and Sons.
- Bruce, Gary. 2010. *The Firm: the Inside Story of the Stasi*. Oxford; New York: Oxford University Press.
- Buchanan, Allen. 2006. "Institutionalizing the Just War." *Philosophy & Public Affairs* 34(1):2–38.
- Buchanan, Allen. 2010. The Legitimacy of International Law. In *The Philosophy of International Law*, ed. Besson S and Tasioulas J. Oxford: Oxford University Press.

- Burnham, T.C. 2003. "Engineering Altruism: A Theoretical and Experimental Investigation of Anonymity and Gift Giving." *Journal of Economic Behavior and Organization* 50(1):133–144.
- Burnham, T.C. and B. Hare. 2007. "Engineering Human Cooperation : Does Involuntary Neural Activation Increase Public Goods Contributions?" *Human Nature* 18(2):88–108.
- Calvert, Randall, Matthew McCubbins and Barry Weingast. 1989. "A Theory of Political Control and Agency Discretion." *American Journal of Political Science* 79:588–611.
- Chaiken, Shelly and Yaacov Trope. 1999. *Dual-Process Theories in Social Psychology*. New York: Guilford Press.
- Chesterman, Simon. 2006. "The Spy Who Came in From the Cold War: Intelligence and International Law." *Michigan Journal of International Law* 27:1071–1130.
- Chesterman, Simon. 2011. *One Nation Under Surveillance: A New Social Contract to Defend Freedom Without Sacrificing Liberty*. Oxford, UK: Oxford University Press.
- Chisholm, Roderick and T.D. Feehan. 1977. "The Intent to Deceive." *The Journal of Philosophy* 74:143–159.
- Connolly, William E. 1993. *The Terms of Political Discourse*. Oxford: Blackwell.
- Cooter, Robert. 2002. *The Strategic Constitution*. Princeton: Princeton University Press.
- Daniels, Norman. 2003. "Reflective Equilibrium." *Stanford Encyclopedia of Philosophy* .

- Davis, James Kirkpatrick. 1992. *Spying on America: the FBI's Domestic Counterintelligence Program*. New York: Praeger.
- Dawes, R.M., J. McTavish and H. Shaklee. 1977. "Behavior, Communication, and Assumptions about Other People's Behavior in a Commons Dilemma Situation." *Journal of Personality and Social Psychology* 35(1):1–11.
- de Lazari-Radek, Katarzyna and Peter Singer. 2010. "Secrecy in Consequentialism: A Defense of Esoteric Morality." *Ratio* 23(1):34–58.
- de Vattel, Emmerich. 1883. *The Law of Nations: Or, Principles of the Law of Nature Applied to the Conduct and Affairs of Nations and Sovereigns*.
- Dennis, Mike. 2003. *The Stasi: Myth and Reality*. London: Longman.
- Devine, Dennis J., Laura D. Clayton, Benjamin B. Dunford, Rasmy Seyer and Jennifer Pryce. 2001. "Jury decision making: 45 years of empirical research on deliberating groups." *Psychology, Public Policy, and Law* 7(3):622–727.
- Dodd, Lawrence and Richard Schott. 1979. *Congress and the Administrative State*. New York: Wiley.
- Dovi, Susan. 2007. *The Good Representative*. Malden: Blackwell.
- Driver, Julia. 2005. "Consequentialism and Feminist Ethics." *Hypatia* 20(4):183–199.
- Driver, Julia. 2007. *Ethics: the Fundamentals*. Malden, Mass.: Blackwell Pub.
- Eisner, Marc and Kenneth Meier. 1990. "Presidential Control versus Bureaucratic Power." *American Journal of Political Science* 34:267–287.
- Epstein, David and O'Halloran Sharyn. 1994. "Administrative Procedures, Information, and Agency Discretion." *American Journal of Political Science* 38(3):697–722.

- Ernest-Jones, M., D. Nettle and M. Bateson. 2011. "Effects of Eye Images on Everyday Cooperative Behavior: A Field Experiment." *Evolution and Human Behavior* 32(3):172–178.
- Fain, Tyler, Katharine Plant and Ross Milloy, eds. 1977. *The Intelligence Community: History, Organization, and Issues*. Public Documents Series New York: RR Bowker.
- Fallis, Daniel. 2009. "What is Lying?" *The Journal of Philosophy* 106(1):29–56.
- Fearon, James. 1999. Electoral Accountability and the Control of Politicians: Selecting Good Types versus Sanctioning Poor Performance. In *Democracy, Accountability, and Representation*, ed. Adam Przeworski, Susan Stokes and Bernard Manin. Cambridge: Cambridge University Press.
- Feinberg, Joel. 1983. "Autonomy, Sovereignty, and Privacy: Moral Ideas in the Constitution?" *Notre Dame Law Review* 58:445–492.
- Ferejohn, John and Charles Shipan. 1990. "Congressional Influence on Bureaucracy." *Journal of Law, Economics, and Organization* 6:1–20.
- Finer, Herman. 1949. *The Theory and Practice of Modern Government*. Revised ed. New York: Holt.
- Fiorina, Morris. 1977. *Congress: Keystone of the Washington Establishment*. New Haven: Yale University Press.
- Fiorina, Morris. 1981. Congressional Control of the Bureaucracy: A Mismatch of Incentives and Capabilities. In *Congress Reconsidered*, ed. Lawrence Dodd and Bruce Oppenheimer. New Haven: Yale University Press.
- Fiorina, Morris. 1982. "Legislative Choice of Regulatory Forms: Legal Process of Administrative Process?" *Public Choice* 39:33–36.

- Fiorina, Morris and Roger Noll. 1978. "Voters, Bureaucrats, and Legislators: A Rational Choice Perspective on the Growth of the Bureaucracy." *Journal of Public Economics* 9:239–254.
- Foreman, D M and C Farsides. 1993. "Ethical Use of Covert Videoing Techniques in Detecting Munchausen Syndrome by Proxy." *BMJ* 307(6904):611–613.
- Foucault, Michel. 1995. *Discipline and Punish: The Birth of the Prison*. New York: Vintage.
- Freeman, Samuel. 2000. "Deliberative Democracy: A Sympathetic Comment." *Philosophy & Public Affairs* 29(4):371–418.
- Freund, Paul. 1971. Privacy: One Concept or Many? In *Nomos XIII: Privacy*, ed. J. Roland Pennock and John Chapman. New York: Atherton Press.
- Fried, Charles. 1970. *An Anatomy of Values*. Cambridge, MA: Harvard University Press.
- Garner, Bryan and Henry Campbell Black. 1991. "Espionage n." *Black's Law Dictionary* .
- Gendron, Angela. 2005. "Just War, Just Intelligence: An Ethical Framework for Foreign Espionage." *International Journal of Intelligence and CounterIntelligence* 18(3):398 – 434.
- Goldman, Jan. 2006. *Ethics of Spying: A Reader for the Intelligence Professional*. Lanham, MD: The Scarecrow Press.
- Goldman, Jan. 2010. *Ethics of spying: a Reader for the Intelligence Professional, v.2*. United States: Scarecrow.

- Goodin, Robert E. 1985. *Protecting the Vulnerable : A Reanalysis of Our Social Responsibilities*. Chicago: University of Chicago Press.
- Goodin, Robert E. 1988. "What Is So Special about Our Fellow Countrymen?" *Ethics* 98(4):663–686.
- Gosseries, Alex. 2005. "Publicity." *Stanford Encyclopedia of Philosophy* .
- Govier, Trudy. 1996. "Trust and Totalitarianism: Some Suggestive Examples." *Journal of Social Philosophy* 27(3):149–163.
- Grant, Ruth and Robert Keohane. 2005. "Accountability and Abuses of Power in World Politics." *American Political Science Review* 99(1):29–43.
- Gwyn, William. 1965. *The Meaning of the Separation of Powers: An Analysis of the Doctrine from Its Origin to the Adoption of the United States Constitution*. New Orleans: Tulane University Press.
- Haley, K.J. and D.M.T. Fessler. 2005. "Nobody's watching? Subtle Cues Affect Generosity in an Anonymous Economic Game." *Evolution and Human Behavior* 26(3):245–256. cited By (since 1996) 166.
- Hamilton, Alexander, James Madison and John Jay. 2009. *The Federalist Papers: Alexander Hamilton, James Madison, John Jay*. New Haven: Yale University Press.
- Hammond, Thomas and Jack Knott. 1996. "Who Controls the Bureaucracy? Presidential Power, Congressional Dominance, Legal Constraints, and Bureaucratic Autonomy in a Model of Multi-Institutional Policy-Making." *Journal of Law, Economics, and Organization* 12:119–166.

- Hardin, Russell. 1990. *Morality within the Limits of Reason*. Chicago: Univ Of Chicago Press.
- Hardin, Russell. 2002. *Trust and Trustworthiness*. New York: Russell Sage Foundation.
- Hare, R. M. 1981. *Moral Thinking: its Levels, Method, and Point*. Oxford; New York: Clarendon Press ; Oxford University Press.
- Hare, R.M. 1973. "Rawls' Theory of Justice - I." *The Philosophical Quarterly* 23(91):144–155.
- Hare, R.M. 1979. "What is Wrong with Slavery?" *Philosophy & Public Affairs* 8(2):103–121.
- Hart, H. L. A. 1997. *The Concept of Law*. Clarendon law series. 2nd ed. New York ; Oxford: Oxford University Press.
- Hayward, Clarissa. 2000. *Defacing Power*. Cambridge: Cambridge University Press.
- Herbig, Katherine, Martin Wiskoff and James Riedel. 2002. *Espionage Against The United States by American Citizens: 1947-2001*. Monterey, CA: Defense Personnel Security Research Center.
- Hilbert, M. and P. Lopez. 2011. "The World's Technological Capacity to Store, Communicate, and Compute Information." *Science* 332(6025):60–65.
- Hippner, Christian. 2009. A Study into the Size of the World's Intelligence Industry PhD thesis Department of Intelligence Studies, Mercyhurst College.
- Hitz, Frederick Porter. 2008. *Why Spy? Espionage in an Age of Uncertainty*. New York: Thomas Dunne Books/St. Martin's Press.

- Hobbes, Thomas. 1994. *Leviathan*. Indianapolis: Hackett Publishing Company.
- Hobbes, Thomas. 1998. *On the Citizen*. Cambridge University Press.
- Hoffman, E., K. McCabe, K. Shachat and V. Smith. 1994. "Preferences, Property Rights, and Anonymity in Bargaining Games." *Games and Economic Behavior* 7(3):346–380.
- Hohfeld, Wesley. 1919. *Fundamental Legal Conceptions as Applied in Judicial Reasoning*. New Haven, CT: Yale University Press.
- Holmstrom, Bengt. 1979. "Moral Hazard and Observability." *Bell Journal of Economics* 10:74–91.
- Hooker, Brad. 2008. "Rule Consequentialism." *Stanford Encyclopedia of Philosophy*.
- Huber, John, Charles Shipan and Madelaine Pfahler. 2001. "Legislatures and Statutory Control of Bureaucracy." *American Journal of Political Science* 45(2):330–345.
- Hurka, Thomas. 1997. *The Justification of National Partiality*. Oxford University Press.
- Hurka, Thomas. 2005. "Proportionality in the Morality of War." *Philosophy & Public Affairs* 33(1):34–66.
- Irons, Jenny. 2010. *Reconstituting Whiteness: The Mississippi Sovereignty Commission*. Nashville: Vanderbilt University Press.
- Jackson, Frank. 1998. *From Metaphysics to Ethics*. New York, NY: Oxford University Press.
- Javers, Eamonn. 2010. *Broker, Trader, Lawyer, Spy: The Secret World of Corporate Espionage*. New York: HarperCollins.

- Jensen, Michael and William Meckling. 1976. "The Theory of the Firm." *Journal of Financial Economics* 3:305–360.
- Jervis, Robert. 2006. "Reports, Politics, and Intelligence Failures: The Case of Iraq." *Journal of Strategic Studies* 29(1):3–52.
- Johnson, Loch. 2004. "Congressional Supervision of America's Secret Agencies: The Experience and Legacy of the Church Committee." *Public Administration Review* 64:3–14.
- Johnson, Loch. 2006. "Secret Spy Agencies and a Shock Theory of Accountability." *Occasional Paper Series: Department of International Affairs, The University of Georgia* pp. 1–14.
- Johnson, Loch. 2007. Intelligence Oversight in the United States. In *Intelligence and Human Rights in the Era of Global Terrorism*, ed. Steve Tsang. Westport: Praeger.
- Kant, Immanuel. 1990. *Foundations of the Metaphysics of Morals and What Is Enlightenment?* Englewood Cliffs, N.J.: Prentice Hall.
- Kant, Immanuel. 1991. *Kant: Political Writings*. Cambridge; New York: Cambridge University Press.
- Kant, Immanuel. 2012. *Groundwork of the Metaphysics of Morals*. Cambridge: Cambridge University Press.
- Karmis, Dimitrios and Wayne Norman. 2005. In *Theories of Federalism: A Reader*. New York: Palgrave MacMillan.
- Kessler, Ronald. 2003. *The Bureau: the Secret History of the FBI*. New York: St. Martin's Paperbacks.

- Knightly, Philip. 1986. *The Second Oldest Profession: Spies and Spying in the Twentieth Century*. New York, NY: W.W. Norton.
- Knobe, Joshua. 2006. "The Concept of Intentional Action: A Case Study in the Uses of Folk Psychology." *Philosophical Studies* 130:203–231.
- Korsgaard, Christine. 1986. "The Right to Lie: Kant on Dealing with Evil." *Philosophy & Public Affairs* 15(4):325–349.
- Kotz, Nick. 2005. *Judgment Days: Lyndon Baines Johnson, Martin Luther King Jr., and the Laws that Changed America*. New York: First Mariner.
- Kurzban, R. 2001. "The Social Psychophysics of Cooperation: Nonverbal communication in a Public Goods Game." *Journal of Nonverbal Behavior* 25(4):241–259.
- Lackey, Douglas. 1989. *The Ethics of War and Peace*. Englewood Cliffs, NJ: Prentice Hall.
- Laertius, Diogenes. 1997. *Hellenistic Philosophy Introductory Readings*. Indianapolis: Hackett.
- Laurence, Stephen and Eric Margolis, eds. 1999. *Concepts: Core Readings*. Cambridge, Mass.: The MIT Press.
- Le Carre, John. 1963. *The Spy Who Came in from The Cold*. New York, NY: Pocket Books.
- Levy, Sanford S. 1994. "The Coherence of Two-Level Utilitarianism: Hare vs. Williams." *Utilitas* 6(02):301–309.
- Locke, John. 1988. *Two Treatises of Government*. Cambridge [England]; New York: Cambridge University Press.

- Lord, Charles, Lee Ross and Mark Lepper. 1979. "Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence." *Journal of Personality and Social Psychology* 37(11):2098–2019.
- Lowenthal, Mark. 2009. *Intelligence: From Secrets to Policy*. 4 ed. Washington: CQ Press.
- Luban, David. 1996. The Publicity Principle. In *The Theory of Institution Design*, ed. Robert Goodin. Cambridge, UK: Cambridge University Press pp. 154–198.
- Lukes, Steven. 2004. *Power: A Radical View*. 2 ed. New York: Palgrave MacMillan.
- Lyon, David. 2001. *Surveillance Society: Monitoring Everyday Life*. Philadelphia, PA: Open University Press.
- Mabbott, J. D. 1937. "Is Plato's Republic Utilitarian?" *Mind* 46(184):pp. 468–474.
- MacDonald, Jason. 2010. "Limitation Riders and Congressional Influence Over Bureaucratic Policy Decisions." *American Political Science Review* 104(4):766–782.
- Mackie, J.L. 1984. Utility and Rights. Oxford: Basil Blackwood chapter Rights, Utility and Universalization.
- Mahon, James. 2007. "A Definition of Deceiving." *International Journal of Applied Philosophy* 47:132–144.
- Mahon, James. 2008. "The Definition of Lying and Deception." *Stanford Encyclopedia of Philosophy* .
- Mansbridge, Jane. 2009. "A Selection Model of Political Representation." *The Journal of Political Philosophy* 17(4):369–398.

- Marquez, Xavier. 2011. "Spaces of Appearance and Spaces of Surveillance." *Polity* 44(1):6–31.
- Marx, Gary T. 1990. *Undercover: Police Surveillance in America*. Berkeley: Univ Of California Press.
- Mason, Andrew. 1997. "Special Obligations to Compatriots." *Ethics* 107(3):427–447.
- Mason, Elinor. 1998. "Can an Indirect Consequentialist Be a Real Friend?" *Ethics* 108(2):386–393.
- Mayer, Jane. 2011. "The Secret Sharer." *The New Yorker* .
- McCubbins, Matthew. 1985. "Regulating the Regulators: A Theory of Legislative Delegation." *American Journal of Political Science* 29:721–748.
- McCubbins, Matthew, Roger Noll and Barry Weingast. 1987. "Administrative Procedures as Instruments of Political Control." *Journal of Law, Economics, and Organization* 3(2):243–277.
- McCubbins, Matthew and Thomas Schwartz. 1984. "Congressional Oversight Overlooked: Police Patrols versus Fire Alarms." *American Journal of Political Science* 28:165–179.
- McMahan, Jeff. 1997. *The Limits of National Partiality*. Oxford University Press.
- McMahan, Jeff. 2005. "Just Cause for War." *Ethics & International Affairs* 19(03):1–21.
- Meier, Kenneth and Laurence O'Toole. 2006. "Political Control versus Bureaucratic Values: Reframing the Debate." *Public Administration Review* 66:177–192.

- Melnyk, Andrew. 2008. "Conceptual Analysis: A Two Step Program." *Nous* 42:267–291.
- Mill, John Stuart. 1979. *Utilitarianism*. Indianapolis: Hackett Pub. Co.
- Mill, John Stuart. 1989. *On Liberty and Other Writings*. Cambridge Texts in the History of Political Thought Cambridge: Cambridge University Press.
- Miller, David. 1988. "The Ethical Significance of Nationality." *Ethics* 98(4):647–662.
- Miller, David. 1997. *On Nationality*. New York: Oxford University Press.
- Miller, David. 2004. "Reasonable Partiality towards Compatriots." *Ethical Theory and Moral Practice* 8(1):63–81.
- Miller, David. 2009. "Democracy's Domain." *Philosophy & Public Affairs* 37(3):201–228.
- Moe, Terry. 1985. "Control and Feedback in Economic Regulation." *American Political Science Review* 79:1094–1116.
- Moe, Terry. 1987. "An Assessment of the Positive Theory of Congressional Dominance." *Legislative Studies Quarterly* 11:475–520.
- Moore, Adam. 1998. "Intangible Property: Privacy, Power, and Information Control." *American Philosophical Quarterly* 35:365–378.
- Moore, Adam. 2003. "Privacy: Its Meaning and Value." *American Philosophical Quarterly* 40(3):215–227.
- Moore, G. E. 1903. *Principia Ethica*. Cambridge: University Press.

- Moynihan, Daniel P. 1998. *Secrecy: The American Experience*. New Haven: Yale University Press.
- Nasheri, Hedieh. 2005. *Economic Espionage and Industrial Spying*. Cambridge, UK; New York: Cambridge University Press.
- Newburn, Tim and Stephanie Hayman. 2002. *Policing, Surveillance and Social Control : CCTV and Police Monitoring of Suspects*. Cullompton, Devon, UK; Portland, Or.: Willan Pub.
- North, Douglass. 1990. *Institutions, Institutional Change, and Economic Performance*. Cambridge, UK: Cambridge University Press.
- Olson, Mancur. 1971. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Cambridge, Mass: Harvard University Press.
- O'Neill, Onora. 1985. "Between Consenting Adults." *Philosophy & Public Affairs* 14(3):252–277.
- Orend, Brian. 2005. "War." *Stanford Encyclopedia of Philosophy* .
- Orwell, George. 2007. *1984*. New York: Signet Classics.
- Parent, W.A. 1983. "Privacy, Morality, and the Law." *Philosophy & Public Affairs* 12(4):269–288.
- Parfit, Derek. 1984. *Reasons and Persons*. Oxford [Oxfordshire]: Clarendon Press.
- Parfit, Derek. 2011. *On What Matters*. Vol. 1&2 Oxford; New York: Oxford University Press.
- Parker, Richard. 1974. "A Definition of Privacy." *Rutgers Law Review* 27:275–296.

- Pearson, James. 1975. "Oversight: A Vital Yet Neglected Congressional Function." *Kansas Law Review* 23:277–288.
- Pettit, Philip. 1988. "The Consequentialist Can Recognise Rights." *The Philosophical Quarterly* 38(150):42–55.
- Pettit, Philip. 1989. "Consequentialism and Respect for Persons." *Ethics* 100(1):116–126.
- Pettit, Philip and Geoffrey Brennan. 1986. "Restrictive Consequentialism." *Australasian Journal of Philosophy* 64(4):438–455.
- Pettit, Philip and Robert Goodin. 1986. "The Possibility of Special Duties." *Canadian Journal of Philosophy* 16(4):651–676.
- Philp, Mark. 2009. "Delimiting Democratic Accountability." *Political Studies* 57:28–53.
- Plutarch. 1951. *Selected Lives and Essays*. New York: Walter J Black, Classics Club.
- Pogge, Thomas. 1992. "Cosmopolitanism and Sovereignty." *Ethics* 103(1):48–75.
- Pogge, Thomas. 2002. "Cosmopolitanism: a Defence." *Critical Review of International Social and Political Philosophy* 5(3):86–91.
- Posner, Richard. 1981. *The Economics of Justice*. Cambridge, MA: Harvard University Press.
- Posner, Richard. 2006. *Not a Suicide Pact: The Constitution in a Time of National Emergency*. Oxford; New York: Oxford University Press.
- Posner, Richard. 2008. "Privacy, Surveillance, and Law." *University of Chicago Law Review* 75(1):245–260.

- Pound, Roscoe. 1915. "Interests in Personality." *Harvard Law Review* 28:343–456.
- Powers, Richard Gid. 1987. *Secrecy and Power: the Life of J. Edgar Hoover*. New York; London: Free Press ; Collier Macmillan.
- Radsan, John. 2007. "The Unresolved Equation of Espionage and International Law." *Michigan Journal of International Law* 28:596–623.
- Railton, Peter. 1984. "Alienation, Consequentialism, and the Demands of Morality." *Philosophy & Public Affairs* 13(2):134–171.
- Rawls, John. 1971. *A Theory of Justice*. Cambridge, Mass.: Belknap Press of Harvard University Press.
- Rawls, John. 1999. *A Theory of Justice*. Revised ed. Cambridge, Mass.: Belknap Press of Harvard University Press.
- Rawls, John. 2001. *Justice as Fairness: A Restatement*. Cambridge, Mass.: Harvard University Press.
- Rawls, John. 2005. *Political Liberalism*. Columbia Classics in Philosophy. expanded ed. New York: Columbia University Press.
- Rehfeld, Andrew. 2005. *The Concept of Constituency: Political Representation, Democratic Legitimacy, and Institutional Design*. Cambridge: Cambridge University Press.
- Richards, Neil. 2008. "Intellectual Privacy." *Texas Law Review* 87.
- Richelson, Jeffrey. 2012. *The US Intelligence Community*. Boulder, CO: Westview Press.

- Rose, Alexander. 2007. *Washington's Spies: The Story of America's First Spy Ring*. New York: Bantam Dell.
- Rourke, Francis. 1976. *Bureaucracy, Politics, and Public Policy*. New York: Little, Brown.
- Rousseau, Jean-Jacques. 2002. *The Social Contract; and, The First and Second Discourses*. New Haven: Yale University Press.
- Rubenstein, Jennifer. 2007. "Accountability in an Unequal World." *The Journal of Politics* 69(3):616–632.
- Ryle, Gilbert. 2009. *The Concept of Mind*. London; New York: Routledge.
- Sagar, Rahul. 2007. "On Combating the Abuse of State Secrecy." *The Journal of Political Philosophy* 15(4):404–427.
- Samuels, M. 1994. "Video Surveillance in Diagnosis of Intentional Suffocation." *The Lancet* 344(8919):414–415.
- Sanders, Lynn. 1997. "Against Deliberation." *Political Theory* 25(3):347–376.
- Scarre, Geoffrey. 1996. *Utilitarianism*. New York: Routledge.
- Schoeman, Ferdinand. 1984. *Philosophical Dimensions of Privacy: An Anthology*. Cambridge: Cambridge University Press.
- Scholz, John. 1991. "Cooperative Regulatory Enforcement and the Politics of Administrative Effectiveness." *American Political Science Review* 85:115–136.
- Searle, John. 1983. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge, UK: Cambridge University Press.

- Searle, John. 1992. *The Rediscovery of the Mind*. Cambridge, MA: The MIT Press.
- Searle, John. 1995. *The Construction of Social Reality*. New York, NY: The Free Press.
- Sepper, Elizabeth. 2010. "Democracy, Human Rights, and Intelligence Sharing." *Texas International Law Journal* 46:151–207.
- Shaw, William H. 1999. *Contemporary Ethics: Taking Account of Utilitarianism*. Malden, Mass.: Blackwell.
- Shipan, Charles. 2004. "Regulatory Regimes, Agency Actions, and the Conditional Nature of Congressional Influence." *American Political Science Review* 98:467–480.
- Shorrock, Tim. 2008. *Spies for Hire: the Secret World of Intelligence Outsourcing*. New York: Simon and Schuster.
- Shue, Henry. 1978. "Torture." *Philosophy & Public Affairs* 7(2):124–143.
- Shue, Henry. 1996. *Basic Rights: Subsistence, Affluence, and U.S. Foreign Policy*. Princeton, NJ: Princeton University Press.
- Sidgwick, Henry. 1891. *The Elements of Politics*. MacMillan.
- Sidgwick, Henry. 1981. *The Methods of Ethics*. Indianapolis: Hackett Pub. Co.
- Sinnott-Armstrong, Walter. 2011. "Consequentialism." *Stanford Encyclopedia of Philosophy*.
- Smart, J. J. C and Bernard Arthur Owen Williams. 1973. *Utilitarianism: For and Against*. Cambridge: Cambridge University Press.

- Smist, Frank. 1994. *Congress Oversees The United States Intelligence Community 1947-1994*. 2 ed. Knoxville: University of Tennessee Press.
- Smith, Adam. 1982. *The Theory of Moral Sentiments*. Indianapolis: Liberty Fund.
- Solove, Daniel. 2003. "Reconstructing Electronic Surveillance Law." *George Washington Law Review* 72:1264–1305.
- Solove, Daniel. 2008. "Data Mining and the Security-Liberty Debate." *University of Chicago Law Review* 75(1):343–362.
- Southall, D P and M P Samuels. 1993. "Ethical Use of Covert Videoing for Potentially Life Threatening Child Abuse: A Response to Drs Foreman and Farsides." *BMJ* 307(6904):613–614.
- Southall, D P and M P Samuels. 1995. "Some Ethical Issues Surrounding Covert Video Surveillance—A Response." *Journal of Medical Ethics* 21(2):104–115.
- Southall, David, Michael Plunkett, Martin Banks, Adrian Falkov and Martin Samuels. 1997. "Covert Video Recordings of Life-Threatening Child Abuse: Lessons for Child Protection." *Pediatrics* 100(5):735–760.
- Staples, William. 2007. *Encyclopedia of Privacy*. Westport, CT: Greenwood Press.
- Steunenberg, Bernard. 1992. "Congress, Bureaucracy, and Regulatory Policy-Making." *Journal of Law, Economics, and Organization* 8:673–694.
- Stich, Stephen and Jonathan Weinberg. 2001. "Jackson's Empirical Assumptions." *Philosophy and Phenomenological Research* 62:637–643.
- Stuntz, William. 1991. "Warrants and Fourth Amendment Remedies." *Virginia Law Review* 77(5):881–943.

- Swain, Stacy, Joshua Alexander and Jonathan Weinberg. 2008. "The Instability of Philosophical Intuitions: Running Hot and Cold on Truetemp." *Philosophy and Phenomenological Research* 76:138–155.
- Taylor, Charles. 1994. *Multiculturalism*. Princeton University Press.
- Thaler, Richard H and Cass R Sunstein. 2009. *Nudge: Improving Decisions about Health, Wealth, and Happiness*. New York: Penguin Books.
- Thomas, T. 1996. "Covert Video Surveillance—An Assessment of the Staffordshire Protocol." *Journal of Medical Ethics* 22(1):22–25.
- Thompson, Dennis. 1999. "Democratic Secrecy." *Political Science Quarterly* 114:181–193.
- Thomson, Judith Jarvis. 1975. "The Right to Privacy." *Philosophy & Public Affairs* 4(4):295–314.
- Thucydides. 1972. *History of the Peloponnesian War*. Harmondsworth: Penguin Books.
- Tsang, Steve, ed. 2007. *Intelligence and Human Rights in the Era of Global Terrorism*. Westport: Praeger.
- Vermeule, Adrian. 2007. *Mechanisms of Democracy: Institutional Design Writ Small*. New York, NY [etc.]: Oxford University Press.
- Vile, M.J.C. 1998. *Constitutionalism and the Separation of Powers*. 2 ed. Indianapolis: Liberty Fund.
- Waldron, Jeremy. 2012. *The Harm in Hate Speech*. Cambridge, MA: Harvard University Press.

- Waltz, Kenneth. 2001. *Man, the State, and War: A Theoretical Analysis*. Revised ed. New York: Columbia University Press.
- Waltz, Kenneth Neal. 1959. *Man, the State, and War: A Theoretical Analysis*. New York: Columbia University Press.
- Walzer, Michael. 1973. "Political Action: The Problem of Dirty Hands." *Philosophy & Public Affairs* 2(2):160–180.
- Warren, Samuel and Louis Brandeis. 1890. "The Right to Privacy." *Harvard Law Review* 4(5):193–220.
- Weingast, Barry. 1981. "Regulation, Reregulation, and Deregulation: The Political Foundations of Agency Clientele Relations." *Law and Contemporary Problems* 44:147–177.
- Weingast, Barry. 1984. "The Congressional-Bureaucratic System: A Principal-Agent Perspective with Applications to the SEC." *Public Choice* 41:141–191.
- Weingast, Barry and Mark Moran. 1982. "The Myth of Runaway Bureaucracy." *Regulation* 6:33–38.
- Weingast, Barry and Mark Moran. 1983. "Bureaucratic Discretion or Congressional Control? Regulatory Policymaking by the Federal Trade Commission." *Journal of Political Economy* 91:765–800.
- Welsh, B. C. and D. P. Farrington. 2003. "Effects of Closed-Circuit Television on Crime." *The ANNALS of the American Academy of Political and Social Science* 587(1):110–135.
- Westin, Alan. 1968. *Privacy and Freedom*. New York: Atheneum.

- Williams, Bernard. 1988. *Hare and Critics: Essays on Moral Thinking*. Oxford University Press chapter The Structure of Hare's Theory.
- Williams, Bernard. 2002. *Truth and Truthfulness: An Essay in Genealogy*. Princeton, N.J.: Princeton University Press.
- Wilson, James Q. 1975. "The Rise of the Bureaucratic State." *Public Interest* 41:77–103.
- Wilson, James Q. 1980. *The Politics of Regulation*. New York: Basic.
- Wood, Dan. 1988. "Principals, Bureaucrats, and Responsiveness in Clean Air Enforcement." *American Political Science Review* 82:213–234.
- Wood, Dan and James Anderson. 1993. "The Politics (or Non-Politics) of U.S. Antitrust Regulation." *American Journal of Political Science* 37:1–40.
- Yoo, John and Glenn Sulmasy. 2007. "Counterintuitive: Intelligence Operations and International Law." *Michigan Journal of International Law* 28:625.
- Zegart, Amy. 2009. *Spying Blind: The CIA, the FBI, and the Origins of 9/11*. Princeton: Princeton University Press.
- Zegart, Amy. 2012. *Eyes on Spies: Congress and the United States Intelligence Community*. Kindle ed. Stanford: Hoover Institution Press.