The Concept of Security

Uncertainty, Evidence and Value

Jonathan Charles Carter Herington

A thesis submitted for the degree of
Doctor of Philosophy
of
The Australian National University

September 2012

THE AUSTRALIAN
LIBRARY
NATIONAL UNIVERSITY

Statement of Originality

The work contained in this thesis is my own and, unless otherwise attributed, is not the result of joint effort or research. No material in this thesis has been previously submitted for the award of a degree at any institution of higher learning.



Jonathan Charles Carter Herington

Abstract

Typically we think of security as a good like any other, something with a fairly determinate content which is enjoyed by individuals or groups. There have been many conceptions of what this good consists in but little agreement. In this thesis I conduct an analysis of the concept of security, arguing that while conceptions of security disagree about the content of this good, they all rely on an underlying thin concept: a mode of enjoying that content 'securely.'

I provide an account what it means for an entity to enjoy a good securely from both a fact-relative and an evidence-relative perspective. From the fact-relative perspective I argue against a commonly held conception of secure enjoyment as protection from the interference of the powerful and instead suggest that it should be understood purely as the objective probability of enjoying the good in the future. From the evidence-relative perspective, I argue that we should understand the security of a good as the minimum degree of credence an agent may justifiably assign to enjoying that good. Securely enjoying a good therefore implies a reasonable guarantee of enjoying that good in the future, even in instances where the available evidence is limited or imprecise.

I argue that the secure enjoyment of goods, from both a fact-relative and an evidence-relative perspective, has an important role to play in moral decision-making. In particular, the importance an agent places on the fact-relative security of a good models the agent's attitude towards outcome risks. Likewise, in situations where there are multiple credence functions compatible with the evidence, placing special weight on the secure expected utility of an act, appears to model the importance of avoiding an epistemic risk. In this respect, the value of security may amount to ensuring the achievement of a minimally decent future despite the limits of the available evidence.

Table of Contents

Abstract iii

Table of Contents iv

Acknowledgements vi

Introduction 1

Motivations 2

Scope and goals 4

Thesis overview 7

Chapter One The History of Security 11

1.1 The Genealogy of Security 11

1.2 Traditional Accounts of Security 18

1.3 Critical Accounts of Security 22

1.4 Radical Accounts of Security 31

Chapter Two A Conceptual Framework 39

2.1 Essentially contested concepts 40

2.2 The Three Concepts of 'Security' 45

2.3 The Structure of Being Secure 51

Chapter Three Subjective and Objective Security 61

3.1 The Standard Account 61

3.2 External and Affective Security 63

3.3 Fact-relative, Evidence-relative and Belief-relative Security 64

3.4 The New Model of Objective and Subjective Security 71

Chapter Four Security and Possible Futures 79

4.1 Preliminaries 79

4.2 Promotional and Protective Security 85

4.3 Promotional Security: Chance 87

4.4 Protective Security: Robustness across Dispositions 92

Chapter Five Evidence and Security **99**

5.1 Evidence, Beliefs and Justification 100

5.2 First Moves 103

5.3 Limited and Imprecise Evidence 108

5.4 Robust Credence 111

Chapter Six Security and Goodness **117**

6.1 The Standard Teleological Model 118

6.2 Security and the Best Possible Future 120

6.3 Fact-relative Security and the Goodness of Acts 127

6.4 Evidence-relative Security and Evidence-relative Oughts 135

Conclusion **147**

Bibliography **151**

Acknowledgements

Whilst any omissions and errors are my own, this thesis would not have been possible without the assistance, intellectual support and patience of a great many people.

The intellectual environments I have encountered at ANU, first at the Centre for Applied Philosophy and Public Ethics and then within the Research School of Philosophy, have been invaluable. Discussions with faculty and students – whether at formal seminars, reading groups or the ritual of afternoon tea – have shaped this thesis in innumerable ways, and I regret that I cannot name everyone who has spent time with me discussing the ins and outs of security. Geoff Brennan, John Broome, Alan Hájek, John Kleinig, Seth Lazar, Michael Smithson, Nic Southwood and David Wiens all contributed in more particular ways – pointing out the limits of my thoughts and ways they could be improved.

My most sincere thanks must go to two individuals in particular. Christian Enemark, who first pushed me to think about the ethics of security, deserves special commendation for attempting to keep me grounded in the day-to-day realities of security. As an advisor and advocate he cannot be matched. Likewise, Michael Selgelid has my lasting gratitude for spending countless hours challenging my assumptions, testing the strength of my arguments and providing advice on the next book to read. Without his patience and supportive attitude, it is likely that I would have given up on the concept of security long ago.

Lastly, I have this thesis to thank for introducing me to the fiercest intellectual companion, and loving partner, that anyone might hope to find. My deepest gratitude goes to Rosa, for having the patience to listen to me explain my latest problem and for helping me solve a great many of them.

"Of course, we shall probably know more tomorrow, and it would be foolish then to disregard the new knowledge. But we cannot take advantage of the new knowledge in advance. We must cross that bridge when we come to it, and accept the consequences of our unfortunate epistemic situation with what composure we can find. Life is hard."

– T. Williamson, *Knowledge and its Limits*, p. 237

Introduction

The future is capricious. Human lives are full of uncertainties: both delightful and disastrous, personal and global. In much of the world, abundance, good health and stable government are stalked by the prospect of famine, disease and civil war. In the relative safety of the developed world we lose jobs, homes, friends and family to events which are neither intended nor foreseen. Terrorism, environmental degradation, pandemics and natural disasters beset the poor and the rich alike. We have little information about the timing and magnitude of many of these events, and yet we must plan our lives regardless. Given the potential for catastrophe, and limited evidence about its likelihood, it should come as no surprise that a large portion of the time, resources and attention of individuals and governments alike is devoted to obtaining a measure of security.

In this thesis I provide an analysis of the concept of security and why it appears to be an independently valuable end. Typically we think of security as a good with a fairly determinate content (e.g. protection from interpersonal violence, State stability etc.) which is enjoyed by individuals or groups. There have been many conceptions of what this good consists in (and whom enjoys it), but little agreement. I argue that while thick conceptions of security disagree about the content of this good, at their heart is a thinner notion of security as a quality of reliability. This thinner notion of security captures the sense in which – though an entity may enjoy a good by virtue of a calculated gamble, a serendipitous confluence of circumstance or through some procedure which provides it reliably – many of our plans depend upon, not the mere enjoyment of goods, but their secure enjoyment.

As this thesis progresses it will become clear that I view secure or reliable enjoyment as a form of protection against uncertain or indeterminate futures. Uncertainty about the future is an inescapable part of our lives, and it complicates and undermines our decision-making at every turn. Sometimes we have good evidence about the future, such that we can be relatively precise with our predictions, other times we remain completely ignorant. Securely enjoying a good suggests that, come

what may, we possess a reasonable guarantee that we will continue to enjoy that good into the future.

Motivations

Before beginning our discussion of the concept and value of security, we might ask a preliminary question: why should moral and political philosophers look at the *concept* of security? There are so many other worthwhile normative projects which could be undertaken on the topic of security. One might evaluate which security threats are particularly worthy of remedy, propose just policy options to address those threats or discuss the ethics of actual security practice. Analysing the concept of security could appear to be philosophical fiddling whilst the world burned. Contrary to this view, I believe that an analysis of the concept of security is fundamental to understanding security's normative significance. This is both because it is often unclear what is meant by the term 'security' and because it is nonetheless assumed to be especially valuable.

First, security is commonly held to possess special normative importance. This importance stems both from its power to politically legitimise action which would, in other respects, be impermissible and secondly, from its status as a genuinely weighty moral consideration. Throughout history, the invocation of security has been used to compel obedience to political authority, legitimise declarations of war and justify the exclusion (or extermination) of undesirable peoples. Enormous sums of public money are directed towards militaries, intelligence agencies and public order organisations based upon their putative role in preserving security. Normally inviolable rights have been swiftly suspended, with little opposition from citizens or courts, by governments who invoke the spectre of threats to national security.[1] Indeed, security could be described as the dominant political virtue of contemporary institutions.

Moreover, security is not just politically valuable, but also appears to be a genuinely weighty moral consideration. A broad sweep of political and moral philosophy considers some modicum of security to be essential to human flourishing. Henry Shue describes it as the most basic right, and the concept is central to historic political documents such as the Declaration of the Rights of Man of 1789 and the

[1] Constitutional courts, particularly those belonging to the common law tradition, are famously deferential to the executive's judgement in matters involving security. See *Liversidge v. Anderson*, AC 206(1942); *Ex parte Quirin*, 317 U.S. 1(1942); *Korematsu v. United States*, 323 U.S. 214(1944); *Hamdi v. Rumsfeld*, 542 U.S. 507(2004). See also the suspension clauses in otherwise strongly rights-oriented constitutions; *U.S. Constitution*, art. I, §9, cl. 2; *Basic Law for the Federal Republic of Germany*, art. 17a, cl. 2; *Constitution of Ireland*, art. 40.4.5°.

Universal Declaration on Human Rights.[2] Likewise, political philosophers consider the provision of security to be one of the key tasks of political institutions. Hobbes, Locke, and Mill all considered security to be one of the primary justifications for the interference of the state in the lives of its citizens.[3] Whilst security has not received the same degree of attention within contemporary political philosophy as the concepts of justice, liberty or equality, it retains a broadly acknowledged role as a weighty consideration in moral and political theorising.[4] Given the political power of security, and its acknowledged position as a genuinely weighty moral consideration, it should be clear that the concept is especially important to our moral and political lives.

Second, the meaning of 'security' is opaque. Security language can include claims about truth (e.g. "The United States is no longer secure"), value (e.g. "the security of our society is the most cherished aspect of our lives") and right action (e.g. "national security requires us to deny terrorism suspects the right to seek *habeas corpus*"). When we regard these statements which deploy the word 'security', however, we find that their content doesn't give us much insight into the actual epistemic and evaluative claims which the speaker is making. Consider, for example, the following declaration from a politician: "the security of our nation's nuclear missiles is threatened!" Whilst this seems like a relatively straightforward claim, the multiple meanings of the word 'security' make discerning the overall meaning of the statement difficult. Does our politician mean that the missiles are not well guarded? Or that they are improperly situated in an earthquake zone? Or does a peacenik politician plan to eliminate them as a cost-saving measure? Could a rogue commander use them without legal authorization? Indeterminacy over the meaning of 'security' obscures the epistemic and evaluative claims being made by actors who use the term.

Moreover, even though security is widely thought to be especially valuable, it is often unclear what kinds of goods are ultimately at stake when the term 'security' is invoked. Many of the goods which are explicitly mentioned in connection with

---

[2] Henry Shue, *Basic Rights: Subsistence, Affluence, and Us Foreign Policy* (Princeton University Press, 1980 [1996]); *Déclaration Des Droits De L'homme Et Du Citoyen De 1789*, (Assemblée nationale, 1790); *Universal Declaration of Human Rights*, G.A. res. 217A (III), art. 3.

[3] I will consider this in more depth in Chapter One. See Thomas Hobbes, "Leviathan," ed. J.C.A Gaskin (Oxford: Oxford University Press, 1996), pp. 133, chap. 17; John Locke, "Two Treatises of Government," (London: J. M. Dent & Sons, 1690 [1924]; reprint, 1990), pp. book 2, chap. 2, sec. 8; John Stuart Mill, "Utilitarianism," in *On Liberty and Other Essays*, ed. John Gray (Oxford: Oxford University Press, 1991), p. 190.

[4] See James Griffin, *On Human Rights* (Oxford: Oxford University Press, 2008), p. 33; Robert E. Goodin, *Political Theory and Public Policy* (Chicago: University of Chicago Press, 1982), pp. 220-41; Jeremy Waldron, "Security and Liberty: The Image of Balance," *The Journal of Political Philosophy* 11, no. 2 (2003); Jeremy Waldron, "Security as a Basic Right (after 9/11)," in *Global Basic Rights*, ed. Charles R. Beitz and Robert E. Goodin (Oxford: Oxford University Press, 2009).

security are valuable only in so far as they promote or maintain more fundamentally valuable goods. In the case of our security-conscious politician, for instance, it would be odd if she thought that nuclear missiles were valuable *qua* nuclear missiles (and even more odd if she stated that view publicly). Well-guarded nuclear missiles might be instrumental to the realisation of other, intrinsically valuable, goods – such as the wellbeing of the nation's citizens or the political independence of their state – only if certain contingent facts obtain. If these missiles do not, in fact, promote these final goods or there is some other more acceptable method of promoting these goods, then the security of those missiles may cease to have any value at all. The goods which a speaker wants to secure and, in exactly what way they are meant to be secure are not well defined by our current understanding of the concept of security. Security is, in this sense, like justice: a powerful word with an opaque meaning.

We should be in no doubt, given these observations, that security is a concept which demands attention from philosophers. The opacity of the concept should trouble those who are concerned to determine what is, and what is not, morally justified by the invocation of security. Likewise, the centrality of the term to debates in political theory and international politics means philosophers wishing to engage with these disciplines ought to have a clear understanding of its connection to moral and political philosophy. Yet, despite persistent reference to security within the literature on human rights and political philosophy, there has been little analysis of the concept itself by philosophers.[5] Work on the concept of security has largely been left to theorists of international politics and criminology, and the result has been a lack of interrogation into the grounds for the value of security. My hope is that by analysing the concept of security and identifying its underlying structure, we can begin to clarify security's role in moral decision-making.

Scope and goals

Whilst my overall goal is to provide an analysis of the concept of security, there are several constraints on the scope of this thesis.

First, because of the paucity of philosophical engagement with the concept of security, there is a great deal of conceptual ground clearing which must be undertaken. The conceptual home of security and the discipline where the concept has been given

[5] Pace Jeremy Waldron, "Safety and Security," *Nebraska Law Review* 85, no. 2 (2006).

the closest attention is International Relations.[6] Unfortunately, this literature often confuses or neglects simple distinctions between subjective and objective security, instrumental and final ends, and the normative and practical importance of security. In this sense, part of the contribution I hope to make with this thesis is to clear the conceptual ground and to identify what value commitments are really at issue within the contest over the meaning of 'security.'

Second, this thesis is not in the business of providing a full conception of what it means to *be* secure as a human being. Over the course of the twentieth century, there has been an enormous array of conceptions of what it means to be secure: from the narrow view that it merely consists in being free from intentional violence,[7] to the ultra-expansive view that it consists in being emancipated from any constraint on one's ability to live a meaningful life.[8] In Chapter One I will explore these many conceptions in order to glean some understanding of how they are underpinned by the concept of secure enjoyment, but I will not seek to provide any alternative to them. A philosophically defensible conception of security would require a level of engagement with the notions of basic goods, human rights and substantive conceptions of the good which this thesis simply cannot contain. Indeed, part of the lesson of this thesis is that thick conceptions of security may obscure a number of important trade-offs which are made between the enjoyment of a good and its secure enjoyment. I shall grapple with the difficult questions posed by secure enjoyment, and leave it to others to provide thick accounts of security.

Third, whilst at the conclusion of this thesis I will provide an account of secure enjoyment's role in moral decision-making, I will not provide an account of just how much weight it ought to be accorded in our all things considered judgements on the goodness of acts. It may be the case that acts intended to secure the enjoyment of a good will come apart from those acts which are intended to maximise the enjoyment of

---

6 For a selection of conceptual work in this discipline, see: David A. Baldwin, "The Concept of Security," *Review of International Studies* 23, no. 1 (1997); Ken Booth, *Theory of World Security* (New York: Cambridge University Press, 2007); Barry Buzan, *People, States and Fear: The National Security Problem in International Relations* (Sussex: Wheatsheaf Books, 1983); Barry Buzan, Ole Wæver, and Jaap De Wilde, *Security: A New Framework for Analysis* (London: Lynne Rienner Publishers, 1998); J Huysmans, "Security! What Do You Mean?: From Concept to Thick Signifier," *European Journal of International Relations* 4, no. 2 (1998); Bill McSweeney, *Security, Identity and Interests: A Sociology of International Relations* (Cambridge: Cambridge University Press, 1999); Emma Rothschild, "What Is Security?," *Daedalus* 124, no. 3 (1995); Arnold Wolfers, "National Security as an Ambiguous Symbol," in *Discord and Collaboration: Essays on International Politics* (Baltimore: Johns Hopkins Press, 1962).

7 Shue, *Basic Rights: Subsistence, Affluence, and Us Foreign Policy*, p. 14.

8 Ken Booth, "Security and Emancipation," *Review of International Studies* 17, no. 4 (1991).

a good. I will not be venturing a view on how to balance these two considerations, but will merely offer a view of the value considerations they represent.

Given my motivations for investigating the concept of security, and these self-imposed constraints, my conceptual analysis will be guided by the following set of goals. My analysis of security is not merely concerned with identifying some part of our conceptual space for the sake of completeness, rather I seek an account of the concept of security which serves some specific purposes.[9] Whilst there are many legitimate purposes for which an analysis of the concept of security might be crafted, three such considerations guide my analysis in this thesis.

First, I seek an analysis which is able to be integrated with theorising about the ultimate goals of our moral and political lives. As we identified above, we typically view security as especially valuable: perhaps as an end in itself. An analysis which identifies the concept's role within moral decision-making might help identify the grounds for security's special value. Whilst more post-structurally inclined scholars have recognised the need for an account of the value of security,[10] the rich tradition of Anglo-American moral and political philosophy has yet to engage fully with the topic of security. In seeking a definition of security which is able to be integrated with such theorising, we should attempt to make clear the grounds, if they exist, for the value of security.

Second, an analysis of the concept of security should be ecumenical towards the variety of different conceptions of the good. Security is a normatively thick concept, in the sense that conceptions of security instantiate a variety of different value judgements on top of the concept's common core. As we shall see, the contestation over the content of security is primarily a contestation over the entity whose security matters, and what kinds of goods constitute their security. Because such debates are primarily debates over what matters, an analysis of the concept of security should be ecumenical towards a variety of differing conceptions of the good.

Third, an analysis of the concept of security should be connected to our ordinary language. Appropriating the term 'security' requires me to be sensitive to its

---

[9] The notion that conceptual analysis of social and political terms should be grounded in a number of legitimate purposes, is borrowed from Sally Haslanger, "What Are We Talking About? The Semantics and Politics of Social Kinds," *Hypatia* 20, no. 4 (2005).

[10] Booth, Theory of World Security, p. 109; Anthony Burke, Beyond Security, Ethics and Violence: War against the Other (New York: Routledge, 2007); Michael Dillon, Politics of Security: Towards a Political Philosophy of Continental Thought (New York: Routledge, 1996).

extension in the world, even as I aim to be slightly revisionary in its meaning.[11] Thus, in analysing the concept I will seek to develop a broad understanding of the competing conceptions actually held by individuals in the world. By investigating these conceptions, a picture of the underlying structure of the concept may emerge which is not merely stipulative, but accords with our intuitions about the meaning of security. As we shall see, it is this underlying structure which helps explain the role that security plays in moral decision-making.

A complete argument for these three considerations would be extensive and is beyond the scope of this thesis. Suffice to say that I believe these purposes are intuitively appealing, *prima facie* plausible and ultimately defensible. My goal is therefore to analyse the concept security in a way which both explains security's role in our moral decision-making, is ecumenical towards an array of reasonable conceptions of the Good and accords with at least some of the ordinary uses of the term.

Thesis overview

This thesis can be broken into two parts. In the first part of the thesis I lay some important groundwork for thinking about the concept of security, identifying common conceptions of security and some important distinctions which have hitherto been overlooked. In the second part I provide an account of secure enjoyment from both a fact-relative and evidence-relative perspective. Ultimately I argue that the secure enjoyment of a good is a distinct consideration with respect to deliberations on how an agent ought to act.

In Chapter One I explore the many different meanings which have been ascribed to the word 'security.' The rich history of the word itself, in Classical and European thought, illuminates many interesting, and sometimes forgotten, conceptions of what it means to be secure. I delve into this history, not to show that security has one correct definition to which we should return, but to show that the word has several features which have become obscured and conflated over time.

In Chapter Two I establish a conceptual framework for thinking about security by distinguishing between the concepts of security practice, being secure, and secure enjoyment. I argue that whilst we sometimes view security as 'essentially contested'

---

[11] see Sally Haslanger, "Gender and Race:(What) Are They?(What) Do We Want Them to Be?," *Nous* 34, no. 1 (2000): p. 34.

we can, by distinguishing between these concepts and their various instances, make progress towards exposing the precise truth and value claims which are at stake when security is invoked. Moreover, I argue that conceptions of being secure – as a state, an individual or other kind of entity – posit that the entity must *securely enjoy* a set of specified goods in order to be secure. It is this structure, where a particular individual securely enjoys a set of goods, which is the hallmark of conceptions of security and which potentially grounds the special value of security.

In Chapter Three I explore the distinction between objective and subjective security, arguing that it obscures two important sets of distinctions. First, the objective-subjective distinction describes a distinction between the secure enjoyment of a set of concrete goods and the secure enjoyment of affective and psychological goods. Whilst some conceptions of being secure refer only to the secure enjoyment of external circumstances, others require that agents possess a "sense" of security. This can be distinguished from the view that the subjective-objective distinction refers to the perspective from which one judges whether a particular good is enjoyed securely. In particular, it appears that one can assess the secure enjoyment of a good from either a fact-relative, belief-relative or evidence-relative perspective. I posit that this more fine-grained set of distinctions allows us to build a model of the connection between the fact-relative, evidence-relative and belief-relative enjoyment of a set of goods, and the possession by an agent of the affective sense of security.

In Chapter Four I begin to give more content to the notion of fact-relative secure enjoyment by interrogating two potential conceptions of secure enjoyment. The promotional conception of security suggests that to securely enjoy a good is to have a high chance of enjoying that good in the future. The protective conception of security, on the other hand, suggests that to securely enjoy a good is to have a high chance of enjoying that good regardless of the dispositions of powerful agents towards you. Whilst the protective conception is intuitively appealing, I provide an error theory which suggests that the legitimate purposes of our conceptual analysis would best be served by adopting the promotional account.

In Chapter Five I take the promotional account from Chapter Four and seek to apply it to circumstances where our evidence about the future is limited or imprecise. I argue that given the available evidence, an agent may be justified in possessing a particular model of the future, which includes a set of epistemically possible futures and a credence function which assigns a probability to each future of being the actual future. Where the available evidence is compatible with multiple models of the future, then the degree to which a good is securely enjoyed is the minimum probability of

enjoying the good were any of these models the actual way in which the future was structured.

In Chapter Six I ask what role fact-relative security and evidence-relative security have in determining which acts an agent ought to perform. Initially assuming a standard model whereby the best act is the act which brings about the best possible future, I ask whether secure enjoyment contributes to the fact-relative goodness of a possible world. I argue that it is likely that it does not, but that, nonetheless it may bear on what an agent ought fact-relatively to do. In particular, I propose that considering the fact-relative security of goods amounts to considering the moral appropriateness of taking risks with respect to unfavourable outcomes. Placing weight on the evidence-relative secure enjoyment of a good, as opposed to the maximisation of its expected enjoyment, amounts to considering the moral appropriateness of taking risks with respect to unfavourable outcomes and with respect to limited or imprecise evidence.

Chapter One
The History of Security

Security language is opaque. When we hear or read the word 'security' we are often left to guess at its precise meaning: at the epistemic and evaluative claims that are being made by the author or speaker. In part, the opacity of the word 'security' is the result of a complex etymological history and a corresponding proliferation in contemporary definitions of what it means to be secure. In the following chapter, I delve into this history – not to argue that 'security' has one correct definition to which we should return – but to identify a set of common threads which run throughout the history of its meaning. Disentangling these threads will help to guide my analysis of the concept in later chapters.

I begin this chapter by interrogating the etymological genealogy of 'security,' tracing it back to two distinct Greek concepts, and their various translations into Latin and pre-Enlightenment English. I then show how one of these Greek concepts became the dominant understanding of security from the Enlightenment through to the modern era. In the second part of the chapter, I begin a critique of contemporary accounts of security. I categorise these accounts – most of which are concerned with the conduct of states in the international system – as traditional, critical or radical accounts of security. I conclude the chapter by cataloguing a number of themes which infuse accounts of the meaning of 'security': including a close connection to notions of necessity and reliability, a tension between the objective and subjective realisation of security, and an ambiguous attitude towards the value of security.

1.1 The Genealogy of Security

Contemporary usage of 'security' includes an incredible diversity of meaning.[12] 'Security' can denote a type of financial instrument, a psychological condition, systems of defence, a physical state of being, and much more besides. Whilst much of this

[12] At least 20 different definitions of the noun 'security', a further 21 definitions of the adjective/adverb 'secure', and 28 of the verb 'secure' can be found in the 2nd Edition of the Oxford English Dictionary.

contemporary diversity appeared only in the twentieth century, it is the product of a complex etymological history. In what follows I lay out a roughly chronological history of the word 'security', highlighting the major shifts in meaning and emphasis from its etymological antecedents in Greek and Latin through to the twentieth century.

The Ancients and *ataraxia*

The direct etymological antecedent of 'security' is the Latin *securitas*: which literally translates as "freedom from care" (from the phrase *sine cura*: without care).[13] Whilst there are some important ambiguities, Roman and Medieval usage of *securitas* referred primarily to a serene state of mind.[14] This state of mind was, in the pre-Christian era at least, intimately connected with a quiet and reflective life disconnected from the baser pursuits of politics, business and society. In this sense, *securitas* and the reflective life which accompanied it were not only seen as valuable but where so virtuous as to be what Cicero called the "object of supreme desire."[15]

This close connection between the early Roman usage of *securitas* and a serene disposition is a product of its association with the Greek concept of 'ἀταραξία' (*ataraxia*), which refers to an "impassiveness (or) calmness."[16] This concept is associated predominantly with the Epicurean tradition, which viewed the possession of a serene disposition as one of the prerequisites for *eudaimonia* ('flourishing').[17] The Epicureans viewed the attainment of *ataraxia* as a purely internal project, not only unbound from the concrete facts of an individual's circumstance, but hindered by too great a participation in civic life. *Ataraxia*, rather than something which could be furthered by living within a functioning political community, was to be found in detaching oneself from religious and political commitments and pursuing a life of quiet reflection.

The notion that there is a direct etymological lineage from the Greek concept of *ataraxia* to the Roman concept of *securitas* is attractive, but there are two important caveats to this simple story. First, *securitas* was not always associated solely with states of mind, but also with the concrete circumstances which made the attainment of

---

[13] *Oxford Latin Dictionary*, vol. 2 (Oxford: Oxford University Press, 1982).

[14] See J. Frederik M. Arends, "From Homer to Hobbes and Beyond - Aspects of 'Security' in the European Tradition," in *Globalization and Environmental Challenges: Reconceptualizing Security in the 21st Century*, ed. Hans Günter Brauch, et al., *Hexagon Series on Human and Environmental Security and Peace* (New York: Springer, 2008), p. 269.

[15] Rothschild, "What Is Security?," p. 61.

[16] Henry George Liddell and Robert Scott, *A Greek-English Lexicon*, ed. Henry Stuart Jones and Roderick McKenzie, 9th ed. (Oxford: Clarendon Press, 1940).

[17] Gisela Striker, "Ataraxia: Happiness as Tranquility," *The Monist* 73, no. 1 (1990).

those states of mind possible. Whilst in early Latin *securitas* straightforwardly referred to an inner calm, in later Roman usage (c. 1st-3rd century CE), *securitas* became associated with the *Pax Romana* which assured the physical safety and political liberty of Roman citizens.[18] For examples, the printing of coins declaring *"securitas publica"* and *"securitas perpetua"* during the tumultuous reign of Gallienus was part of a propaganda effort to remind citizens of the empire which underwrote their serene and carefree lives.[19] The Epicurean association between *securitas* and a retreat from the state thus began to dissolve.[20]

Second, *securitas* was not always viewed as a valuable trait. The literal Latin meaning of 'freedom from care' imbues the term with an ambiguous value: sometimes denoting an admirable quality of calm wisdom, other times denoting foolhardy self-assuredness or carelessness.[21] As the meaning of *securitas* shifted to denote the external circumstances which underpinned an individual's carefree state of mind, it further reinforced that an individual could be unjustifiably carefree. Indeed, early Christian usage of *securitas* referred to a sinful certitude in the face of God, and this meaning carried over into early English usage of 'security.'[22] Thus, as the connection between *securitas* and the external world began to strengthen, so did the term acquire a negative connotation.

Whilst these points complicate the picture, the meaning of *securitas* is dominated during the pre-Enlightenment period by an association with a sense of internal calm and freedom from fear. The ancient association between *ataraxia* and *securitas* inaugurates the long association between the concept of security and being free from fear. This is security as a psychological disposition – manifested as a defeasibly valuable state of mind that individuals hold largely without reference, or in opposition, to their external circumstances. Some of this sense remains in standard usage throughout the Latinate languages. In modern French, *sécurité*, principally denotes the feeling of being safe and the Spanish, *seguridad*, also holds a similar sense. Since the

---

[18] Arends, "From Homer to Hobbes and Beyond - Aspects of 'Security' in the European Tradition," p. 270.

[19] Gervase Mathew, "The Character of the Gallienic Renaissance," *The Journal of Roman Studies* 33(1943); Mason Hammond, "Res Olim Dissociabiles: Principatus Ac Libertas: Liberty under the Early Roman Empire," *Harvard Studies in Classical Philology* 67(1963).

[20] Arends, "From Homer to Hobbes and Beyond - Aspects of 'Security' in the European Tradition," p. 267.

[21] ibid., p. 269.

[22] see the use in Macbeth: "...security; Is mortals' chiefest enemy". William Shakespeare, "The Tragedy of Macbeth," in *William Shakespeare: The Complete Works (Compact Edition)*, ed. Stanley Wells and Gary Taylor (Oxford: Clarendon Press, 1988), pp. act 3, scene 5, line 31.

great works of Enlightenment political philosophy, however, the English language concept has lost much of this meaning. We now turn to consider this shift.

The Enlightenment and *asphaleia*

Although pre-Enlightenment usage of *securitas* is profoundly influenced by the concept of *ataraxia*, during the Enlightenment, *securitas* (and in English: 'security') came to be associated with an entirely different Greek concept: "ἀσφαλεία" (*asphaleia*). In the Greek, *asphaleia* is a negation of *sphallô*, which is a term associated with ancient wrestling, meaning "make to fall, overthrow, (properly) by tripping up."[23] Early Greek usage of *asphaleia* appears to track the literal meaning, implying steadfastness or the physical stability of an individual or object.[24] As Arends notes, however, *asphaleia* is also associated in Greek with the stability of city-states and empires, and during the Enlightenment the meaning of 'security' shifts away from its association with a state of mind and towards this more concrete understanding.[25]

This shift may owe much to the peculiar intellectual trajectory of Thomas Hobbes, and the profound influence of his thought on later political philosophers. Hobbes' first substantial work, written well before his influential political philosophy, was his 1628 English translation of Thucydides' *The History of the Peloponnesian Wars.* In the famous dialogue between the Athenians and the Melians, *asphaleia* is repeatedly used by Thucydides to denote the stability of the Athenian Empire. [26] The preoccupation of the Athenians is to avoid the bloody collapse of their empire and ultimately the 'overthrow' (here, a version of *sphallô*) of their city state by its former colonies in violent war.[27] In his 1628 translation of *The History*, Hobbes translated *asphaleia* as, variously: 'assurance,' 'protection,' 'safety' and 'security.' [28] The preoccupation with stability and protection within Thucydides work, particularly from physical violence and war, infuses the later Hobbes' work and his usage of the word 'security.'[29] In his first substantial original work, *The Elements of Law* (1640), we see the emergence of a political philosophy preoccupied with physical safety, which he equates with security:

---

[23] Liddell and Scott, *A Greek-English Lexicon.*

[24] Arends, "From Homer to Hobbes and Beyond - Aspects of 'Security' in the European Tradition," p. 265.

[25] ibid., p. 263.

[26] see, in particular, Thucydides, "The Peloponnesian War," pp. Book V, chap. 91-99.

[27] ibid., pp Book V, Chapter 99.

[28] See Thucydides, "History of the Peloponnesian War," In *Perseus Digital Library*, ed Thomas Hobbes. (1843), http://www.perseus.tufts.edu/hopper/text?doc=Perseus:text:1999.01.0247.

[29] See Richard Schlatter, "Thomas Hobbes and Thucydides," *Journal of the History of Ideas* 6, no. 3 (1945).

> The end for which one man giveth up, and relinquisheth to another, or others, the right of protecting and defending himself by his own power, is the security which he expecteth thereby, of protection and defence from those to whom he doth so relinquish it. And a man may then account himself in the estate of security, when he can foresee no violence to be done unto him...[30]

Thus, writing *Leviathan* in 1651, he supposes that the "end of common-wealth" is "security," – by which he meant the mechanism by which citizens get "themselves out from that miserable condition of war."[31] Furthermore, in Hobbes' subsequent Latin translation of *Leviathan* (published in 1668), he uses *'securitas'* to refer to just the same concept of physical safety. Hobbe's use of 'security' and *'securitas'* to denote physical safety is an important milestone in the history of the word, and highlights that for Hobbes 'security' (and *securitas*) refer to the Thucydidean concept rather than the Epicurean affect.

Hobbes' interpretation of security not only firmly establishes the primacy of the concrete over the psychological in the concept of security; it also inaugurates the Enlightenment belief that security can only be *guaranteed* by a political authority. For Hobbes, the legitimacy of the state is dependent on its ability to protect its citizens "from the invasion of foreigners, and the injuries of one another."[32] Hobbes famously contends that citizens seeking security can only achieve that end by subordinating their natural right to self-defence to the state.[33] Writing in a similar vein, Locke, in 1690, considers that the tie which binds men and women in political community is one which "secure(s) them from injury and violence."[34] Later figures, such as Condorcet, Paine and Rousseau, likewise saw security as a guarantee, given by the state, that citizens will be protected from violence against their person and property.[35] This is put into even starker terms by Leibniz, who *defines* the state "a great society of which the object is *la seureté commune* (the common security)."[36] The motif, running throughout Enlightenment political philosophy, is that a citizen should not merely be

---

[30] Thomas Hobbes, "De Corpore Politico," in *Human Nature and De Corpore Politico*, ed. J.C.A Gaskin (Oxford: Oxford University Press, 1994), p. 111, chap. XX.

[31] Hobbes, "Leviathan," p. 133, ch. 17.

[32] ibid., ch. II, 18, §13.

[33] For an excellent commentary on the moral basis of Hobbes' political theory, see Tom Sorrell, "Hobbes's Moral Philosophy," in *The Cambridge Companion to Hobbes's Leviathan*, ed. Patricia Springborg (Cambridge: Cambridge University Press, 2007).

[34] Locke, "Two Treatises of Government," p. 120 (book 2, chap. 2, sec. 8).

[35] Rothschild, "What Is Security?," pp. 63-65.

[36] Leibniz, writing in French, employs the use of *sûreté* to denote an objective state of being safe, rather than the psychological feeling of being safe denoted by *sécurité*. Leibniz writes: 'Ma definition de l'Estat, ou de ce que chez les Latins est appellé *Respublica* est: que c'est une grande societé don't le but est la seureté commune.' (My definition of the State, or of what the Latins call Respublica is: that it is a great society of which the end is common security) Gottfreid Wilhelm Leibniz, "Letter of 1705," in *Die Werke von Leibniz*, ed. Onno Klopp (Hannover: Klindworth, 1864-1873), p. 143.

free from moment to moment from violence: she should be *assured* that she will be safe. The association between 'security' and a guarantee is noted later by Bentham:

> ...man is not like the brutes, limited to the present time, either in enjoyment or suffering, but that he is susceptible of pleasure and pain by anticipation, and that it is not enough to guard him against an actual loss, but also to guarantee to him, as much as possible, his possessions against future losses. The idea of his security must be prolonged to him throughout the whole vista that his imagination can measure.[37]

By grasping the sense of security as an assurance or a guarantee from one entity to another, Bentham illustrates another important connection: that between security and the domain of rights.

In many Enlightenment and Industrial Era accounts, security is enshrined as a right which is one of the preconditions for the enjoyment of all other rights. As Condorcet puts it: the "natural rights of humanity" begin with "the security of one's person, a security which includes the assurance that one will not be troubled by violence, either within one's family or in the use of one's faculties."[38] For Mill, it was the case that security was "the most vital of all interests" since:

> ...on (security) we depend for all our immunity from evil, and for the whole value of all and every good, beyond the passing moment; since nothing but the gratification of the instant could be of any worth to us, if we could be deprived of anything the next instant by whoever was momentarily stronger than ourselves.[39]

The French revolution, echoing this concern for the neccesity of security, enshrined the rights to *"la liberté, la propriété, la sûreté et la résistance à l'oppression."*[40] Whereas the State guaranteed each individual's interest in *sûreté,* the wider freedom from anxiety denoted by *sécurité* was the preserve and responsibility of each individual.[41] For political philosophers of this era, *sûreté* is thus a fundamental objective of all individuals which can only be gained through the State. Rational actors seeking protection from the violence of others form the State to safeguard this fundamental interest. More importantly, the state is seen as *necessary* for the maintenance of security, in stark contrast to the inner calm of Roman *securitas* to which it is at best irrelevant or, at worst, an impediment. The necessity of the state to security (as *sûreté*) thus becomes enshrined within the concept itself.

---

[37] Jeremy Bentham, "Principles of the Civil Code," ed. John Bowring (1843), ch. VII.
[38] Condorcet quoted in Rothschild, "What Is Security?," p. 67.
[39] Mill, "Utilitarianism," p. 190.
[40] *Déclaration Des Droits De L'homme Et Du Citoyen De 1789,* art. 2.
[41] Rothschild, "What Is Security?," p. 63.

The modern period

Further changes in conceptions of security occur in the 18th, 19th and 20th centuries, guided as always by the political and historical context in which they flourished. The political philosophies of the Enlightenment period coalesced around a rigid form of rationalism; a tendency which may have excluded the original affective meaning of *securitas* in favour of a meaning of 'security' which referred to the objective characteristics of an individual's circumstance. So then did the increasing political importance of the state, particularly in the Napoleonic era, solidify security's close association with the power of the state and its instruments.[42] Furthermore, what initially began as an instrumental question – how best do we obtain/protect the security of individuals? – slowly morphed during the 19th century into the essentialisation of certain *practices* of security. As McSweeney notes regarding the use of the word 'security' in the past two centuries:

> The noun became attached to, and interchangeable with, property, land, money, fortifications – these things are said to have, or to be, 'security' – and to the means by which such things are made secure: armies, weapons.[43]

In McSweeney's view, the continued historical success of particular modes of creating physical security for citizens, meant that usage began to consider those modes (a sturdy gate, a well funded army, competent leaders, fiscal health) as synonymous with the word 'security.' No longer viewed as instruments for achieving security for an entity, the objects and practices of security, came to be seen, in some senses, as the essence of the concept of security itself. As the study of international relations became a formalised discipline in the early twentieth century, the association between the State, military power and security was further entrenched. By the mid twentieth century most scholars had a conceptualisation of security which either saw it as "a condition of the international community of states, deriving from interstate cooperation and the essential interdependence of international relations" or "a property of the state, deriving from the balance of power consequent upon state behaviour."[44]

Reflecting upon the genealogy of 'security' we can identify four fundamental notions which seem to reside at the heart of the concept. The first is a tension between the subjective and objective aspects of security. Early association with the subjective and

---

[42] McSweeney, *Security, Identity and Interests: A Sociology of International Relations*, pp. 17-18; Rothschild, "What Is Security?," p. 64.

[43] McSweeney, *Security, Identity and Interests: A Sociology of International Relations*, p. 18.

[44] ibid., p. 28.

affective concept of *ataraxia* gave way during the Enlightenment to the objective and corporeal concept of *asphaleia.* Whilst a concern for the subjective aspects of security has never disappeared, it has been subsumed by conceptions of security as freedom from violence and physical danger. This tension will re-emerge as we explore the different contemporary accounts of security which place different emphasis on security as a state of mind and security as a concrete state of being.

The second is the connection between security and future reliability. Hobbes work established the idea that security was not a good which could be obtained moment to moment, but was something that we had to have reliably. For Hobbes and later Enlightenment figures it was the state who provided such reliability. As we engage with contemporary conceptions of security, we will find that these too rely upon an implicit notion of reliability.

The third is the close association between security and necessity. Whenever security has been invoked, from the Epicurean concept to modern theories of rights, it has been accompanied by the idea that the content of the concept is necessary for the relevant entity to flourish. Whilst rights theorists associate the necessity of security with the flourishing of individuals, its provision is also seen as necessary for the state. For Enlightenment scholars, states which do not provide security to their citizenry are illegitimate, and hence the security of the state itself, its protection from war and foreign domination, is seen as a necessary component of being a state.

The fourth is a departure from viewing security as a good with ambiguous moral value, to viewing it as an indefeasible (sometimes overriding) good. Whilst the pre-Enlightenment association with 'carelessness' and sinful certitude imbued the word with a negative connotation in certain contexts, Hobbes association between security, necessity, and the state rehabilitates security. No longer was there the possibility of a disjunction between an individual's state of mind and their external state of affairs, security just was the external state of affairs. In this sense, the concept of security loses, under the Hobbesian and Enlightenment interpretation, the ambivalence it once bore in the Roman and early Christian eras. As we shall see, these four notions (subjectivity and objectivity; reliability; necessity, and; ambiguity over value) are also at the heart of contemporary re-conceptualisations of the meaning of security.

1.2 Traditional Accounts of Security

Contemporary Security Studies offers us a wealth of definitions of 'security.' Since the end of the Cold War, International Relations has spawned a small cottage industry

aimed at defining, conceptualising, critiquing and deconstructing security.[45] In this section I will broadly categorise contemporary definitions as either traditionalist, critical, or radical definitions of security.[46] In what follows, I begin by exploring some of the traditional and critical accounts of security. Whilst there is considerable diversity amongst these definitions, they almost all provide an account of what it means for a particular type of entity to be secure. Indeed, what is perhaps most striking about these accounts is that they all conceive of security as a state of being where some set of conditions *reliably obtain.* These accounts of security then differentiate themselves in three ways. The first is the type of entity whose security the conception describes: typically the state, societies, or individual human beings. The second is the conditions which the conception suggests should obtain for a referent object (i.e. Australia) to be considered secure as an entity of the specified type (i.e. a state). The final axis of differentiation is the emphasis on the objective realisation of security versus the subjective sense of security for a referent.

My investigation of contemporary accounts of security will finish with a consideration of some of the more radical approaches to the concept of security: particularly those held by social constructivists, post-structuralists and so-called Critical Theorists. In contrast to the traditional and critical strands of theorising, such accounts typically don't forward a substantive conception of the conditions for security. These radical definitions are therefore sometimes difficult to precisely characterise, but they provide important insights into a strand of security theorising which focuses on the things which are done (or the objects which are used) in the name of security.

Traditionalists

Security Studies arose as a sub-branch of the study of International Relations, and as such has traditionally been concerned with the relationships between States. For theorists of security in this mould, the primary question is how the State, as an institution, survives in an anarchic world where the threat of war is ever-present. While other branches of International Relations investigate the diplomatic process, international political economy and the role of political systems in State behaviour,

---

[45] For an excellent overview and a proto-sociology of academic Security Studies since 1945, see Barry Buzan and Lene Hansen, *The Evolution of International Security Studies* (Cambridge: Cambridge University Press, 2009).

[46] This categorisation violates some of the longstanding labels applied by Security Studies scholars themselves. The disciplinary movement known as Critical Security Studies, for instance, normally understands itself to include most of the conceptions I have termed 'radical' and few of those I term 'critical'.

Security Studies is primarily concerned with the "the study of the threat, use, and control of military force" by and against States.[47] According to this traditional account, the word 'security' refers to the security of the State, and consists only in being free of the threat of military force from other States. This tight focus on just one part of State functioning forms the basis of the traditional account of security within Security Studies: what is sometimes labelled 'national security.'[48]

Although the basis of national security is freedom from military interference, some have sought to expand the notion to include the state's freedom from other threats:

> ...national security is freedom from foreign dictation.[49]

> ...national security must be defined as integrity of the national territory and of its institutions[50]

> ...a threat to national security is an action or sequence of events that (1) threatens drastically and over a relatively brief span of time to degrade the quality of life for the inhabitants of a state, or (2) threatens to significantly narrow the range of policy choices available to the government of a state or to private, nongovernmental entities (persons, groups, corporations) within the state.[51]

> The objective of national security is to sustain freedom from foreign dictation and improvement of living standard in an environmentally sustainable fashion.[52]

Thus, the scope of national security might be expanded to include threats to political autonomy, the state's populace and, at the limit, its natural environment. In essence, these are claims about what kinds of goods are necessary for the state to function as a state. These expansions do not, however, change the fundamental association between the word 'security' and the security of the State.

One important dilemma which definitions of national security face is the extent to which they include, as Ullman does, consideration of "the inhabitants of a state."[53] There seems to be something perverse in seeking to secure a set of goods for the state's sake alone. The state is an institutional framework whose agency is a matter of deep controversy and its unclear why we should think its security (however defined) is

---

[47] Stephen M. Walt, "The Renaissance of Security Studies," *International Studies Quarterly* 35, no. 2 (1991): p. 212.

[48] On the ambiguity of the phrase 'national security' see Wolfers, "National Security as an Ambiguous Symbol,".

[49] Harold D. Lasswell, *National Security and Individual Freedom* (New York: McGraw-Hill, 1950), p. 51

[50] Hans Morgenthau, *Politics among Nations: The Struggle for Power and Peace*, 4th ed. (New York: Alfred A. Knopf, 1965), p. 562.

[51] Richard H. Ullman, "Redefining Security," *International Security* 8, no. 1 (1983): p. 133

[52] Joseph J. Romm, *Defining National Security: The Nonmilitary Aspects* (New York: Council on Foreign Relations, 1993), p. 85.

[53] Ullman, "Redefining Security," p. 133.

morally important without some account of its relationship to moral agents. In this sense, the traditional account of security frequently implies another account: which is the freedom of individuals from actual or threatened physical violence perpetrated by other individuals. Henry Shue, in his study of the right to security, takes individuals to be secure when they are not "subjected to (or threatened with) murder, torture, mayhem, rape, or assault".[54] This is a right which is held by an individual against other individuals, and protected by the state in the mode of a diffuse public good.[55]

Although definitions of national security make it appear as if they view the state as an end in itself, a more charitable reading is that such definitions view the security the state provides as *necessary* for the realisation of individual form of security. Even Hobbes, the pre-eminent champion of State authority, thought that the State's ultimate end was not self-perpetuation but the protection of its citizens from violence and coercion.[56] The plausibility of this position rests on three claims. The first is the common notion that "unless men enjoy some measure of security against the threat of death or injury at the hands of others, they are not able to devote energy or attention enough to other objects to be able to accomplish them."[57] The second is the Hobbesian claim that outside of the protection of a state, nefarious contemporaries, bandits and the soldiers of other states consistently threaten violence upon individuals. Lastly, the only stable solution to the violence and insecurity of the state of nature is the protection afforded by a political authority like the state. If these three claims hold true, then the State's protection from military attack and foreign political domination could be plausibly held to be necessary to the security of individuals.

If there is a unifying theme to traditional accounts of national and individual security, it is that they are concerned with freedom from harms intentionally inflicted by other agents. For individuals, security is protection from assault and murder, for the State this is protection from military conflict and the violation of its territorial integrity. The intentionality which characterises these threats is crucial to the traditional account of security, and is commonly held to ground a sharp distinction between safety and security.

---

[54] Shue, *Basic Rights: Subsistence, Affluence, and Us Foreign Policy*, p. 20.

[55] For a discussion on the notion of security as an individual or diffuse good in Shue, see Waldron, "Security as a Basic Right (after 9/11),".

[56] See the discussion in Glyn Morgan, *The Idea of a European Superstate: Public Justification and European Integration* (Princeton: Princeton University Press, 2007), p. 98.

[57] Hedley Bull, *The Anarchical Society: A Study of Order in World Politics* (London: Macmillan, 1977), p. 5.

> Safety is a more general concern than security because safety requires prudent protection from all probable dangers, whereas *security* is protection from dangers arising from the illegal aggression of others.[58]

Whilst the distinction between intended and unintended harms need not carry moral weight, when used to bound the concept of security, which is said to be of special importance, it is frequently taken to imply that intentionally inflicted harms are more important than unintentionally inflicted harms. In the next section, we will see how some critical accounts of security reject the intentionality distinction, and suppose that security consists in a broader set of goods than freedom from intentionally caused violence.

1.3 Critical Accounts of Security

At the conclusion of the Cold War, a sense of dissatisfaction over traditional conceptions of security settled over the policy and academic community. The end of the threat of nuclear annihilation and the disintegration of many post-Soviet states exposed what many saw as the narrow focus of traditional conceptions and prompted a wholesale critique of the traditional conception of security. These critical efforts focused on a number of questions around the "broadening and deepening" of the concept of security.[59] From these critiques, at least three major changes have been suggested to the traditional conception of security:

1. Shift the focus from the security of the State to another entity
2. Broaden the set of goods which constitute security
3. Emphasise the subjective realisation of security

These three axes, individually or in combination, provide the basis of critiques of the traditional accounts of security.

From the state to other entities

The first major critique sought to re-orient the meaning of 'security' away from considering only the security of the state and towards considering the security of other entities. This critique operates in two strands, each of which seeks to question the tight association between security and the state.

The first strand argues that though the state may be necessary to individual security, individual security is not reducible to state security. As we saw above, many traditional accounts reduce 'security' down to state security because, it is argued, the

[58] Naomi Zack, *Ethics for Disaster* (Lanham: Rowman & Littlefield Pub Inc, 2009), p. 91.
[59] Booth, *Theory of World Security*, p. 149.

state is indispensable to protecting the security of individuals. However, as Buzan notes, although "the security of individuals is inseparably entangled with that of the state" we should not assume that it always provides security to individuals.[60] The security of the state may be in direct opposition to the security of at least some individuals: such as when state institutions suppress the citizenry in order to prevent revolution or political change. Reducing 'security' to state security, as the traditional conception does, obscures the serious harm that some states do to their citizens, and the fact that "even those which are producers of security represent the means and not the ends."[61] This strand does not deny that state security is a constitutive meaning of 'security', just that we should recognise individual citizens as the "fundamental referent" or "subject" of 'security.'[62] In this sense, whilst the state might remain necessary to individual security, this fact does not totally obviate the need to talk about the security of individuals.

The second strand is those who want to challenge the idea that the state is necessary for the security of individuals. The traditional account, following the social contract model of Hobbes, Locke and other Enlightenment political philosophers, suggests that the horror of the state-of-nature compels us to form at least a minimal state and renders inconceivable any other form of social organisation.[63] Although the necessity of the state is generally popular, it is by no means clear that the state *is* empirically necessary for individual human flourishing. Anarchistic communities, such as the Catalonian anarchist movement during the 2nd Spanish Republic and the Kibbutz movement in Palestine, have proved stable for short periods. As such, anarchist political philosophers have argued that a State which monopolises violence may not only be unnecessary, but also unjust.[64] Likewise, in the discipline of International Relations, work by Booth, Wyn-Jones, Linklater and others inspired by the Frankfurt School of social theory have argued that human flourishing might best be served by "emancipatory political communities."[65] While the State remains the dominant force in global politics, and the predominant mode of ensuring individual security, it is nonetheless a contingent entity.

---

60 Buzan, *People, States and Fear: The National Security Problem in International Relations*, p. 21.

61 Booth, "Security and Emancipation," p. 320.

62 See RBJ Walker, "The Subject of Security," *Critical Security Studies: Concepts and Cases* (1997); Buzan, Wæver, and De Wilde, *Security: A New Framework for Analysis.*

63 Robert Nozick, *Anarchy, State, and Utopia* (Malden: Blackwell Publishing, 1974), p. 52.

64 See Robert Paul Wolff, *In Defense of Anarchism* (Berkeley: University of California Press, 1998).

65 Booth, *Theory of World Security* ; Andrew Linklater, "Dialogic Politics and the Civilising Process," *Review of International Studies* 31(2005); R Wyn Jones, "Security, Strategy, and Critical Theory," *London: Lynne* (1999).

For these two reasons, there has been a shift towards accounts of 'security' which endorse the possibility of referring to the security of a plethora of entities (e.g. individuals, communities, society, humanity, economies, states). The UNDP's concept of "personal security", for instance, is defined by reference to instances where "human life is...threatened by sudden, unpredictable violence."[66] These physical threats might include:

> Threats from the state (physical torture); threats from other states (war); threats from other groups of people (ethnic tension); threats from individuals or gangs against other individuals or gangs (crime, street violence); threats directed against women (rape, domestic violence); threats directed at children based on their vulnerability and dependence (child abuse); (and) threats to self (suicide, drug use).[67]

Whilst the security of the state and the security of the individual are thus decoupled under this definition, the focus remains on the traditional constitutive good of security: freedom from violence. As many have noted, however, once we shift away from the close connection between security and the state, the focus on intentionally inflicted violence seems inadequate. That a person is secure when they are free from intentional violence, even if they are threatened by a large number of environmental and structural dangers, seems intuitively wrong. The intuitive inadequacy of this account, particular in the domain of individuals, has motivated some to suggest that a broader set of threats and goods, not just intentionally inflicted violence, should be included in the concept of security.

Expanding threats, broadening goods

The second major critique seeks to broaden the range of goods which are required for security. Perhaps the most influential and widely cited contribution to this movement has been Barry Buzan's *People, States and Fear*:

> An enormous array of threats, dangers and doubts loom over everyone, and although the better-off can distance themselves from some of these (starvation, preventable/curable disease, physical exposure, criminal violence, economic exploitation, and such like), they share others equally with the poor (incurable disease, natural disasters, nuclear war), and create some new ones for themselves because of their advantages (air crashes, kidnappings, diseases of excessive consumption, and so forth).[68]

Such an expansion of the goods which constitute security has inspired a proliferation of work on environmental security, food security, the security of health, and even

---

[66] United Nations Development Program, "Human Development Report 1994," (New York: United Nations Development Program, 1994), p. 30.

[67] ibid., p. 30.

[68] Buzan, *People, States and Fear: The National Security Problem in International Relations*, p. 19.

cultural security.[69] The justification for this expansion is varied, but two major strands can be observed. The first is a concern with how traditional security is affected by such threats. For instance some scholars note that the spread of diseases amongst a states citizenry might create tension with neighbouring states.[70] In these accounts of security the entity remains the state, and its security remains defined by the content of a traditional account: it is merely the kinds of threats to the military security of the state which have expanded.

The second strand is concerned with the deficiencies of the traditional account in capturing what constitutes a secure state. In this case, freedom from pandemic diseases, a stable climate and a prosperous economy are viewed as constitutive of the security of the state itself. Likewise, recall Ullman's suggestion that state security is constituted by the maintenance of its inhabitant's quality of life.[71] This expansion in the constitutive goods of security seems to be predicated on expanded notions of which goods are necessary for the functioning of the state.

Some have argued that such an expansion of the goods which constitute security make the concept "inscrutable" and a "loose synonym for 'bad things that can happen.'"[72] In particular, these critics argue, the broadening move creates an equivalency between threats: "if we expand the idea of national security in its own right, then global warming becomes a threat to security that is to be avoided as much as, say, nuclear war."[73] Whilst not all security threats should be treated equally, including a wider array of goods in our definition of security does not necessitate an equivocation between threats to those goods. Furthermore, to say that a particular good is constitutive of security does not presuppose the size of its contribution to security. Some goods, such as bodily integrity, may contribute more heavily to the realisation of security than others, such as a clean environment. The move towards broadening the set of goods which are required for security cannot be resisted on the

---

[69] D. Deudney, "The Case against Linking Environmental Degradation and National Security," *Millenium: journal of international studies* 19, no. 3 (1990); L. Elliott, "Imaginative Adaptations: A Possible Environmental Role for the UN Security Council," *Contemporary Security Policy* 24, no. 2 (2003); and Christian Enemark, *Disease and Security: Natural Plagues and Biological Weapons in East Asia* (New York: Routledge, 2007); Sandra J. Maclean, "Microbes, Mad Cows and Militaries: Exploring the Links between Health and Security," *Security Dialogue* 39, no. 5 (2008); and Buzan, Wæver, and De Wilde, *Security: A New Framework for Analysis*, pp. 119-39.

[70] See, for instance, the work in health security by Harley Feldbaum, "U.S. Global Health and National Security Policy," (Washington: Centre for Strategic and International Studies, 2009); Susan Peterson, "Epidemic Disease and National Security," *Security Studies* 12, no. 2 (2002).

[71] Ullman, "Redefining Security," p. 133.

[72] Keith Krause, "The Key to a Powerful Agenda, If Properly Delimited," *Security Dialogue* 35, no. 3 (2004); Roland Paris, "Human Security: Paradigm Shift or Hot Air?," *International Security* 26, no. 2 (2001).

[73] Romm, *Defining National Security: The Nonmilitary Aspects*, p. 83.

grounds that it implies an *equivalence* between the various goods.[74] If these goods have a distinctive character which unifies them in some way, perhaps that they are necessary to the entity under consideration, then it seems natural to extend our account of the security of that entity to include them.

Subjective security

Finally, there has been a move toward greater recognition of the subjective dimension of security. Recalling the Greek notion of *ataraxia*, some have suggested that security should not be understood solely in terms of the concrete circumstances faced by individuals, but also as the states of mind those individuals form. An association between insecurity and fear is common, and its absence is seen as constitutive of the value of security.[75] Following this line of thinking, fear of a particular threat is noted by some authors as a key determinant of whether a particular issue is likely to be regarded as a threat to security.[76]

Although it is commonplace to cite the link between the subjective and objective components of security, few have attempted to provide a substantive conceptualisation of the subjective realisation of security. One exception is work by Aries and Leaning, who suggest that although "minimum material inputs to sustain survival" are an important component of human security, little attention has been paid to the importance of "core psychosocial supports."[77] These core psycho-social supports are ultimately constituted by "a sustainable sense of home and safety...a network of constructive social or family support...(and) an acceptance of the past and a positive grasp of the future."[78] Reliably possessing these psychological goods is critical to the ability of individuals to cope with physical vulnerability and constitutive of their security. Similarly, the work of sociologist Anthony Giddens has been particularly important in furthering the claim that security relates fundamentally to states of mind. Giddens develops the concept of 'ontological security' to explain the fundamental need

---

[74] There is, however, a pragmatic worry, which is the extent to which including more goods within an account of security undermines the usefulness of the category in decision-making and policy advocacy. See, for instance, the criticism of the UNDP's account of 'human security' in Nicholas Thomas and William T. Tow, "The Utility of Human Security: Sovereignty and Humanitarian Intervention," Security Dialogue 33, no. 2 (2002).

[75] See Robert E. Goodin and Frank Jackson, "Freedom from Fear," *Philosophy and Public Affairs* 35, no. 3 (2007); Waldron, "Safety and Security," p. 315.

[76] Enemark, *Disease and Security: Natural Plagues and Biological Weapons in East Asia.*

[77] Jennifer Leaning and Sam Arie, "Human Security: A Framework for Assessment in Conflict and Transition," (Washington: United States Agency for International Development/Complex Emergency Response and Transition Initiative, 2000), p. 16.

[78] ibid., p. 18.

of individuals to have "a sense of continuity and order in events."[79] This is, at heart, a strategy of coping with the chaos and unquantifiable uncertainty which infuses our social world by "bracket(ing)...possible events which could threaten the bodily or psychological integrity of the agents".[80] Without such a strategy, we are overwhelmed by the multitude of catastrophic possibilities which permeate our world. As Mitzen elaborates:

> Ontological security is security not of the body but of the self, the subjective sense of who one is, which enables and motivates action and choice. To say that individuals need security of this self is to say that their understandings of it must be relatively stable...Individuals value their sense of personal continuity because it underwrites their capacity for agency.[81]

Indeed, Giddens views the rationality of agents as not given by nature, but rather as a product of practices which generate their ontological security. Without ontological security, we lose the stable basis by which we can make rational determinations, and are consumed by the "deep, incapacitating state of not knowing which dangers to confront and which to ignore."[82] Although Giddens' approach is designed to explain individual behaviour, some have suggested that a drive for ontological security infuses not only individual lives but also underpins state action.[83] According to these authors, states operate on the basis of routinised understandings of themselves and their adversaries in order to bracket off certain possibilities from the decision-making process. States, in this sense, seek certainty about themselves and the political environment, mirroring the pursuit of subjective security by individuals. In essence, the approaches of Learning, Arie, Giddens and Mitzen seek to identify the value of security with states of mind. In doing so, they rehabilitate an important psychological aspect of the concept of security.

Broad Human Security – Three axes in unison

Drawing these three axes of critique together is the notion of human security, which finds its first clear expression in the 1994 UNDP Human Development Report:

---

[79] Anthony Giddens, *Modernity and Self-Identity: Self and Society in the Late Modern Age* (Cambridge: Polity Press, 1991), p. 243.

[80] ibid., p. 36.

[81] Jennifer Mitzen, "Ontological Security in World Politics: State Identity and the Security Dilemma," *European Journal of International Relations* 12, no. 3 (2006): p. 344.

[82] ibid., p. 345.

[83] See ibid; Brent J. Steele, *Ontological Security in International Relations: Self-Identity and the Ir State* (New York: Routledge, 2008).

> Human security can be said to have two main aspects. It means, first, safety from such chronic threats as hunger, disease and repression. And second, it means protection from sudden and hurtful disruptions in the patterns of daily life – whether in homes, in jobs or in communities.[84]

The UNDP embraces the shift from the state towards the individual, a broadening of the goods which are required for security, and recognition of the subjective realisation of security. The UNDP report, although essentially a political document, inaugurated a small debate amongst development economists and security theorists. Although a myriad of definitions have been provided by scholars, I will focus on one particular definition as being an exemplar of the field.[85] For Alkire: "the objective of human security is to safeguard the vital core of all human lives from critical pervasive threats, in a way that is consistent with long-term human fulfilment."[86] This illuminates two key questions for human security: what is the *vital core* of human lives? and, what is it to *safeguard* that core from critical and pervasive threats?

The first question interrogates the broadness of the set of goods which are necessary for human flourishing. This is the central debate within human security, and pivots on the extent to which security should emphasise the UN goals of universal "freedom from fear" or "freedom from want." Proponents of a narrow perspective on the set of goods, advocate limiting human security to physical violence against the individual. Proponents of a broad perspective are concerned with expanding the set of goods beyond physical violence to include the social, political and economic goods they view as the necessary for human flourishing. On some broad accounts, these goods are purely the objective goods required for human functioning, whilst others include certain states of mind. An example of a synthesis of these two considerations is the account held by Waldron, for whom security:

> ...comprises protection against harm to one's basic mode of life and economic values, as well as reasonable protection against fear and terror, and the presence of a positive assurance that these values will continue to be maintained into the future.[87]

---

[84] United Nations Development Program, "Human Development Report 1994," p. 23.

[85] See, for instance, Fen Osler Hampson et al., *Madness in the Multitude: Human Security and World Disorder* (Ottawa: Oxford University Press, 2002); Gary King and Christopher J.L. Murray, "Rethinking Human Security," *Political Science Quarterly* 116(2001); Leaning and Arie, "Human Security: A Framework for Assessment in Conflict and Transition," ; Edward Newman, "Critical Human Security Studies," *Review of International Studies* 36, no. 01 (2010); Taylor Owen, "Human Security-Conflict, Critique and Consensus: Colloquium Remarks and a Proposal for a Threshold-Based Definition," *Security Dialogue* 35, no. 3 (2004); Thomas and Tow, "The Utility of Human Security: Sovereignty and Humanitarian Intervention,".

[86] Sabina Alkire, *A Conceptual Framework for Human Security: Working Paper 2*, Crise Working Paper (Oxford: Centre for Research on Inequality, Human Security and Ethnicity (CRISE), 2003), p. 2.

[87] Waldron, "Safety and Security," p. 320.

Waldron's insistence that security includes an aspect of assurance relates to our second question: what is it to safeguard or protect these core values? Whilst there is considerable nuance to this question, and it is one which I will pursue in detail throughout this thesis, I believe that at its core, what it means to safeguard those goods is to have them obtain *reliably* across a wide range of possible futures. Indeed, it seems that part of what is distinctive about human security, as opposed to human wellbeing or flourishing, is that it connotes an assurance, or guarantee, that certain goods will be reliably accessible to the entity.[88] This was recognised in the original UNDP report, which made a clear distinction between development and security:

> Human development is...a process of widening the range of people's choices. Human security means that people can exercise these choices safely and freely-and that they can be relatively confident that the opportunities they have today are not totally lost tomorrow.[89]

This distinction reinforces the idea that the concept of security relates to the reliability of necessary goods. Indeed, greater specificity regarding the meaning of 'protection' reinforces this connection.

> Security is not synonymous with the average level of future well-being, but instead focuses on the risks of being severely deprived. My security today is not only a function of my well-being today, but also the prospects of avoiding states of great deprivation in the future...We define an individual's human security as his or her expectation of years of life without experiencing the state of generalized poverty.

Although I ultimately think the notion of reliability has a more complex relationship with uncertainty than is captured by the concept of 'expectation,' this is a plausible redescription of the concept of human security. Indeed, as this thesis will bear out, I believe that the notion of reliability (and cognate concepts, such as robustness) lie at the core of security.

Security as a thin concept

Although the bulk of scholarly work in recent years has been devoted to quite thick accounts of security, which include reference to a number of determinate goods, there has been some notable work on security as a thin concept. This strand of theorising on security is small, but it attempts to capture the sense in which:

---

[88] Indeed, those accounts of human security which include a very broad set of goods in their conception are prone to the critique that they equate human security with human wellbeing. See Krause, "The Key to a Powerful Agenda, If Properly Delimited,".

[89] United Nations Development Program, "Human Development Report 1994," p. 23.

> 'Security' is not something we can have more of or less of, because it is not a thing at all. It is...the name we use for a temporally extended state of affairs characterized by the calculability and predictability of the future.[90]

In this vein, perhaps the foundational account of a thin concept of security is Arnold Wolfers', who simply defines security as both the "absence of threats to acquired values" and the "absence of fear that such values will be attacked."[91] This, in and of itself, doesn't seem to give us much to go on; but it lays the foundations of a thin concept of security. In particular, it highlights that the subjective and objective realisation of security is as important to the thin concept as it is to thick accounts. Moreover, it illustrates the sense in which the value of security might not be something which can be evaluated in the abstract, but is dependent on the goods which are secure.

Two recent works have expanded on these themes. Graham Smith considers security to be a relational term, which does not have "value in and of itself" and "has meaning only insofar as it presupposes values which are considered ends in themselves."[92] Security in a political context then becomes a practice of "orders" whose members "seek to realise their values (and especially core values), and to protect the institutions and arrangements which they understand to best preserve and promote those values."[93] In this sense, the value of security is related to whether or not "just" orders are secured.[94] In a similar vein, Waldron (who forwards his own thick account of security, detailed below) recognises that the word security has an "adjectival" meaning – wherein it is understood to be "a mode of enjoying other goods" rather than as a good in and of itself.[95] This mode is constituted by an assurance that the goods held today will not be gone tomorrow. In discussing Bentham's conception of security, Waldron suggests that:

> To sustain security...it is not enough that threats of this kind be repelled. There must be an assurance that they will be repelled, an assurance that people can count on and build upon in advance of the outcome of any particular attack.[96]

---

[90] Mariana Valverde, "Governing Security, Governing through Security," in *The Security of Freedom: Essays on Canada's Anti-Terrorism Bill*, ed. Ronald J. Daniels, Patrick Macklem, and Kent Roach (Toronto: University of Toronto Press, 2001), p. 85.

[91] Wolfers, "National Security as an Ambiguous Symbol," p. 150.

[92] GM Smith, "Into Cerberus' Lair: Bringing the Idea of Security to Light," *British Journal of Politics & International Relations* 7, no. 4 (2005): p. 488.

[93] ibid., p. 488.

[94] Paul Roe, "The 'Value' of Positive Security," *Review of International Studies* 34(2008)

[95] Waldron, "Safety and Security," p. 318.

[96] ibid., p. 317.

Moreover, this assurance isn't just "a matter of probability" it also includes "confidence" in the goods it protects.[97] Whilst there is little explicit philosophical work on the thin concept of security beyond Waldron, there are a host of related concepts, such as robustness, stability and resilience which have been given thorough philosophical treatments. Chief amongst these is Philip Pettit's concept of modally robust goods, which we will deal with in great detail in Chapter Four.[98] As we will explore in Chapters Two, Four and Five, characterising this quality of reliability or assurance is not an easy task, but it lies at the heart of thick accounts of security, and hence is vital to understanding security's value.

1.4 Radical Accounts of Security

These critical conceptions of security are a crucial step forward, but for some theorists, they do not go far enough. In particular, the notions of reliability and necessity which underlie the traditional and critical accounts are precisely those things which are seen as problematic by those with more radical aims. So far, we have focused on contemporary accounts of security which lend themselves to re-description in terms of the reliability of a set of goods for a particular entity. There are, however, a number of contemporary definitions which are altogether different to the standard concept.

Emancipatory security

Whilst many conceptions of security view it as a defeasible good which can be weighed against other goods (notably liberty), there are some conceptions of security which embrace it, when conceptualised in a particular way, as an all-encompassing good. In particular, scholars influenced by the Frankfurt School of social theory (so-called Critical Theory) have endeavoured to develop a conception of security as an emancipatory political process.[99] The leading figure in this movement is Ken Booth,[100] for whom this emancipatory process can be defined as:

> ...the securing of people from those oppressions that stop them carrying out what they would freely choose to do, compatible with the freedom of others....Emancipation is the philosophy, theory, and politics of inventing humanity.[101]

---

[97] ibid., pp 318-19.

[98] Pettit characterisation of these is, as yet, unpublished. His most expansive discussion is given in Philip Pettit, "Making Good: The Challenge of Robustly Demanding Values," *Uehiro Lectures*(2011), http://www.practicalethics.ox.ac.uk/lectures/2011_uehiro_lectures.

[99] See Booth, *Theory of World Security*; Wyn Jones, "Security, Strategy, and Critical Theory,".

[100] See Booth, "Security and Emancipation,"; Booth, *Theory of World Security.*

[101] Booth, *Theory of World Security*, p. 112.

Although emancipation owes much to the Marxist tradition, Booth associates it with a Kantian account of autonomy.[102] In this sense he conceives of security as a *means* towards the end of an emancipatory political order which safeguards an individual's autonomy. Thus, although Booth conceives of emancipation as a process, it also points towards an ideal, and thus straddles accounts of security as a state of being and security as a political practice.[103] If we were to redescribe emancipatory security in the terms of a state of being, it would likely boil down to a concern for the reliability of an individual's capability for autonomous choice, particularly around the central project of realising their identity. The more security for any one individual: the more robust their capability for autonomous, undetermined action and identity-formation. From this perspective, emancipatory security seems merely to be a version of human security – one which places all of its emphasis on the autonomy of individuals.

The emancipatory account of security is peculiar, however, because it conceives of security as *always* valuable. Emancipation is the sole goal of a Critical Theory approach to moral and political philosophy and encapsulates the theory's vision of a just political order.[104] It thus subsumes underneath emancipation other things which we normally think of as intrinsically and incommensurably valuable – such as equality, wellbeing and the satisfaction of desert claims. Although Booth considers security to be constitutive of emancipation, he associates it so tightly (as "two sides of the same coin") with emancipation that it loses independent meaning.[105] Security, on Booth's account, just is an emancipatory political order and we therefore lose much of the need for a separate concept at all. If we are to abandon the idea that security is a separate concept, however, it renders common value judgements regarding security practices implausible. For example, there is a *prima facie* tension between the values of security and liberty which, on the emancipatory account becomes merely an empirical matter of determining what level of autonomy for each is compatible with the same level of autonomy for all.[106] If we come to view security *just as* the project of human emancipation, we lose sight of the possibility that pursuing that project might require us to make value tradeoffs. Furthermore, others have noted that the emancipatory security project has difficulty in providing an internally coherent account of when violence, even in self-defence, is justified: making its application to questions regarding

---

[102] ibid., p. 113.
[103] ibid., p. 114.
[104] This order might ultimately be structured by a Habermasian dialogic ethic. See Linklater, "Dialogic Politics and the Civilising Process,".
[105] Booth, "Security and Emancipation," p. 319; Booth, *Theory of World Security*, p. 107.
[106] On the tension, see, for instance Waldron, "Security and Liberty: The Image of Balance,".

the justifiable use of force problematic.[107] The potential tradeoffs against other things which we hold to be valuable, such as equality or wellbeing, are made impossible by the insistence that security (as emancipation) is the only political project which has intrinsic value.[108]

Constructivist accounts

Within the emancipatory security literature there is a strong emphasis on the processes of security, and this suggests an altogether different reading of the concept of security. One approach to security is the idea, forwarded primarily by the so-called Copenhagen School, that 'security' does not refer to a state of being at all. Utilising a broadly constructivist approach, such scholars suggest that "'security' has become the indicator of a specific problematique, a specific *field of practice.*"[109] This field of practice is concerned with the way specific issues become 'securitised' by a polity, and the kinds of measures they use to secure a referent. According to securitisation theory, when an actor invokes 'security,' they typically do so in order to legitimise the use of emergency measures. As Buzan and Wæver note, securitisation occurs when:

> [A]n intersubjective understanding is constructed within a political community to treat something as an existential threat to a valued referent object, and to enable a call for urgent and exceptional measures to deal with the threat.[110]

The Copenhagen School utilises J.L. Austin's speech act theory to animate this claim, focusing on the idea that speaking security is an act, where "something is done."[111] In the case of a securitising speech act, the intent (in speech act parlance: the illocution) is the legitimation of exceptional measures in order to combat the threat to the referent. The speech act is successful if the claim is accepted by its audience; thus socially constructing what is – and by omission, what is not – a security issue.[112] The concept

---

[107] See Columba Peoples, "Security after Emancipation? Critical Theory, Violence and Resistance," *Review of International Studies* 37, no. 3 (2011).

[108] While the emancipatory account is untenable if security is synonymous with the end of emancipation, it may forward an important empirical claim: that emancipatory politics is the most fruitful way of creating a sustainable condition of security for the global population. Understood in this way, the Critical Theory approach avoids the conceptual problems above and instead forwards an interesting political theory. Taken in such a way, we are free to consider the concept of security separately, while potentially endorsing the empirical claim that an emancipatory political order would also secure individuals.

[109] Ole Wæver, "Securitization and Desecuritization," in *On Security*, ed. Ronnie D. Lipschutz (New York: Columbia University Press, 1995), p. 50.

[110] Barry Buzan and Ole Wæver, *Regions and Powers: The Structure of International Security* (Cambridge: Cambridge University Press, 2003), p. 491.

[111] Buzan, Wæver, and De Wilde, *Security: A New Framework for Analysis*, p. 26; J. L. Austin, *How to Do Things with Words*, ed. J. O. Urmson and Marina Sbisà, 2nd ed. (Oxford: Oxford University Press, 1980), p. 94ff.

[112] Buzan, Wæver, and De Wilde, *Security: A New Framework for Analysis*, p. 25.

of security is thus defined by its relationship to 'exceptional measures', and in particular is demarcated by a zone of 'special politics':

> "Security" is the move that takes politics beyond the established rules of the game and frames the issue either as a special kind of politics or as above politics.[113]

Rather than offering a conceptualisation of security as a state of being for some entity, securitisation theory captures a quality of the acts performed in security's name (e.g. emergency measures) and uses instances of their use to define the set of security issues.

Whilst this approach has been enormously influential, it has also raised some problems for those interested in the normative dimension of security practice. Firstly, merely because an actor possesses the political legitimacy and wherewithal to apply security practices is no reason to suppose the soundness of their claim that the security of a particular referent is threatened.[114] Moreover, some of the most insecure individuals and groups do not possess the required political or rhetorical legitimacy to make such security claims and have them accepted by an audience.[115] In effect, securitization theory risks becoming hostage to the claims made by powerful political actors because it lacks a clear standard from which to question the validity of their securitizing moves or to argue for the inclusion of the insecurities of the less powerful.

Postmodern accounts: Neocleous, Dillon, Burke

Sharing the view that security refers to a specific kind of practice is recent post-modernist work. Again, the term 'security' refers to a specific kind of politics:

> (Security is) an interlocking system of knowledges, representations, practices and institutional forms that imagine, direct and act upon bodies, spaces and flows in certain ways – to see security not as an essential value but as a *political technology*.[116]

On this approach, security is seen as a mode of thinking and of acting and reaches beyond specific speech acts to infuse the very logic which governs all state action. In its political role, it is a politics of emergency which is grounded in the prerogatives of the state.[117] Although it is unclear whether these authors wish to collapse security down to practice alone or whether they think it can also refer to a state of affairs, there is a sense in which they deploy the thin concept of security discussed above. In particular, much of their criticism of security practices rests on the idea that security is

---

[113] ibid., p. 23.

[114] Buzan et al. admit as much in ibid., p. 206.

[115] See the critique of the Copenhagen School offered by Lene Hansen, "The Little Mermaid's Silent Security Dilemma and the Absence of Gender in the Copenhagen School," *Millenium: journal of international studies* 29, no. 2 (2000).

[116] Burke, *Beyond Security, Ethics and Violence: War against the Other*, p. 28.

[117] Mark Neocleous, *Critique of Security* (Montreal: McGill-Queen's University Press, 2008), p. 42.

a state of certitude. Dillon's work in particular deploys the idea of security as a process of *fixing* reference points (whether they be value judgements or metaphysical doctrines) to bound our philosophical and political reflections.[118] Similar strands of thought run through Neocleous and Burke's concern that alternative forms of political community are excluded by the role of security within our political discourse.[119]

Although the relationship of these scholars to the concept of security is complex and ambiguous, a common current is the questioning of the value of such a project. On Burke's analysis, the discourse of security "animates forms of geopolitical control, desire and violence" while simultaneously stifling "alternative possibilities of society, justice and existence."[120] For Neocleous, the move to broaden the notion of security "is to hand (the social order) over to the key entity which is said to be the ground of security, namely the state" without interrogating whether security was "little more than a semantic and semiotic black hole allowing authority to inscribe itself deeply into human experience." [121] The political technology of security, as currently constituted, thus becomes an inherently disvaluable process.

The work of scholars on the harmful components of the politics of security reminds us that security shouldn't be valued above all else. Yet, the wholesale dismissal of security, particularly by Neocleous, seems too quick: outside the rarefied language of the academy, citizens conceive of and talk about security as a positive aspect of their existence. Although the post-modern analysis of the practice of security illuminates important ways in which it closes off potentially fruitful political possibilities, it ignores the potential value of closing off some possibilities – those which are catastrophic. Rather than dismiss security as valueless, with an eye to the dangers of over-emphasising being secure (and a critical stance towards security as a political practice), we might best be served by viewing security as a defeasible good – that there is a presumption in favour of security, but that it can be overridden by countervailing considerations.

Conclusion

This chapter has endeavoured to provide a history of the concept of security. While the history of the concept cannot provide us with definitive reasons for accepting one

---

[118] Dillon, *Politics of Security: Towards a Political Philosophy of Continental Thought*, p. 14.

[119] In particular, security politics is seen as "anti-political", because it makes certain possibilities unthinkable. See Burke, *Beyond Security, Ethics and Violence: War against the Other*, p. 54; Neocleous, *Critique of Security*, p. 185.

[120] Burke, *Beyond Security, Ethics and Violence: War against the Other*, p. 53.

[121] Neocleous, *Critique of Security*, pp. 3-4.

account over another, it does point to a number of themes which intersect the concept of security.

Firstly, security is not always valuable. The early usage of the term was infused with an ambiguity over its value. During the Enlightenment, this ambiguity was lost as security came to be seen as necessary, both to individuals and also for the legitimation of the state. Contemporary thick accounts of security reflect their author's specific moral and political commitments, but overwhelmingly view security as an extremely important good. Indeed, on many of these accounts, maintaining security is seen as overridingly important, regardless of the other values at stake. One of the main aims of this thesis is to reject this view of the value of security, in favour of the view that its promotion is dependent on what other values are at stake. In Chapter Six, I will argue more specifically for the value of security and its role in our all things considered judgements about the goodness of acts.

Secondly, thick accounts of security make claims regarding the goods which are necessary for an entity. These constraints are specified as criteria (perhaps a specific amount of a good, or protection from a specific threat) which are required – by practical necessity, morality or political claims – for a referent to function as that type of entity. These claims, particularly about what goods are necessary to be human, are ultimately mixed up with moral claims, and, if we are serious about an ecumenical concept of security, it seems we should be in favour of moving discussions about the desirability or necessity of particular goods away from narrow discussion on the concept of security and back towards the domain of moral and political philosophy.

Thirdly, security may be realised either subjectively or objectively. For some, an entity is secure only in so far as it actually (objectively) possesses the requisite goods specified by a substantive account. For others, the subjective sense of security is foundational. For still others, some mixture of the two is required for an entity to be secure. In Chapter Four, I will investigate how we might formalise the relationship between subjective and objective security.

Fourthly, it may be the case that many accounts of security will turn out to be compatible with one another: describing distinct conceptual spaces. Many radical concepts of security focus on security practices and these appear to be importantly distinct to security as a state of being. Likewise, many thick accounts of security for a particular entity (e.g. the State) appear to be compatible with thick accounts of security for a different entity (e.g. citizens). As I will explore in Chapter Two, 'security' might therefore refer to several distinct concepts.

Finally, and most importantly, there is a close relationship between the notions of security and reliability. Indeed, the idea of a thin concept of security, of what it means to hold a good securely, seems intimately tied to the idea of reliability, assurance and guarantee. Unfortunately, conceptual and philosophical analysis of what it means to enjoy a good securely is sparse. As Chapter Two will introduce, this might be particularly troubling because it appears that many contemporary definitions of security as a state of being rely on an implicit deployment of what it means to enjoy something securely.

Chapter Two
A Conceptual Framework

One of the goals of this thesis is to clarify the internal structure of the concept of security. As I suggested in the Introduction, the political power of the word 'security' and its weight as a moral consideration is too great for us to be guessing at the actual claims being expressed by political actors and theorists. Additionally, as Chapter One illustrated, the use of the word 'security' can obscure an incredible array of empirical and evaluative claims. We thus need some mechanism with which to clarify the claims made within each account of security and to identify why they appear to be so weighty. With a clear picture of the internal structure of the concept we might be able to identify a ground for the special value of security. In this sense, clarifying the structure of the concept of security is a key task in the broader project of determining what the value of security might justify.

Whilst in the previous chapter I identified a number of common themes within historical and contemporary accounts of the meaning of 'security,' this is not a substitute for cataloguing their structure and substantive content systematically. Merely having an understanding of the various accounts of 'security' on offer does not alleviate the opacity of the term. This chapter establishes a new framework for thinking about the concept of security so as to further clarify some of this opacity. First, I lay some analytic groundwork by introducing a nuanced account of 'essentially contested concepts.' In particular, I argue that identifying the common core of essentially contestable concepts helps elucidate the internal structure and implicit valuations within conceptions. Second, I apply this nuanced view of essential contestability to security. I argue that the word 'security' is poly-semantic and denotes at least three distinct and compatible concepts: security as a practice, security as a state of being for an entity, and security as a quality of reliability. Third, I fix on the concept of security as a state of being as one of the most normatively important concepts. I argue that instances of security as a state of being are, most likely,

essentially contestable. Finally, I introduce one of the central contributions of this thesis: the notion that substantive conceptions of security as a state of being are constituted by reference to the security (reliability) of certain conditions, and that the concept of security as a quality of reliability constitutes the core of the concept of security as a state of being.

2.1 Essentially contested concepts

How then should we think of the concept of security? Whilst little conceptual analysis of security has been undertaken by philosophers, the dominant approach within Security Studies has been to view it as a single 'essentially contested concept.'[122] In this sense it is claimed to be similar to other value-laden concepts – such as health, freedom, justice, race or gender – whose proper definition is contested by a number of competing conceptions (see Figure 1, p. 58). As a way of thinking about concepts, essential contestability can help identify the ways in which different value-commitments constitute different understandings of a concept. Thus, it might provide a framework for thinking about what is really at stake amongst the varying definitions of 'security' – whether they conflict, whether they are compatible, and, most interestingly of all, whether they rely upon a shared conceptual structure. The literature on essentially contested concepts is complex and extensive, but unfortunately, few scholars of the concept of security have bothered to engage with it in any systematic fashion.[123] Below I review some important distinctions which arise from the literature on essentially contested concepts, in order to apply a more nuanced understanding to the word 'security.'

The notion of an essentially contested concept was first outlined by W.B. Gallie, who suggested that there are some terms (such as "work of art" or "democracy") for which "there is no one clearly definable general use...which can be set up as the correct or standard use."[124] Contested concepts arise when a number of different definitions are applied to a single term and where the users of these definitions claim that their definition "is the correct or proper or primary, or the only important, function which the term in question can plainly be said to fulfil."[125] Essentially contested concepts are

---

[122] See Buzan, *People, States and Fear: The National Security Problem in International Relations*, p. 6; Steve Smith, "The Contested Concept of Security," in *Critical Security Studies and World Politics*, ed. Ken Booth (Boulder: Lynne Rienner Publishers, 2005), pp. 27-28.

[123] Very limited exceptions are Baldwin, "The Concept of Security," p. 10; McSweeney, *Security, Identity and Interests: A Sociology of International Relations*, pp. 83-84.

[124] W. B. Gallie, "Essentially Contested Concepts," *Proceedings of the Aristotelian Society* 56(1955): p. 168.

[125] ibid., p. 168.

those where this dispute is irresolvable by virtue of characteristics or features of the concept itself. Gallie proposes that essentially contested concepts have five features:

(I) The concept must be "appraisive." It must signify or accredit something valued.
(II) The concept must be "internally complex," even if it is valued as a whole.
(III) The concept must be "initially variously describable." Any explanation of the concept's value "must therefore include reference to the respective contributions of its parts or features," where the importance of its component parts may legitimately be ordered in various ways.
(IV) The concept must be "open in character." It thus should be able to accommodate modification in the light of changed circumstances, and this modification shouldn't be able to be predicted.
(V) The various users of the concept recognise that they are competing over the same concept.[126]

The heavy emphasis on value-commitments within these conditions have lead most commentators to suggest that a concept is essentially contested, rather than merely contingently contested, when "rival uses of it express conflicting moral and political commitments between which reason cannot arbitrate."[127] On a straightforward reading of Gallie's account then, we have good reason to suppose that security is an essentially contested concept. It signifies something valued,[128] is internally complex and its constituent parts are held to contribute to its value in a diverse set of ways. It is also able to accommodate modification in virtue of changed circumstances (i.e. technological changes such as the introduction of nuclear weapons). Finally, echoing commentators on Gallie, definitions of security are seen as "derivative" of conflicting visions of the "character and purpose of politics."[129] I argue, however, that this is a simplistic understanding of the essential contestability of 'security' and that a more nuanced view is possible. In particular, since Gallie's original paper there have been a suite of modifications to his formulation of essentially contested concepts, clarifying its content and implications.

---

[126] ibid., pp 171-72.
[127] John N. Gray, "On the Contestability of Social and Political Concepts," *Political Theory* 5, no. 3 (1977): p. 334.
[128] cf. Baldwin, who suggests that security cannot be appraisive because it is not the sole or primary goal of states. This seems to me a misreading of (I), since something can obviously be valued even if it isn't lexically prior to all other values or is sometimes traded for some other value. See Baldwin, "The Concept of Security," pp. 10-12.
[129] Booth, *Theory of World Security*, p. 109.

First, clarification of (V) highlights that terms "which are simply radically confused" (e.g. catachresis, malapropisms), "whose uses conceal a diversity of distinguishable concepts" (e.g. homonyms, polysemes), or which refer to vague concepts are not contested concepts.[130] To say that an individual is tall is to use a vague concept, and the phrase "the singer literally brought the house down" is to engage in a form of idiomatic catachresis. Polysemy, however, is perhaps the most important (and deceptive) trait for the theorist of contested concepts. To call someone a 'child' illustrates polysemy well, since the word can refer to someone of a younger age, somebody's offspring or an immature person. Each of these meanings deploy distinct though related concepts, and are not conceptions of a single disputed concept of child. Moreover, for a definition to count as a conception of a concept it must purport to define the concept itself and not some instance of the concept. Thus, if we consider an account of brotherly love alongside an account of romantic love, it should be clear that they are not conceptions of the concept of love *per se* but accounts of different instances of a shared supra-concept. There is no contest between the concepts of brotherly and romantic love unless each is taken to properly define love in and of itself.[131]

The notion of a contested concept thus relies on a distinction between a concept, *m*, and conceptions of *m*. The concept *m* is contested if there is more than one conception of *m* (m_1, m_2, m_3, etc.) and each of these conceptions claim to properly define *m*.[132] Conceptions of a contested concept therefore all point to the same 'common core:' which might be a shared schema, form, or paradigmatic exemplar which all conceptions take the concept to include.[133] The concept of 'race' is a paradigmatic example of a contested concept since it has many different conceptions – e.g. race as skin colour, race as lineage, race as socio-cultural construct, etc. – and these are all claimed to define the same thing.

Second, merely because a concept is contested does not imply that the concept itself has the quality of *essential contestability*. Recall that essentially contested concepts, according to Gallie, are single concepts whose deployment "*inevitably* involves endless

---

[130] Gray, "On the Contestability of Social and Political Concepts," p. 337; Barry Clarke, "Eccentrically Contested Concepts," *British Journal of Political Science* 9, no. 1 (1979).

[131] cf. The dialogue between Aristophanes (190a) and Socrates (201d) in Plato's *Symposium*.

[132] The distinction between conception and concept, which is ultimately rejected in favour of ideals and conceptions, comes from Christine Swanton, "On the "Essential Contestedness" of Political Concepts," *Ethics* 95, no. 4 (1985).

[133] Some contested concepts, like justice, may not even be said to have common cores, instead the conceptions wrestle over the *ideal* of justice itself. See ibid., p. 818

disputes about their proper uses on the part of their users."[134] Most nuanced understandings of essential contestability therefore make a distinction between actually contested concepts and concepts that are, by their nature, essentially contestable.[135] Many concepts are contested even when conceptualisers agree on the criteria by which conceptions should be judged because of, for instance, a lack of evidence or an unwillingness on the part of conceptualisers to accept available evidence. The proper definition of many concepts in natural science (e.g. electron, virus, acceleration of gravity, etc.) were contested at one point but, by virtue of a shared set of criteria of what would constitute the best conception, these contests have been resolved. Essentially contested concepts, however, are supposed to necessarily involve an irresolvable conflict.

Third, the exact nature of this conflict is disputed, but there is general agreement that it derives from some combination of the opacity of the concept and the appraisive claim which the concept makes.[136] For concepts which satisfy (II) and (III) their internal structure is opaque either because there is no fact of the matter regarding the internal structure of the concept or because the structure is epistemically inaccessible.[137] This can result in a conflict amongst conceptions even in the absence of an appraisive aspect to the concept (e.g. the concept of biological species). The special claim about essentially contested concepts is that they have an appraisive quality (and thus satisfy (I)). This gives the contest a special character: it becomes a contest over value commitments. Most contemporary moral and political philosophers hold that there are multiple reasonable interpretations of what is good, and no decisive criteria for distinguishing the best conception of the good.[138] This general feature of our moral landscape imposes itself on concepts which satisfy (I), (II) and (III); meaning that different conceptions of the good will provide different accounts of the contribution of a concept's internal parts to its overall appraisive claim. Essential contestability then, as a property of a concept, can be seen as a susceptibility to the irresolvable contest between conceptions of the good.

Fourth, this last point might lead commentators to surmise that, to the extent that the influence of a substantive conception of the good can be extracted or

---

[134] Gallie, "Essentially Contested Concepts," p. 169.
[135] Clarke, "Eccentrically Contested Concepts," p. 124.
[136] See Swanton, "On the "Essential Contestedness" of Political Concepts," p. 820.
[137] ibid., p. 814.
[138] This is one of the central motivations for contemporary political philosophy in the liberal tradition. See J Rawls, *A Theory of Justice*, Rev. ed. (Cambridge: Belknap Press, 1999); Joseph Raz, *The Morality of Freedom* (Oxford: Oxford University Press, 1986); T. M. Scanlon, *What We Owe to Each Other* (Cambridge: Belknap Press, 1998).

bracketed from a conception of a social and political concept, then this 'bare-bones' conception will resolve the contestability of the concept. Recalling that essentially contestable concepts are appraisive, however, suggests that extracting references to a substantive conception of the good will provide a definition which does not fully capture the concept's appraisive character. A value-free account of an essentially contested concept, even if it is possible, cannot provide a definition of that concept. A barebone's conception may, however, capture an aspect of the 'common core' of the concept which unifies conceptions and, if it identifies where valuations occur within conceptions, then it may help clarify the value claims which are implicit in each conception. This kind of clarificatory role may be invaluable if an essentially contestable concept is especially opaque and if, upon reflection, some of the conceptions of the good which inform particular conceptions may be unreasonable. Thus, even when a concept is found to satisfy the criteria of essential contestability, there might still be a wealth of interventions which are useful for clarifying the value commitments which underpin each conception.

Lastly, although there is no way to ascertain the best conception of an essentially contested concept, this does not commit us to complete relativism over *all* the possible conceptions of a concept. For instance, a contest between conceptions might be inevitable because there is no fact of the matter about which conception is the best conception (the relativist thesis), or there is no way to ascertain which conception is the best conception (the sceptical thesis).[139] Therefore, for a concept to be essentially contestable, on either the sceptical or relativist interpretation, there must be no generally agreed upon criteria or external standard by which to settle upon a *best* conception. The lack of a generally agreed upon standard by which the best conception is judged does not, however, imply that there are no standards. Indeed, part of the idea of sharing a 'common core' implies that conceptions must, at the very least, identify that core as being part of the concept. Moreover, that it is possible to identify *reasonable* interpretations of what is good, even if we cannot know which interpretation is the best, is well recognised in discussions on the concept of justice.[140] If we hold that through the satisfaction of (I), (II) and (III) features of our moral landscape, such as the ability to identify reasonable but not finally authoritative conceptions of the good, impose themselves on essentially contestable concepts, then it

---

[139] Swanton, "On the "Essential Contestedness" of Political Concepts," p. 814.

[140] Of course, the precise demarcation of which conceptions of the good are reasonable, and which unreasonable, may nonetheless be very difficult. See John Rawls, *Justice as Fairness: A Restatement* (Cambridge, MA: Belknap Press, 2001), pp. 191-92.

seems plausible that analysts of essentially contestable concepts should be able to exclude obviously ludicrous, incoherent or infelicitous conceptions even if they cannot completely settle the contest between reasonable conceptions.[141]

Taken together, these six points provide a more nuanced view of essentially contested concepts than Gallie's original formulation. Moreover, they suggest that the utility of labelling a concept as essentially contestable is that it motivates us to adopt a certain picture of that concept: one where there is a central, uncontested, core to the concept, around which conceptions of the good impose sets of contestable valuations. In the following sections I apply these nuances to the debate over the essential contestability of security.

2.2 The Three Concepts of 'Security'

The traditional view of 'security' is that it refers to a single essentially contested concept.[142] By the lights of a more nuanced view of essential contestability, however, it is unclear whether security really is a single contested concept, let alone whether it fulfils the criteria of essential contestability. Contrary to the traditional analysis, I believe that we should view 'security' as a polyseme which refers to three distinct, and compatible, concepts: (i) security as practice, (ii) security as a complex state of being for an entity, and (iii) security as a quality of reliability. Many of the different accounts of security offered by analysts should not, therefore, be seen as contesting the same concept of security, but rather describing distinct concepts which all use the word 'security.' Importantly, this is not a claim that all the ways in which we actually use the word 'security' are captured by these concepts, only that this is a profitable way of categorising the distinct meanings of 'security' which are most important for moral and political theory. I now provide an extended defence of this account of the term 'security.'

The distinction between these concepts is subtle, but some initial understanding of it can be gleaned from the different syntaxes of the word 'secure'. In the first instance, 'secure' can be used as a verb – as in, "X will secure Z" – to describe an action or process undertaken by X. In the second instance, 'secure' can be used as an adjective – as in, "X is secure" – to describe a quality of X's state of being. Finally,

---

[141] As Swanton memorably remarks: "While affording no criterion for determining the best wheat, contests have point (sic) because they separate the wheat from the chaff; and this separation is an endless process because political theorists are always in the habit of adding to the chaff" Swanton, "On the "Essential Contestedness" of Political Concepts," p. 815.

[142] Buzan, *People, States and Fear: The National Security Problem in International Relations*, p. 10.

'secure' can be transformed into the adverb 'securely' – as in "X securely holds Y" or "X can securely Y" – to describe a quality of the relation between X and Y.[143] These different uses of the word 'secure' could merely be syntactic transformations of a single concept of security, but when we reflect upon the different extensions which we associate with them, it appears that they refer to distinct definitional projects.

Security as a practice

First, focusing on the verbal form of 'secure' (e.g. securing someone or something) brings to mind closed circuit cameras, alarms, patrolling guards, martial law, checkpoints and fences; the paraphernalia of private security firms and state security organisations.[144] Thinking about these actors and objects motivates definitions of 'security' which capture the things people and institutions have traditionally *done* in the name of security. Indeed, as we saw in Chapter One, constructivist and post-modern scholars believe 'security' refers to "a specific field of practice" which is differentiated from other practices by a "special" kind of political action.[145] For such scholars, security practice is defined, not by states of affairs, but by reference to the practices of actors traditionally associated with 'security.' Let us call this use of the word 'security', Security$_P$.[146] The definitional project of Security$_P$ is thus constituted by answering a descriptive question about these practices:

1) What practices are associated with 'security'? Is there a property which unifies these practices?

In answering (1), scholars have variously identified the use of militaries, coercive force and other 'emergency measures' as the unifying property of security practices.[147] More post-structurally inclined scholars suggest that these practices are unified by a particular threat-defence logic which determines the way in which we engage with issues labelled as threats to security.[148] Whilst these practices clearly have some

---

[143] We might also use 'secure' as an adjective – as in "X is secure from Y" or "Y is secure for X" – to describe a quality of such a relation.

[144] Bill McSweeney makes a similar point in McSweeney, *Security, Identity and Interests: A Sociology of International Relations*, pp. 13-22.

[145] Wæver, "Securitization and Desecuritization," p. 50.

[146] Throughout this chapter whenever I refer to a concept of security I will capitalise (e.g. Security$_P$) and where I refer to an actual instance of security in the relevant sense I will use normal font (e.g. secured$_P$, security$_P$).

[147] For a representative sample, see Buzan, Wæver, and De Wilde, *Security: A New Framework for Analysis* ; Huysmans, "Security! What Do You Mean?: From Concept to Thick Signifier,"; Walt, "The Renaissance of Security Studies,".

[148] See again, Huysmans, "Security! What Do You Mean?: From Concept to Thick Signifier," pp. 232-33.

traditional relationship to security as a state of being – in the sense that they almost always worked to secure someone or something in the past, and perhaps still have the stated intention of doing so – the states of affairs which these practices actually create may be quite different from what we consciously associate with the state of being secure.[149] Indeed, this disconnect has been an important site for the critique of security as a practice – especially where security practices have induced more rather than less fear, and have resulted in the erosion of civil liberties, the persecution of minority groups or the escalation of conflict between States.[150] For scholars who focus on (1), however, what differentiates security from other concepts isn't the properties of these states of affairs but the character of the practices themselves.

Security as a state of being

Secondly, focusing on the adjectival form of 'secure' (e.g. of someone or something being secure) brings to mind qualities of an entity's state of being: i.e. that it is free from violence, has effective control of its political destiny, feels safe, and has a stable lifestyle. Thinking about these qualities seems to motivate definitions of 'security' which capture properties of an entity's state of being. Let us call this Security_B. Typically, these properties are identified by reflection upon what we consciously think it means for an entity to be secure. The definitional project is thus constituted by answering a conceptual question:

2) What does it mean for an entity to be secure? What are the properties of these states of being?

Importantly, answers to (2) are dependent on the kind of entities whose state of being we are interrogating. When asked what it means for a human individual to be secure, for instance, we might think that it means they are in a carefree state of being, where they can be assured that "the opportunities they have today (will not be) totally lost tomorrow."[151] Likewise, when thought of in the context of institutions like the state, we might think of a secure state of being as involving political and territorial control, and the maintenance of internal law and order. Security_B is thus defined, not in general, but by reflection on the secure states of being of a particular entity.

---

[149] Theorists sometimes ask whether *securing* someone always has the intention or result of making them *secure*. One might argue that if you don't intend to make someone secure then you aren't securing them, but as will become apparent, I think it is more profitable to treat this as a perfectly sincere and meaningful phrase.

[150] See, for instance, Hansen, "The Little Mermaid's Silent Security Dilemma and the Absence of Gender in the Copenhagen School,"; Neocleous, *Critique of Security*.

[151] United Nations Development Program, "Human Development Report 1994," p. 23.

Importantly, these states of being are not necessarily those which are protected by traditional security practice; the idea of security as a state of being might include goods beyond those that security forces might be able to secure, such as access to adequate shelter and food.

Security as a quality of reliability

Finally, focusing on the adverb 'securely' (i.e. of something had or done securely) brings to mind something being guaranteed, assured, or protected. Thinking about these qualities seems to motivate definitions of 'security' which capture a quality of reliability. Let us call this Security_R. Perhaps the paradigmatic use of security in this context is when we contemplate whether or not an individual enjoys a certain good securely. Consider for instance, what we might mean when we say that "Jon enjoys his liberty securely", or that "Jon's access to water is secure." The definitional project is thus constituted by answering the following question:

3) What does it mean to enjoy a good securely? What are the properties of this quality?

Strangely, very few extended discussions of (3) exist. As I showed in Chapter One, where there has been some consideration of Security_R it has been merely to note that "security is not so much a good in and of itself, but... a mode of enjoying other goods, an underwriting of other values, a guarantor of other things we care about."[152] Whilst such discussions underscore the important distinction between Security_B and Security_R, I believe we require a much more extended discussion of Security_R.

One important point to consider is that the quality of Security_R appears to attach to many things besides the secure enjoyment of goods. It makes perfect sense, for instance to talk of a rain *event* next Tuesday being more or less secure, depending on how likely it is that it might occur. Moreover, it appears that the *knowledge* that the late Cretaceous extinction was caused by an asteroid impact can be more or less secure, depending on how robust it is to changes in the evidence.[153] Importantly then, lots of things besides goods – events, knowledge, relations, facts about the world etc. – seem to be able to be secure.

---

[152] Waldron, "Safety and Security," p. 318.

[153] See, for instance, Brett Calcott, "Wimsatt and the Robustness Family: Review of Wimsatt's Re-Engineering Philosophy for Limited Beings," *Biology and Philosophy* 26, no. 2 (2011); Kent W. Staley, "Robust Evidence and Secure Evidence Claims," *Philosophy of Science* 71, no. 4 (2004).

Perhaps the simplest way of abstracting away from the objects of security$_R$ is to encapsulate them within propositions, such that we can talk of the "security of *p*" to mean the security of any properly specified object. Say I am interested in the security of an individual's freedom from violence: whether or not an individual, *i*, actually enjoys this good can be captured by the truth of the proposition "*i* is free from violence." Propositional encapsulation of the objects of security$_R$ allows for a wide range of things to be felicitous and, importantly, facilitates the adoption of a ready-made modal architecture (possible worlds talk) for the determination of whether or not the object obtains securely.[154] For instance, whilst the truth of a proposition in the actual world may capture actual enjoyment, it does not yet capture secure enjoyment. Whether or not something obtains securely seems dependent, not just on its contingent truth in the actual world, but in some range of possible worlds. At its most basic, Security$_R$ seems to be an alethic modal property (similar to probability, possibility and necessity) which tells us about the truth of a proposition, not just in the actual world, but within some subset of possible worlds. I will explore these thoughts in greater detail in Chapter Four.

One might object, of course, that Security$_R$ is not a distinct concept, rather it is merely a transformation of Security$_B$ from a property of agents into a property of objects. Some understanding of the difference between the two concepts can be gleaned, however, from the way Security$_B$ characterises a good which is constituted by a complex set of subsidiary goods (freedom from violence, from hunger etc.), whilst Security$_R$ characterises just a mode of enjoying any particular good. Consider the phrases: (a) "Jill is secure" and (b) "Jack's rights are secure". Phrase (a) deploys 'secure' to describe a complex good. This good is a complex of other goods – i.e. that Jill is free from physical violence, that she is safe from arbitrary detention and unlikely to be psychologically tortured, etc. Phrase (b), on the other hand deploys 'secure' to describe the way in which Jack holds just one good: his rights. When we express phrases such as "Jack's *y* is secure" or "*y* is secure for Jack" we are typically taken to be saying something about the mode in which Jack enjoys *y*: that there is a reasonable guarantee, or a reliable expectation, or an assurance that Jack will enjoy *y* in the future.

The objection itself also suggests a reason to suspect that these might be distinct concepts: we can describe security$_B$ for an entity by referring to the entity's secure$_R$ enjoyment of certain goods. Thus I might be secure$_B$ because I am secure$_R$ in

[154] See Chapter Four, p. 81.

my job, $secure_R$ in my access to quality food, and $secure_R$ from physical violence.[155] To avoid confusion it seems that the first use of the word 'secure' should refer to a distinct concept from the subsequent uses of the word. We will revisit the precise way in which instances of $Security_B$ can be cashed out through reference to $Security_R$ shortly, but for now it should be clear that because some instances of security as a state of being might be explained via reference to security as reliability we should regard them as informed by distinct definitional projects.

'Security' and essential contestability

What implications does this diversity of definitional projects have for the thesis that 'security' is an essentially contested concept? Recall that a contest over the concept of security can only occur between conceptions which claim to articulate the same concept of security. One might want to view each of these different ways of defining 'security', with their different definitional questions and foci, as informing conceptions which contest a single concept of security. This is the view which would lead us to regard security as a single essentially contested concept. Whilst this view is popular amongst Security Studies scholars, we have three reasons to reject it in favour of regarding these distinct definitional projects as three distinct concepts which all use the same word.

First, essential contestability requires that the various conceptions compete over the same part of our conceptual space. The questions each definitional project asks are not, however, the same, nor are they obviously in competition. Answers to (1) which identify practices which are traditionally associated with 'security' aren't incompatible with answers to (2) or (3) which identify the states of affairs which we consciously believe to be those which constitute being $secure_B$ and what it means to enjoy something $securely_R$. In this sense, it seems like a mistake to characterise them as contesting the same concept.

Second, the tripartite distinction allows us to categorise the claims made within security statements more precisely and gives us greater conceptual space to interrogate the relationship between these concepts. It is important to know, for instance, whether $security_P$ practices always promote or maintain valuable forms of $security_B$. Were we to view each of these definitional projects as contesting a single

[155] The objects which are secure for the entity need not be *goods* for the entity in the traditional sense. We can cash out how it is that a prisoner is secure by referring to the way in which his continued incarceration is secure, that his prison term is secure etc. The property of $security_R$ (as reliability) holds between the prisoner and these objects despite the lack of value they hold for him.

concept, rather than simply aiming at different things, we would impoverish our ability to theorise such questions adequately.

Third, the tripartite distinction is an important step towards clarifying the grounds for the special normative significance of security. For instance, though some forms of Security$_P$ clearly have value, they are almost all contingent and instrumental forms of value. Military hardware, political practices of prioritisation and the derogation of rights might be exceptionally important in certain circumstances, as instruments of promoting of maintaining other values, but it seems likely that none of these are valuable in and of themselves. On the other hand, it is plausible that being secure$_B$ and/or holding something securely$_R$ might be valuable, not just instrumentally, but as ends in themselves.

Given these three considerations, it seems clear that 'security' should be viewed as a poly-semantic word which can express at least three different concepts of security: security as a practice (Security$_P$), security as a state of being (Security$_B$) and security as a quality of reliability (Security$_R$).

2.3 The Structure of Being Secure

Of the three concepts which 'security' can refer to, perhaps the most historically important is Security$_B$. Use of the term 'security' – from the Roman *securitas* to national security and human security – is dominated by references to what it means to be secure. In this section I interrogate the structure of the concept of Security$_B$, by addressing two questions which underpin definitions of Security$_B$: (i) *who or what is secure$_B$?* and (ii) *what does it mean for that entity to be secure$_B$?* I discuss each of these questions in turn, slowly building a picture of the structure underlying accounts of Security$_B$.

Referent objects, entities and security$_B$

The first question which every account of Security$_B$ must address is the type of entity whose security it purports to define. Whilst we sometime speak as if there is an account of what it means to be secure$_B$ *simpliciter*, in reality, accounts of Security$_B$ are always accounts of what it means to be secure as a particular type of entity. For instance, accounts of 'national security' should be viewed as accounts of what it means to be secure$_B$ as a state, whilst accounts of 'human security' are accounts of what it means to be secure$_B$ as an individual human being. Thus, for any class of entities, x, we might imagine an account of what it means to be secure as an entity of that kind, Security$_{B(x)}$. We might therefore have different accounts of what it means to be a secure human being, Security$_{B(h)}$; a secure prisoner, Secure$_{B(p)}$; or a secure state,

Secure$_{B(i)}$, etc. without a fully-determinate account of Security$_B$ *simpliciter*. This insight suggests two important points.

First, whilst it is not often recognised, we can talk about multiple instances of Security$_{B(x)}$ applying to a single referent object.[156] Consider a female prisoner, Maja, who is kept in solitary confinement and refused access to visitors because of fears she may escape. She is regularly brutalised by guards, a practice which has instilled a sense of helplessness and passivity. Whatever we think of the permissibility of Maja's treatment at the hands of her captors I think it should be intuitively clear that she is a secure prisoner, since she is highly unlikely to escape or attack her captors. Yet it also seems clear that Maja is not a secure human being, since she herself could be violently beaten at any moment. There appears to be no 'holistic' sense of Maja's security: she is only more or less secure as a particular kind of entity (e.g. "prisoner", or "human"). We might therefore say that an individual referent object, i, may or may not be secure$_B$ as a type of entity, x_1, and may or may not be secure$_B$ as a different type of entity, x_2, with no holistic sense in which that referent is secure$_B$.

Second, a great many things may substitute for i and thus occupy the role of the referent object. We might, for instance, say that the Vietnamese economic system is a secure$_B$ economy, or that the Artic may be more or less environmentally secure (i.e. is secure$_B$ as a natural environment). For the purposes of moral and political philosophy, however, the kinds of referent objects whose security seems most relevant are things like states, individuals, institutions or communities.

Third, whilst many accounts might profess to properly define the concept of Security$_B$ *simpliciter*, there are good reasons to interpret them only as defining a particular *instance* of Security$_B$. First, interpreting human security, national security etc., as conceptions of the same concept limits our ability to theorise the relationship between the security$_B$ of different types of entities. If human security, for instance, is claimed to properly define Security$_B$ then it becomes difficult to provide a comprehensive account of the security$_B$ of states. Moreover, the security$_B$ of individuals might be partly constituted by the security$_B$ of the states to which they belong, yet the conception interpretation treats this awkwardly: insisting that we either cannot make such claims or that we are using 'security' idiomatically. Second, it

[156] For those familiar with Security Studies, the use of the term "referent object" in this way will appear slightly idiosyncratic. Whereas theorists of security consider the "referent object of security" to be a *class* of entities (such as states, societies or individuals) which may or may not be secure, philosophers use the term to identify specific objects which are referred to by a sentence: such that the man who was Prime Minister of the United Kingdom in 1942 is the referent object of the name "Winston Churchill." In this thesis, I will use the philosophical sense.

is likely that the conception interpretation of human security and national security is founded on a contest over which of these things is most important. Arguments for and against a particular account of Security$_B$ aren't normally framed as attempts to identify a natural kind, but rather as claims about the relative importance of each type of entity and the utility of a particular account to policymakers.[157] It seems advantageous, however, to be able to describe the security of many different entities, even if we don't think that their security is particularly important. By interpreting these accounts as instances of Security$_B$, rather than conceptions of Security$_B$, we may discuss the importance of each of these states of being outside a narrow conceptual contest and within a broader moral and political dialogue. Finally, it is highly doubtful that there is a single definition of Security$_B$ which covers all the various instances in which the word can be applied. (Do we really think, for instance, that there is any commonality between the Security$_B$ of information and that of a human being?). An account of Security$_B$ only seems to take on determinate content when we explicitly or implicitly conceive of it in a particular instance. We thus seem best served by thinking of Security$_B$ only for specific classes of entities, rather than seeking a concept which unifies the states of being of such a diverse set of entities. For these three reasons it seems appropriate to regard instances of Security$_{B(x)}$ as distinct though related concepts, and not conceptions of Security$_B$ *simpliciter.*

The conditions for security

The second major component of accounts of Security$_{B(x)}$ are the goods which a referent must enjoy in order to be secure as a particular kind of entity. As I noted in Chapter One, an account of human security might, for instance, suggest that an individual must enjoy access to adequate food, shelter and clean water, and enjoy freedom from violence in order to qualify as a secure$_B$ human being. Likewise, an account of national security might suggest that a state must enjoy territorial integrity, freedom from domination and provide a basic level of welfare to its citizenry in order to qualify as a secure$_B$ state. Each conception of Security$_{B(x)}$ differentiates itself through the goods which it suggests must be enjoyed by the referent in order for it be secure in the relevant way.

Importantly, however, it is not just that the referent must enjoy those goods, they must enjoy them securely$_R$. An individual is secure$_B$ not merely because he is free from violence today, but because he can be *reasonably assured* of being free from

[157] For a representative sample, see Paris, "Human Security: Paradigm Shift or Hot Air?,"; Thomas and Tow, "The Utility of Human Security: Sovereignty and Humanitarian Intervention,".

violence tomorrow, the next day and so on. It is important that there is some assurance, or guarantee, or reliable expectation of the particular goods, not merely the possession or access to those goods *per se*. To be specific, when we think of whether or not someone is $\text{secure}_{B(h)}$ we think about whether or not they enjoy the individual goods securely_R. Indeed, what seems to distinguish an account of a entity's security_B from an account of some other component of that entity's wellbeing is the quality of Security_R which attaches to the referent's enjoyment of the goods which constitute Security_B as that entity.

Recalling the discussion above, a great many things can be secure_R, and these things can be captured by well-specified propositions. Whether or not some proposition, *p*, is secure_R is dependent, not just on its contingent truth in the actual world, but in some range of possible worlds. Security_R is thus an alethic modal property which describes the truth of a proposition, not just in the actual world, but within some subset of possible worlds. In Chapter Four I will revisit the precise nature of Security_R but in the context of its role in constituting conceptions of $\text{Security}_{B(x)}$ there are three important points.

First, whilst propositions can be claims about many things, actual conceptions of Security_B are concerned with a much more restricted set of propositions: those that capture whether a referent enjoys some good. These conditions can be thought of as propositions of the form "*i* enjoys *y*", where *y* is a good and *i* is a referent which could hold or enjoy that good.[158] The truth of the proposition in the actual world tells us whether the referent will enjoy the good, whereas the security of the proposition tells us whether the referent will enjoy the good in a particular class of possible futures.

Second, it is important to note that not much hinges on the idea of enjoying a good. I'm not claiming, for instance, that a referent must actually value the good in order to enjoy it, nor that the referent must even be aware of their relation to the good.[159] The point is merely that the referent must in some sense 'have' the good in question – whether that requires legally possessing it, having access to it, or possessing the capability to use it are things which can be packed into the specification of the good itself.[160]

---

[158] In what follows, I will sometimes simplify the relation of enjoyment between a referent and a good such that it is expressed as a proposition, *p*. However, where I discuss the security of *p* one should typically assume that I am talking about propositions of the form "*i* enjoys *y*."

[159] For an account of enjoying goods which does associate enjoyment with valuation, see Shelly Kagan, "Well-Being as Enjoying the Good," *Philosophical Perspectives* 23, no. 1 (2009).

[160] For instance, an individual may, variously, enjoy "access to shelter", "legal possession of shelter" and/or "the capability to find shelter when required."

Finally, a great many different things can act as goods within the proposition, but they must be well specified. Speaking very generally, a good is just a thing or class of things which individuals could have reason to value.[161] Goods can be concrete objects (e.g. a car), abstract objects (e.g. transportation), social phenomena (e.g. diversity, cultural tradition), particular relations (e.g. friendship), events (e.g. a sunny day), and expectations (e.g. expected utility). Any of these things can substitute for *y*, with the only requirement being that they are felicitously specified. What counts as a felicitous specification will, of course, be heavily contingent on the good. In general, however, the specification of the good should include consideration of the following questions:

1. Is the good a simple object (e.g. apples) or a quantity of some object (e.g. five apples)?
2. Is the good an exact quantity (e.g. exactly five apples) or merely the satisficing of some threshold (e.g. *at least* five apples)?
3. Is the good enjoyed in a particular mode (e.g. *access* to apples)?
4. Are there temporal components to the good (e.g. an apple *a day*)?

Getting the specification of a good right is important since it will affect both the degree to which the good is secure and the degree to which its security is valuable. For instance, our *basic* liberties will, in general, be more secure than our liberty *per se*, and the security of enjoying those basic liberties will, in general, be more valuable than the security of enjoying all our liberties. For the time being, however, let us abstract away from the specification of goods, noting only that felicitously specified goods are essential to a well specified conception of $\text{security}_{B(x)}$.

The structure and essential contestability

Given our discussion in the two preceding sections, we can now provide a model framework for conceptions of being secure:

> An individual referent, *i*, may be secure_B as a type of entity, *x*, if and only if a set of relevant propositions, $(p_1, p_2, \ldots p_n)$ are secure_R to a sufficient degree.

This is just a preliminary account of the structure of conceptions of $\text{Security}_{B(x)}$ and there may well be many more specific ways of capturing the relationship between the concepts of $\text{Security}_{B(x)}$ and Security_R. For instance, one way of cashing out the sense in which Maja is a secure_B prisoner may be via reference to the security_R of the set of

[161] I remain neutral between objectivist and subjectivist accounts of value. For a comprehensive account of these different accounts, see James Griffin, *Well-Being: Its Meaning, Measurement and Moral Importance* (Oxford: Clarendon Press, 1986).

necessary conditions for Maja to be a prisoner.[162] Likewise, one might seek to restrict the relevant goods to only those which are not strictly necessary but are, in some sense, the "core" goods for the particular kind of entity.[163] Regardless of the ultimately correct formulation, this structure provides a framework for systematising claims about what it means to be secure as a particular kind of entity. Two important points spring directly from it.

To begin with, the structure suggests that conceptions of $\text{Security}_{B(x)}$ requires several different kinds of value judgement. First, the model requires that each conception of $\text{Security}_{B(x)}$ identify those propositions which are relevant to the security of the referent entity. This may result in a simple list of goods but could include an enormous array of conditions, each of which will be more or less valuable according to different conceptions of the Good. Second, if we think that a good can be enjoyed *more or less* securely_{R}, then each conception must define how securely each good must be enjoyed in order to meet the threshold of security_{B} for that particular entity. This is not simply a matter of empiricial investigation, it requires the formulator of the conception to make a value judgement on how much secure_{R} enjoyment of each good is enough for the referent to be secure_{B} in the relevant way. Third, each conception must weigh the importance of the secure enjoyment of each good to the security of the entity as a whole. This may be done in several ways. Some conceptions, for instance, might suggest that secure enjoyment of a good (e.g. freedom from violence) was *necessary* for $\text{Security}_{B(x)}$. Still others might think that secure enjoyment of some set of goods was *sufficient* for $\text{Security}_{B(x)}$. Others may simply apply weightings to the secure enjoyment of each good, such that no set of goods is either necessary or sufficient for $\text{Security}_{B(x)}$ *per se*, but that increases in the secure enjoyment of each good increase or decrease $\text{Security}_{B(x)}$ in some proportion. This weighting process requires a further set of value judgements.

Moreover, the structure suggests that the each instance of $\text{Security}_{B(x)}$ might be essentially contestable. Consider that one conception of what it means to be secure_{B} as an individual human ($\text{Security}_{B(h)}$) might require access to adequate shelter, adequate food and a high level of bodily integrity, whilst another conception might only require freedom from intentionally inflicted violence. Importantly, this seems to be a genuine contest between two conceptions of the same concept of $\text{Security}_{B(h)}$. There is no

---

[162] This formulation was suggested by David Wiens.

[163] See Alkire, *A Conceptual Framework for Human Security: Working Paper 2* ; Matt McDonald, "Securitization and the Construction of Security," *European Journal of International Relations* 14, no. 4 (2008).

ambiguity over how we are using the word 'security', nor is there any confusion over the instance of security we are describing. Each account is seeking to define exactly the same part of our conceptual space. Moreover, this contest appears to satisfy the criteria for essential contestability, since the concept of an individual's security is (I) appraisive, (II) internally complex, (III) variously describable, and (IV) open to modification. Echoing our earlier discussion of essential contestability, the concept of an individual's security also seems susceptible to the irresolvable contest over the Good. This susceptibility isn't just a phenomena of individual security, it appears to be generally true of other instances of $\text{Security}_{B(x)}$. Consider that one conception of national security might argue that the state is secure only when it is free from foreign domination, whilst another conception might insist that it should also be able to recover from natural disasters, still another conception might see the protection of its citizens from harm as necessary for the state to be secure. The suite of conceptions of state security, $\text{Security}_{B(s)}$, and human security, $\text{Security}_{B(h)}$, explored in Chapter One bear out this initial observation.[164]

We can conclude, therefore, that though there is no contest over the meaning of the word 'security' *per se*, there are numerous contests over what it means to be secure as a particular kind of entity. Furthermore, these instances of $\text{Security}_{B(x)}$ are essentially contested by virtue of the large number of value judgements required to constitute each conception. The "common core" of each instance of $\text{Security}_{B(x)}$ is the model framework we established, and in particular, its reliance on the notion of Security_{R}. This establishes a clear framework for thinking about how the notion of secure enjoyment is crucial to accounts of what it means to be secure, and how each of these accounts relate to one another.

Conclusion

This chapter has sought to provide a conceptual framework for interrogating the concept of security. When we interrogate the nuances of essential contestability, I believe we have good reason to reject the standard interpretation of security as a single essentially contested concept. In particular, I have argued that the word 'security' can be used to refer to the distinct concepts of (1) security practices, (2) security as a state of being for an entity, and (3) security as a quality of reliability.

[164] To be clear, I am not claiming that instances of Security_{B} are the only uses of the word 'security' which match the criteria of essential contestability. It is highly plausible, for example, that Security_{P} also fits the criteria of essential contestability, since there seems to at least be a contest, between "emergency measures" and "defence-threat logic", over the defining characteristic of security practice. The tension is evident even within single articles. See Wæver, "Securitization and Desecuritization,".

Furthermore, I claimed that theorists interested in whether security is especially valuable should focus on the latter two of these concepts, and their relationship. This yielded three major conclusions.

First, it seems then that a more nuanced understanding of the concept of security reveals that many accounts that purport to define 'security' do not, in fact, compete with one another. Much of what is claimed to define 'security' should be seen instead as attempts to define a particular instance of $\text{Security}_{B(x)}$. We have thus moved from a conceptual schema (represented in Figure 1) where a single concept of security is contested by a number of different conceptions to a more nuanced schema (represented in Figure 2) which holds that there are multiple concepts represented by the word 'security', that these concepts can have multiple instances and that in the case of $\text{Security}_{B(x)}$, it is these instances themselves which are essentially contestable concepts.

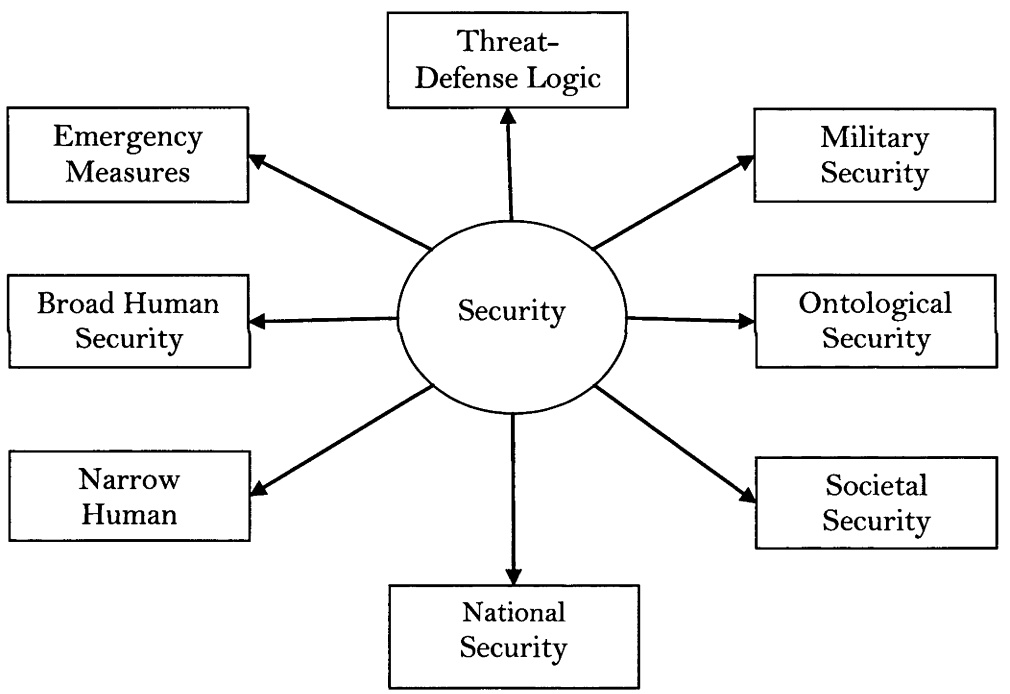


Figure 1: A graphical representation of 'security' as a single essentially-contested concept around which multiple conceptions are arrayed.

Second, this more nuanced account of essential contestability also provided a framework for thinking about conceptions of $\text{Security}_{B(x)}$ – namely that they have an uncontested core, around which different conceptions of the good substitute different conditions for an entity's security. This chapter has introduced a preliminary account of that core – the secure enjoyment of a good by the referent object. As I argue later in this thesis, the value of Security_R appears to be irreducible to other forms of value,

and if it acts as a common core to conceptions of $\text{Security}_{B(x)}$, then it may provide an explanation for the common view that there is something especially and irreducibly valuable about being secure.

Third, the essential contestability of conceptions of $\text{Security}_{B(x)}$ should motivate us to talk about the importance of the secure_R enjoyment of *particular goods* for *particular referents* rather than complex concepts such as national security and human security. When we purport to make trade-offs between "security" and other values such as liberty or equality we should be clear about exactly what set of value judgements constitute our conception of being secure. If we specify which particular goods must be securely enjoyed, then much of the opacity and confusion which seems to accompany the value of security will dissipate. Furthermore, we will no longer be forced to make a judgement on, for instance, the value of secure_R access to food and secure_R access to healthcare, considered together. Rather, we may make trade-offs between the security_R of our access to food, the security_R of our access to basic healthcare etc. Only after sustained normative and philosophical reflection upon the value of the secure enjoyment of particular goods should we attempt to construct a conception of $\text{Security}_{B(x)}$ and deploy it in our moral decision-making.

Figure 2

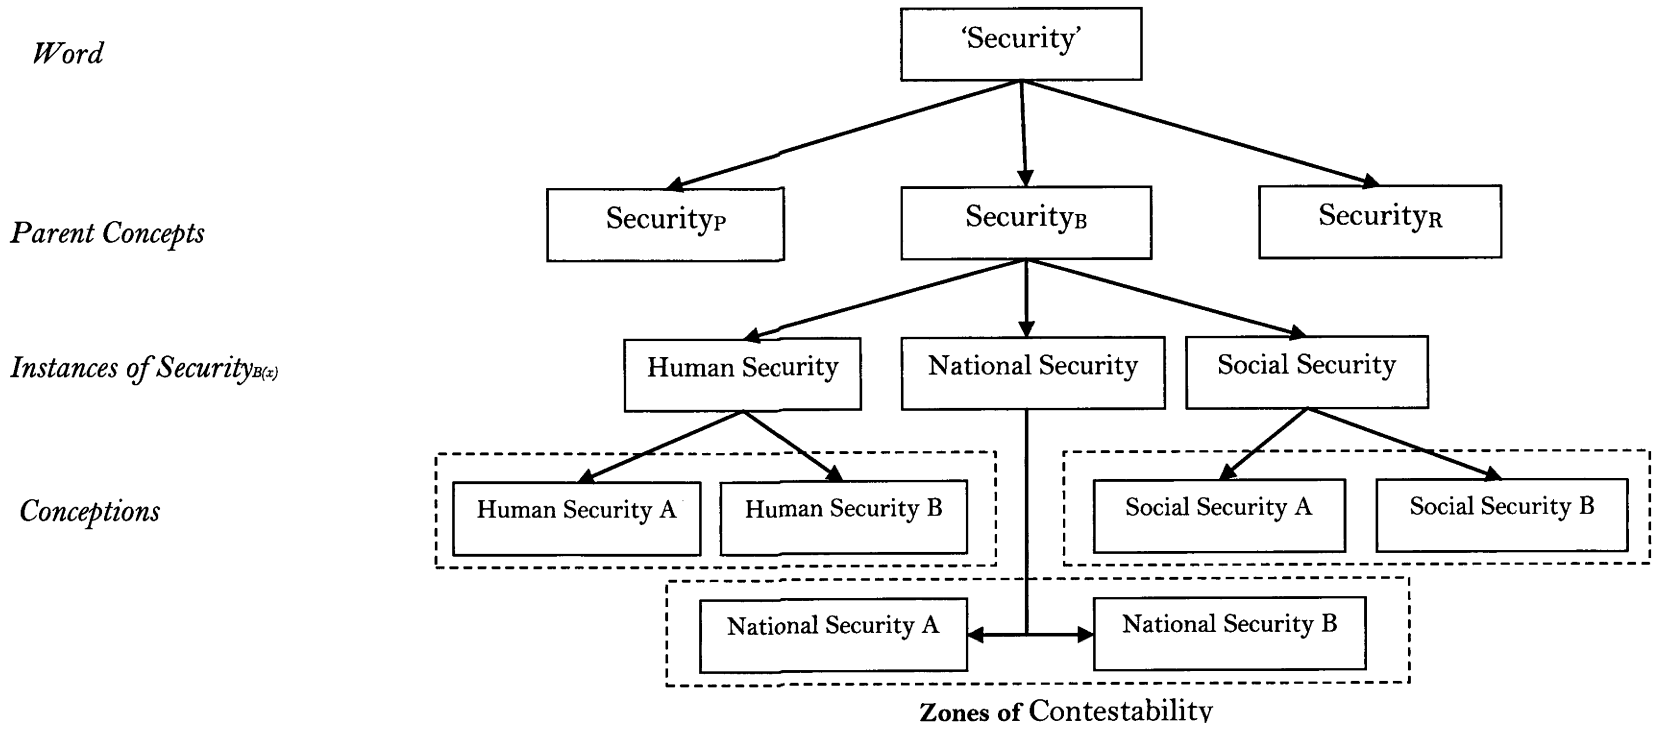


Figure 2: A graphical representation of 'security' as a polyseme which may denote security as a practice (Security$_P$), security as a state of being (Security$_B$) or security as a quality of reliability (Security$_R$). The concept of Security$_B$ can be further sub-divided into various instances (Human Security, National Security, etc.). These instances are themselves essentially-contested concepts with multiple conceptions (Human Security A, Human Security B, etc.).

Chapter Three
Subjective and Objective Security

In clarifying the meaning of the term 'security' it is important to distinguish between the subjective and objective dimensions of security. In Chapter One, I noted that the etymology of the word 'security' tends to makes its use in modern English ambiguous between being secure and feeling secure. The same kind of ambiguity seems to exist within the various concepts of security – especially within Security_B and, as we shall see, Security_R. In this chapter, I will seek to disentangle these two concepts of security from this ambiguity by getting clear on exactly what is at issue when we talk of the objective and subjective dimensions of security. In particular, I will argue that there are two versions of this distinction that can be made: one between the kinds of goods which make an entity secure_B, and one between the kinds of judgements that can be made about the security_R of a proposition. Getting to grips with these two different sets of distinctions allows us to have a better grasp of the relationship between being and feeling secure, the value of each and the role that evidence plays in judgements about security.

3.1 The Standard Account

The standard account of the subjective/objective distinction is best characterised by Arnold Wolfers, who suggests that "security, in an objective sense, measures the absence of threats to acquired values, in a subjective sense, the absence of fear that such values will be attacked."[165] The distinction seems simple enough: objective security is the actual absence of threats, whilst subjective security is the absence of fear about those threats. Obviously, these two dimensions are related, since the enjoyment of objective security may bear on whether you feel secure. For instance, being able to walk freely down a city street with only a remote possibility of being assaulted is likely to give you a far greater sense of security than being constantly at risk of harm.

[165] Wolfers, "National Security as an Ambiguous Symbol," p. 150, similar thoughts can be found in Booth, *Theory of World Security*.

Enjoying objective security should, we suppose, ensure that an entity feels subjectively secure.

Importantly though, there is nothing necessary about the connection between being and feeling secure. Consider two individuals, Bao and Yen, whose disconnect between their objective security and their sense of security are surprisingly common. [166] Bao lives in a gated-community far from any of the dangerous neighbourhoods in town and, as such, violent crime in his community is exceptionally rare. Nonetheless, Bao is incredibly anxious about the possibility of being attacked, and is constantly looking over his shoulder on his morning walk. Bao thus has all of the material goods of security but fails to possess a sense of security. Yen lives in the rougher part of town, where violent crime is frequent and random. Nonetheless, Yen has a serene outlook on her circumstances, trusting in fate and the essential goodness of those around her to keep her safe. She is perfectly at ease walking around late at night by herself. Whilst Yen is objectively insecure she clearly enjoys a strong sense of security – what we would probably call a "false sense of security."

The standard distinction between objective and subjective security thus seems to highlight something important, which is that individuals and states are sometimes prone to over-estimate the threats they face because they feel insecure. Consider Buzan's claim that:

> Efforts to achieve security can become self-defeating, even if objectively successful, if their effect is to raise awareness of threats to such a pitch that felt insecurity is greater than before the measures where undertaken.[167]

Or an entity may be prone to underestimate the threats they face because they feel secure. Consider Baldwin's discussion of Wolfers' distinction:

> Wolfers distinguished between objective and subjective dimensions of security...to allow for the possibility that states might overestimate or underestimate the actual probability of damage to acquired values. In the former case, reducing unjustified fears might be the objective of security policy; while in the latter case, a state might perceive itself as secure when it was not.[168]

The model that Wolfers, Buzan and Baldwin offer is a distinction between being and feeling secure. These two states may come apart, but in general, when things are operating as they should, being secure should directly engender a feeling of security.

---

[166] See, for instance, the large disparity between fear of crime and objective risk of crime Emily Gray, Jonathan Jackson, and Stephen Farrall, "Reassessing the Fear of Crime," *European Journal of Criminology* 5, no. 3 (2008); C. Hale, "Fear of Crime: A Review of the Literature," *International Review of Victimology* 4, no. 2 (1996).

[167] Buzan, *People, States and Fear: The National Security Problem in International Relations*, p. 24.

[168] Baldwin, "The Concept of Security," p. 14.

The problem with this model is that it conflates two underlying distinctions. In particular, it conflates (i) the distinction between the external circumstances of security and an affective sense of security, with (ii) the distinction between fact-relative, belief-relative and evidence-relative judgements about security. In what follows I discuss these two distinctions, their relationship and why thinking about each distinction separately can shed light on the various ways in which the subjective and objective dimensions of security can come apart.

3.2 External and Affective Security

One thing we might mean when we talk of the objective and subjective dimensions of security is that there are both external and affective goods which may, jointly or individually, constitute what it means to be secure. Recall that the Roman usage of *securitas* was ambiguous between an affect of calm assuredness and the external circumstances – an ordered society, the Legions etc. – which underwrote that feeling. Likewise, some advocates for human security include both material and psychological goods in their conception of what it means to be secure as a human being.[169] Others deny that security should have any psychological content at all, that what it means to be secure is simply to enjoy a set of external goods.[170] These appear to be disagreements over the kinds of goods which must be securely_{R} enjoyed in order for an entity to be secure_{B}, and are therefore disagreements over the content of $\text{Security}_{B(x)}$.

On the one hand, an entity could be secure because of certain external features of the world. On some accounts of individual security, for instance, an individual must be free from violence and have access to sufficient food, water and shelter. In still others, they should enjoy the right to vote, to express their opinions and to assemble peacefully without interference from the state. Whilst these conditions vary significantly, they are all features of the world which are, in some sense, external to the entity in question. This is not to say that all of these features are concrete aspects of the world, since whilst they do include material goods (e.g. water, food, etc.), they also include more abstract goods such as equal respect before the law. The claim is merely that being secure consists, either wholly or partially, in enjoying certain features of the world.

On the other hand, being secure could be constituted by the enjoyment of a certain attitude or feeling internal to the individual: a sense of security. We can all

---

[169] See Leaning and Arie, "Human Security: A Framework for Assessment in Conflict and Transition," .

[170] See Kanti Bajpai, "Human Security: Concept and Measurement," in *Kroc Institute Occasional Paper* (Notre Dame: Joan B. Kroc Institute for International Peace Studies, 2000).

identify with the tingling sensation of fear that accompanies a menacing figure looming out of the dark, or the anxiety which might accompany uncertainty over the origin of our next meal. To be without this feeling of anxiety – to be calm, acceptant and tranquil – is precisely what the Epicurean notion of *ataraxia* sought to capture. Importantly, when we speak of a sense of security we do not typically refer to something cognitive, like a belief, but rather something affective. Whilst there are obvious cognitive inputs into affective responses, a sense of security is precisely the thing which is felt rather than the beliefs which might inspire the feeling. The claim is that being secure consists, either wholly or partially, in enjoying a certain affect.

Each conception of $\text{Security}_{B(x)}$ will have a different view on how these two sets of goods interact to constitute security_B. On some *externalist* accounts, being secure consists only in enjoying certain features of the world. On other *affectivist* accounts, being secure ultimately consists in attaining a sense of security. Other *mixed* accounts might claim that affective and external goods are jointly necessary for security_B. Finally, some accounts might claim that there is no overarching concept of being secure, just two states of being which are typically, but not necessarily, connected. Regardless of which conception of $\text{Security}_{B(x)}$ we think turns out to be correct (for each entity, x), they clearly all conform to the same basic structure which we detailed in Chapter Two.[171] An externalist conception of an entity's security_B suggests that the relevant goods which the entity must securely_R enjoy are external goods, whilst an affectivist conception suggests that they are affective goods. Each conception of $\text{Security}_{B(x)}$ thus places different weight on external circumstances and an agent's affective response to those circumstances, and by so doing, makes a claim about the kinds of goods which an individual must enjoy securely_R in order to be secure_B.

3.3 Fact-relative, Evidence-relative and Belief-relative Security

The distinction between objective and subjective security can alternatively be viewed as three different perspectives from which judgements about security can be made. Regardless of which conception of $\text{Security}_{B(x)}$ we hold to be true, and regardless of whether it includes external or affective goods, we can evaluate the security_B of an entity from a fact-relative (i.e., "how secure *is* Bao?"), belief-relative (i.e., "how secure do *you think* Bao is?") and evidence-relative perspective (i.e., "how secure is Bao given the *available evidence*?"). Likewise, the secure enjoyment of a good can also be

[171] i.e. An individual referent, i, may be secure_B as a type of entity, x, if and only if a set of relevant propositions, $(p_1, p_2, \ldots p_n)$ are secure_R to a sufficient degree. See p. 56.

evaluated from a fact-relative, belief-relative and evidence-relative perspective. The distinction between these three perspectives is exceptionally important, as it grounds discussions in both Chapter Four and Five. Because of its importance to the thesis as a whole, in what follows, I give a little more content to each of these three perspectives.

Fact-relative Security

Talk of "objective security" sometimes seems to refer to how secure some proposition is "really" or "actually" or "given all the facts". In this sense of the phrase it is supposed to capture some property of a state of affairs, independent of what any actual agent thinks about that state of affairs. Whilst no actual agent may possess epistemic access to this property, we can think of this property as how an impartial, fully informed observer might judge an individual's secure enjoyment of a good in a given state of affairs. In the case of Bao, for instance, an impartial, fully-informed observer would judge that he is secure – even though Bao himself may think that he is insecure. For Yen, an impartial, fully informed observer would judge that she is radically insecure. For any proposition *p*, let us call this property the *fact-relative* security of *p*.

Whilst the basic notion of fact-relative security should be fairly intuitive, its details are complex. First, it may be tempting to think that fact-relative security refers only to external goods. The beliefs and emotions of individuals are, however, perfectly admissible as propositions which can be more or less secure since the state of mind of individuals within a particular state of affairs may be part of a fact-relative description of that state of affairs. For instance, whilst a sense of security is internal to individuals, it is plausible that there is a fact of the matter about whether or not a particular individual feels secure at any point in time. Likewise, whilst beliefs about secure enjoyment might be internal to agents, they can nonetheless be evaluated from an fact-relative perspective: such that the proposition "Jon believes that he enjoys freedom from violence securely" may be more or less secure depending upon some ontic fact about whether Jon retains that belief in the future. We should therefore remain open to the idea that any proposition – including propositions whose content refers to the states of mind of individuals, and to beliefs about security – can be evaluated from a fact-relative standpoint.

Second, talk of fact-relative security implicitly assumes that at least some facts about the future are unsettled. When we consider the fact-relative security of events like rain on Thursday we are implicitly assuming that the truth of the proposition "it will rain on Thursday" is not yet decided. The claim is not just that we cannot know what will be true of the future, it is that there may be *no fact of the matter* regarding

what the future looks like given the current state of the world and the laws that govern it. This view is known as the "open future" thesis and suggests that though there is only one way in which the future will actually turn out, at any particular moment in time, which possible future will be the actual future is ontically unsettled.[172]

Of course, some contend that all of the facts about the future are determined by the current state of the world and a set of deterministic laws.[173] If the world does indeed have deterministic laws then fact-relative security would be easy to characterise: in a world with only one open future, whether or not an entity enjoys a good securely is dependent on whether or not it will actually enjoy it in that future. If some version of indeterminism is true, however, then our account of fact-relative security must be able to characterise the secure enjoyment of a good given more than one open future. Given that indeterministic physical laws are plausible, and that we do not know the status of the future, we are best served by pursuing the more difficult task of contemplating security in situations with more than one possible future.

Third, there may be many reasons why the future is open for our impartial observer. It may be because there is an irreducible metaphysical property which creates indeterminacy, such as in Popper's "propensity" theory of objective chances. Or, it may be because the "best account" of the universe includes indeterministic laws.[174] Or there may even be indeterminacy over which account of the laws is the best account.[175] Whatever the correct account of why the future is open, I assume that it may be open in at least three ways. First, the future could be precisely chancy, such that some or all of the facts about the future have a precise chance of being true given the current state of the world. Second, the future could be imprecisely chancy, such that some or all of the facts about the future have a vague or imprecise chance (i.e. $0.5 \leq \mathrm{ch}(p) \leq 0.7$) of being true, with no fact of the matter regarding the precise

---

[172] This indeterminacy may be due to a metaphysically irreducible propensities, or it may be the product of indeterminacy in which laws govern the world. The point is that fact-relative security is not what any particular agent actually knows, nor what the best available science knows, but what a fully informed agent could ever know about the future. For a more complete picture, see the discussion of the "open futures thesis" in Elizabeth Barnes and Ross Cameron, "The Open Future: Bivalence, Determinism and Ontology," *Philosophical Studies* 146, no. 2 (2009).

[173] See David Lewis, "A Subjectivist's Guide to Objective Chance," in *Studies in Inductive Logic and Probability*, ed. Richard C. Jeffrey (University of California Press, 1980), p. 118; Carl Hoefer, "Determinism, Causal," in *Stanford Encyclopedia of Philosophy*, ed. Edward Zalta (2004); J. Earman, *A Primer on Determinism* (Dordrecht: D. Reidel Publishing, 1986), p. 13.

[174] See Lewis's Best System Account (BSA) of the natural laws, for instance, which allows that the natural laws could be chancy: such that at any given world at a given time there is some irreducible non-trivial chance (given the current state of the world, F, and the chancy laws, L) that a proposition, p, will be true or false. Paraphrased from David Lewis, "Humean Supervenience Debugged," *Mind* 103, no. 412 (1994).

[175] See Barnes and Cameron, "The Open Future: Bivalence, Determinism and Ontology,".

chance. Finally, the future could be possibilistic, such that though many different futures remain possible, there is no fact of the matter regarding the chance that some or all of them will be the actual future.[176]

In this and subsequent chapters I will assume that the future is precisely chancy. This is for two reasons. First, it is plausible to think that an impartial observer with access to a fully complete science and all of the facts about the actual world would be able to assign precise chances to a large number of facts about the future.[177] Second, if precisely chancy laws govern at least some aspects of the world, as seems plausible, then positing a precisely chancy future will make plain the distinctive differences between our two conceptions of secure enjoyment in Chapter Five. In particular, it will allow us to distinguish between the contributions to security of both chance and robustness. In Chapter Five we will relax this requirement, and see that our intuitions about security in contexts where the precise chances are not known may be markedly different.

Fourth, we are now in a position to identify what it means to suggest that the fact-relative security of p is what an impartial, fully informed observer might think of the security of p. In particular, evaluating fact-relative security requires abstracting away from certain epistemic limitations such that our hypothetical observer has access to the ways things are presently arranged and the various laws which actually govern the world. In a sense, our impartial observer plays the same role as LaPlace's demon: an intellect which has access to all of the information about the world as it is presently arranged as well as access to a set deterministic laws, who may then calculate all of the relevant facts about the future. The difference, of course, is that our observer is not omniscient – since assuming that the future remains unsettled by virtue of indeterministic laws – it cannot know what will in actual fact occur. Our impartial observer knows only the content of the various possible futures and the degree to which they are more or less likely to be the actual future.

Belief-relative Security

Talk of "subjective security", on the other hand, sometimes refers to the judgements that particular agents possess about the security of a proposition. In contrast to the fact-relative security of p, these judgements are relative to each agent's beliefs. Bao,

---

[176] Strictly speaking, possibilism is a special case of an imprecisely chancy future, where some or all facts are assigned the interval $0 \leq \mathrm{ch}(p) \leq 1$.

[177] Indeed, many current physical theories – including Quantum Statistical Mechanics – operate on the understanding that there are a set of precisely chancy laws which govern the universe. See Aidan Lyon, "Deterministic Probability: Neither Chance nor Credence," *Synthese* 182, no. 3 (2011).

for instance, might not know the relevant facts which indicate that he enjoys freedom from violence securely, whereas Yen may have ignored the abundant evidence regarding her insecurity, and mistakenly come to believe that the likelihood of being attacked is low. Broadly speaking, the belief-relative security of *p* for an agent can be analysed as the fact-relative security of *p* were the agent's beliefs true. There are some implications from thinking about belief-relative security in this way.

First, belief-relative judgements can be made about the security$_B$ of an entity or the secure enjoyment of a good. Bao may make a holistic judgement about his security$_B$ as a citizen or he may make a judgement only about the degree to which he securely enjoys freedom from violence. In what follows I will largely focus on belief-relative judgements about the security$_R$ of propositions, but much of what I will say ought to be compatible with judgements about the security$_B$ of a referent.

Second, belief-relative security may be self, other or non-directional. Whilst it is true that in many instances the agent forming the judgement about the secure enjoyment of a good is the entity which will enjoy that good, there is nothing necessary about this connection. Whilst Bao may judge that his enjoyment of freedom from violence is more or less secure, his friends may also have a judgement about the security of his freedom from violence. Agents may therefore have judgements about the security of their own enjoyment of a good (e.g. "I enjoy freedom from violence securely"), or the security of some other entity's enjoyment of a good (e.g. "Bao enjoys freedom from violence securely") or even the security of some non-agent-directed proposition (e.g. "it is secure that it will rain tomorrow"). In this sense, the propositions which can be more or less secure according to a particular agent may contain a variety of content, including but not limited to self and other-directed content.

Third, judgements about the security of some proposition *p* might be actual or hypothetical. An agent may actually hold a particular judgement, or we may talk hypothetically about the judgements that an agent would adopt were they offered new evidence, were fully rational or were in an otherwise different set of circumstances. A large part of the work of Chapter 5 will deal with delineating how agents should judge the security of some proposition, given the available evidence, and this may differ radically from the actual judgements that agents have about the security of a proposition.

Fourth, the belief-relative security of *p* might be more or less rational. An agent may adopt beliefs which are more or less rational, depending on whether those beliefs

conform to norms of epistemic rationality. There are a variety of accounts of the norms of rationality, some more substantive than others, but the most common norms rule out inconsistent beliefs (e.g. to be certain both that "I am free from violence" and "I am not free from violence"), require belief in logical truths, and govern the updating of beliefs given changes in sets of related beliefs.[178] Whilst no actual agent's beliefs conform to these constraints consistently, we can still make comparisons between agents such that they will depart more or less egregiously from these constraints.

Finally, the belief-relative security of some proposition *p* can be more or less justified by that agent's evidence. Bao, for instance, may hold the beliefs he does because he does not know the relevant crime statistics which would show that violent crime is very rare in his community. Importantly, it is not the case that the beliefs he actually holds are inconsistent with his beliefs about his own security, since this would make him epistemically irrational, what we mean is that he is not appraised of relevant evidence which is immediately accessible to him. The degree to which a judgement about the security of *p* is justified by the evidence suggests that there may be a third sense of security: evidence-relative security.

Evidence-relative security

Straddling the traditional distinction between objective and subjective security is the degree to which a good is securely enjoyed given the evidence available to an agent. Broadly speaking, the evidence-relative security of *p* can be analysed as the fact-relative security of *p* were those beliefs justified by the available evidence true. On this view, the evidence-relative security of *p* for an agent is distinguishable from the belief-relative security of *p* for that same agent by the set of propositions upon which it is based: the beliefs which would be justified by the evidence available to the agent rather than the agent's actual beliefs. Unfortunately, there is little agreement on what constitutes the evidence which is available to an agent. On one account the evidence available to an agent is simply those things about which the agent is certain – regardless of how the agent came to be certain of them.[179] Others contend that an agent's evidence is the set of propositions that the agent *ought* to be certain of in a

[178] Perhaps the most influential of these is the Bayesian account of constraints on beliefs, which also requires that one's degree of belief in a proposition should conform to the probability axioms, such that one cannot determinately hold a credence that *p* of 0.4 and a credence that not-*p* of 0.4 simultaneously. see William Talbot, "Bayesian Epistemology," in *The Stanford Encyclopedia of Philosophy*, ed. Edward Zalta (2008).

[179] See ibid.

particular set of circumstances.[180] Still others contend that an agent's evidence is the total set of propositions that the agent *knows* – i.e., the set of propositions about which an agent is justifiably certain and is, in actual fact, true.[181] In lieu of any clear guidance on what counts as the evidence which is available to an agent, and because what I have to say is compatible with a wide range of theories of evidence, I will adhere to a relatively pre-theoretical notion of evidence as a set of propositions about which the agent ought to be certain given their circumstances.

Importantly, the relevant sense of "ought" in this account of evidence is epistemic rather than moral. It is tempting to believe that the evidence available to an agent includes those things which, with a little effort on the part of the agent, they would come to know. If the relevant crime statistics are available from the government, it might seem as though those crime statistics should form part of the evidence available to Bao, even if Bao is never aware that they exist. This, however, confuses a moral or prudential obligation which Bao might have to act in a certain way – i.e. to act so as to be well informed about matters as weighty as his risk of being mugged – with an epistemic fact about what a reasonable person in Bao's position would know.[182] Instead of positing some moralised notion of evidence, we might wish to say that acts based on evidence-relative judgements about the security of some proposition p are more or less morally justified given the particular circumstances in which those acts are performed. If Bao fails to be know the relevant crime statistics, even though they are printed in the paper each week, we might declare him *culpably* ignorant if continue to act in ignorance.[183] If, on the other hand, he fails to know these statistics because the police refuse to release them, then his ignorance may be morally justified. If his knowing the relevant statistics is necessary for him to avert some great disaster then we might think he morally ought to steal them from the police. Obviously, what will count as morally unjustified will turn on a variety of considerations, both epistemic (e.g. the relevant evidence is inaccessible) and moral (e.g. coming to know the information would entail an unacceptable moral cost). On the view of evidence that I am suggesting, it is perfectly available to us to say that an

---

[180] See D.H. Mellor, *Probability: A Philosophical Introduction* (New York: Routledge, 2005), p. 83; Michael J. Zimmerman, *Living with Uncertainty: The Moral Significance of Ignorance* (Cambridge: Cambridge University Press, 2008), contra Richard Feldman, "Having Evidence," in *Evidentialism: Essays in Epistemology*, ed. Earl Conee and Richard Feldman (Oxford: Oxford University Press, 2004).

[181] See Timothy Williamson, *Knowledge and Its Limits* (Oxford: Oxford University Press, 2000), p. 185.

[182] see. Hilary Kornblith, "Justified Belief and Epistemically Responsible Action," *Philosophical Review* 92, no. 1 (1983); William G. Lycan, "Evidence One Does Not Possess," *Australasian Journal of Philosophy* 55, no. 2 (1977)

[183] For a discussion of culpable ignorance, see Holly Smith, "Culpable Ignorance," *The philosophical review* 92, no. 4 (1983)

agent morally ought not to have acted on the evidence he currently possessed if it was open to that agent to investigate more thoroughly. This, however, is a discussion which requires a balancing of moral, as well as epistemic, considerations.

3.4 The New Model of Objective and Subjective Security

After a more thorough investigation of the objective/subjective distinction, it appears that there are at least two ways in which it operates: first on the content of conceptions of $\text{Security}_{B(s)}$, and second on the perspective from which secure enjoyment is evaluated. We can summarise these distinctions as follows:

	Objective	Subjective
Content of Security_B	External goods	Affective goods

	Objective		Subjective
Evaluation of Security_R	Fact-relative security of p	Evidence-relative security of p	Belief-relative security of p

Recalling our earlier discussion, we are now in a position to articulate a more nuanced model of the connection between the external goods of security and a sense of security. Ideally, we might say, a sense of security should develop in all and only those who securely enjoy the relevant external goods. Consider again our two individuals: Yen and Bao. Bao enjoys all of the external circumstances of security without also possessing a sense of security. The more fine-grained account of the subjective and objective dimensions of security provides several possible explanations for this disparity. First, Bao may have judged that he does not securely enjoy the external circumstances of security. This may be because he lacks the relevant evidence or, even if he possesses the relevant evidence, he may have an irrational reading of the evidence. Second, he may have successfully judged that he securely enjoys the relevant goods, but nonetheless failed to experience a sense of security, because he is disposed towards feeling anxious. On the other hand, Yen enjoys few of the material goods of security whilst possessing a strong sense of security. Once again she may either have insufficient evidence to form the correct belief, be unresponsive to the evidence she does possess or otherwise predisposed towards a feeling of calm assurance.

These two cases suggest that we may have a model of the relationship between fact-relative security, evidence–relative, belief-relative and affective security (see Figure 3).

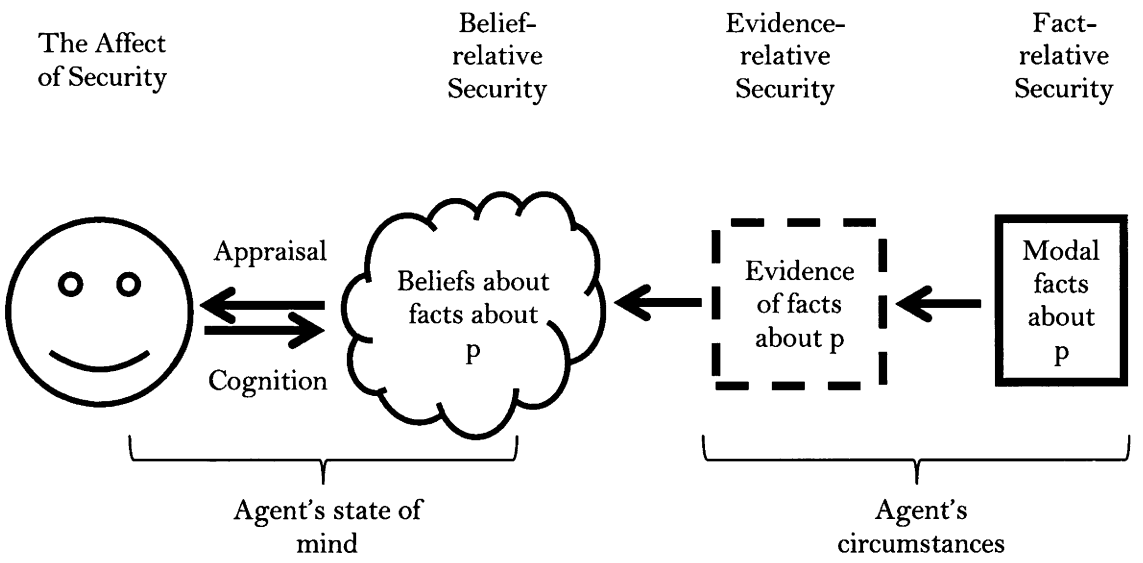


Figure 3: A model of the connection between the fact-relative, evidence-relative and belief-relative perspectives on the security of a proposition *p*, and the affect of security.

This is a vast departure from the simple distinction between objective and subjective security. On this model, there are three crucial steps in the connection between the security of external goods and a sense of security. The first is the connection between the fact-relative security of an entity's enjoyment of external goods and the security of the entity's enjoyment of those goods relative to that agent's evidence. The second is the connection between the evidence-relative secure enjoyment of the goods and the agent's beliefs. Finally, there is the connection between an entity's belief-relative judgements and the affect of security. I will discuss each in turn.

Fact-relative and evidence-relative security

The first departure from the simple distinction between objective and subjective security is that this model explicitly identifies the connection between fact-relative and evidence-relative security. In particular, it posits that the evidence-relative security of *p* can come apart radically from the fact-relative security of *p* by virtue of hard or soft evidential constraints. Hard evidential constraints are those which make knowledge of the fact-relative security of *p* permanently outside the reach of actual agents. We might think, for instance, that there is a set of laws which govern complex social relations, but that the precise structure of these laws will never be knowable by actual agents – no matter how sophisticated their social scientific methods become.[184] Soft

[184] This may be for a variety of reasons: perhaps social scientists cannot observe the system properly without interfering in it, or perhaps the agent's themselves can change the laws once they identify

evidential constraints are those things which currently restrict the evidence available to agents but which it is possible to surmount. Soft constraints might include a lack of currently available scientific evidence (which may be surmounted by performing relevant experiments, collecting the relevant data etc.), institutional censorship, or the agent's epistemic laziness. We might therefore distinguish between judgements which are formed based on the best currently available evidence and those which are based on an agent's incomplete understanding of the evidence.

The upshot of this is that the fact-relative security of most propositions is inaccessible to actual agents. Even if we assume knowledge of the best current science, the fact-relative security of only a fraction of propositions – that "Venus will transit the Sun on 5 June 2012" or "an average male Australian will live to 85" – will be knowable.[185] Most goods that we care about – freedom from violence, the right to vote, adequate incomes – arise from complex social processes which are governed, if at all, by unknown or unknowable laws. We might have some evidence relevant to the fact-relative security of these propositions, but in general we will not have epistemic access to the fact-relative security of most goods. In this sense, as a criterion of actual moral decision-making, we require some account of the way security operates according to the evidence available to agents.

Evidence-relative and belief-relative security

The second departure from the simple distinction between objective and subjective security is that the model allows for the evidence-relative and belief-relative security of p to come apart for the same agent. In particular, the belief-relative security of p will be subject to a host of cognitive biases. These may both privilege particular pieces of accessible evidence and distort their interpretation. For instance, individuals tend to systematically overestimate the risk of flying versus the risk of driving because of the cognitive availability of instances of plane crashes, the illusion of control over motor vehicle accidents and the affective punch of thinking about the large loss of life in an aircraft disaster. Likewise, a phenomena known as confirmation biasing suggests that agents tend to search for and interpret new information such that it confirms their

---

them. It is plausible that the "laws" of social practice are socially constructed and mutable across time and changes in our knowledge of them.

[185] Indeed, the fact-relative security of the proposition "*This* Australian male will live to 85" is likely to be unknowable for almost every Australian male.

preconceptions.[186] There are a host of other biases – including the conjunction fallacy, the gamblers fallacy, backfire effects, and probability neglect – which undermine our interpretation of the evidence. In general, many of these biases stem from the tendency to use heuristics (such as the *availability heuristic*, which suggests that individual estimations of the probability of some event will be affected by the ease with which actual or hypothetical instances of that event come to mind).[187] In general, we cannot assume that the available evidence will be accepted by agents or that perfectly rational inferences from the evidence will be made by agents.

The relationship between evidence-relative and belief-relative security therefore suggests that working to increase the evidence-relative security of some good may not necessarily increase the belief-relative security of the good for many agents. Agents may not be properly apprised of the evidence that is available. Likewise, even if we can apprise agents of evidence, they are likely to ignore some of it, privilege other parts of it, and use it to make inferences which are not strictly rational.

Belief-relative security and the affect of security

The third, and perhaps most significant, departure from the simple distinction between objective and subjective security is that this model allows us to interrogate the way in which well-informed belief-relative judgements may or may not result in a sense of security. As we noted with Bao and Yen, forming the judgement that you securely enjoy a certain good may not necessarily give you a sense of security. This may be for a variety of reasons. Whilst you may judge your enjoyment of freedom from violence to be secure, there may be other goods – such as good health, a steady income etc. – which you have judged to be insecure. Or, one might be unable to shake off an anxious disposition towards a particular circumstance: such as the feeling of panic that an arachnophobic might suffer, even when they know the spider in question is harmless. These kind of disconnects between belief-relative security and a sense of security are importantly different to the evidential and cognitive deficits which separate fact-relative, evidence-relative and belief-relative security.

To understand such disconnects, the psychological and sociological literature on the connection between cognitive and affective responses to insecurity is useful.[188] There are a range of theories which explore this link, but let me introduce just one –

---

[186] Eliezer Yudkowsky, "Cognitive Biases Potentially Affecting Judgement of Global Risks," in *Global Catastrophic Risks*, ed. Nick Bostrom and Milan M. Cirkovic (Oxford: Oxford University Press, 2008), p. 98.

[187] ibid., p. 92.

[188] See, for instance, Ute Gabriel and Werner Greve, "The Psychology of Fear of Crime. Conceptual and Methodological Perspectives," *British Journal of Criminology* 43, no. 3 (2003).

appraisal theory – which seeks to explain differences in the affective responses that individuals have to the same set of circumstances (as well differences in the responses of the same individual over time).[189] Consider the response of similarly well-off individuals to losing their job: some might be anxious, others might be angry, still others might be relieved or despondent. Broadly speaking, appraisal theory suggests that these divergent affective responses arise from different appraisals of the circumstance in relation to each individual's broader configuration of needs, resources and goals.[190] These appraisals may be unconscious or conscious, are multi-dimensional, and may involve a large set of ancillary belief-relative judgements regarding, *inter alia*, the degree of control that an agent has over a situation, the resources at their disposal should the threat materialise and the centrality of the good to their wellbeing.

Importantly, the appraisal theory approach to the production of a sense of security accords with the standard view of ontological security. As we explored in Chapter One, for an entity to enjoy ontological security is for that entity to have a secure and stable sense of self. On the view of Giddens and Mitzen, ontological security is thus constituted by a stable, "routinized" connection between an individual's identity and the social and material order which envelops them.[191] The idea is that some threats may be routine for an individual whereas other individuals may find their understanding of the ordinary world and their place in it disrupted. On the other hand, if some good is not enjoyed securely, then it may not produce a sense of anxiety if the individual values or otherwise accepts that uncertainty as part of their overall conception of self. In this sense, the affect of security may be produced when the judgement that p is secure accords with, or does not disrupt, the individuals own sense of identity and purpose.

None of this is to endorse either appraisal theory or ontological security as the correct account of the connection between belief-relative judgements and a sense of security. More work needs to be done, in psychology, cognitive science and philosophy of mind, to determine the appropriate connection between the cognitive and affective responses to secure enjoyment. These two accounts do, however, draw

---

[189] See I.J. Roseman and C.A. Smith, "Appraisal Theory: Overview, Assumptions, Varieties, Controversies," in *Appraisal Process in Emotion: Theory, Methods, Research*, ed. K.R. Scherer, A. Schorr, and T. Johnstone (Oxford: Oxford University Press, 2001) and the critique by Michael McEachrane, "Emotion, Meaning, and Appraisal Theory," *Theory & Psychology* 19, no. 1 (2009).

[190] C. A. Smith and L. D. Kirby, "Putting Appraisal in Context: Toward a Relational Model of Appraisal and Emotion," *Cognition & Emotion* 23, no. 7 (2009): pp. 1353-4.

[191] Giddens, *Modernity and Self-Identity: Self and Society in the Late Modern Age*, p. 32.

attention to the fact that if an individual may fail to enjoy the affect of security even when they judge that they securely enjoy all of the relevant goods, then intervening to increase the evidence-relative security of those goods and educating individuals about those interventions may not be enough to engender a sense of security. If the affect of security is to be pursued, it may have to be done, not through increasing the evidence-relative or belief-relative security of a good, but by increasing an individual's sense of security more directly. For instance, Bao may buy a gun to keep on his person in order to give him a sense of security, even if he knows that the statistics suggest that this may actually increase his risk of harm. Or governments may place thermal scanners at airports to increase the community's feeling of security from infectious diseases, even when it is clear that it would have a negligible impact on the ferocity of a pandemic.[192] Of course, policy based on increasing enjoyment of the affect of security, even if this decreases evidence-relative security, should be treated with care. The point, however, is that we should not assume that the affect of security is entirely rational.

Conclusion

It should now be clear why the standard account of the distinction between objective and subjective security is inadequate. First, it provides an inadequate model of the way in which an agent's sense of security is dependent on its fact-relative (objective) security. With only a basic distinction between being and feeling secure scholars will under-appreciate the various epistemic, cognitive and affective barriers which operate to create false senses of security and insecurity. By distinguishing between the fact-relative, evidence-relative and belief-relative secure enjoyment of goods I showed how hard and soft epistemic barriers, cognitive traits and dispositional factors all combine to frustrate or enable an agent's feelings of security.

Secondly, the standard account obscures two sets of distinctions which are important for understanding what is at stake in discussions about the value of security. On the standard account it is difficult to distinguish between the value of being fact-relatively secure, the value of judging oneself to be secure and the value of feeling secure. Once we get clear about the distinction between evidence-relative security and the affect of security we may be able to identify more clearly what is valuable about each and what role they play in our moral decision-making. For instance, it is

[192] There is good evidence to suggest that this is actually the case. See D. Bitar, A. Goubar, and J.C. Desenclos, "International Travels and Fever Screening During Epidemics: A Literature Review on the Effectiveness and Potential Use of Non-Contact Infrared Thermometers," *Eurosurveillance* 14, no. 6 (2009).

plausible that feeling secure is a component of the goodness of badness of a *state of affairs*, whereas the secure enjoyment of a proposition, given a particular action, is a component of the rightness or wrongness of that *act*. This distinction between security as a virtue of states of affairs and security as a virtue of actions will be at the heart of the discussion in Chapter Six.

Finally, the new model allows us to interrogate the notion of secure enjoyment more fully. One of the most fundamental claims which this thesis makes is that the special importance of secure enjoyment is grounded in the role of evidence-relative judgements in moral decision-making. In Chapter Four I will show how there appears to be nothing particularly distinctive about secure enjoyment when considered from a fact-relative perspective and that it boils down simply to the chance of some proposition being true in the future. Whilst the concept of fact-relative security thus has content, it is content that is well known and articulated elsewhere, and does not explain nor justify the intuition that there is something especially valuable about the secure enjoyment of a good. In Chapter Five, on the other hand, I will show that evidence-relative security does have distinctive content and that it captures an important and over-looked consideration in moral decision-making.

Chapter Four
Security and Possible Futures

Enjoying a good right now provides small comfort if it could be lost at any moment. We do not want to rely upon good fortune for our good health, decent incomes, stable communities and opportunities for political expression; rather we want some mechanism in place to ensure that we get these things regardless of many of the vagaries of future circumstance. This chapter investigates the relationship between secure enjoyment and related phenomena such as chance and robustness. I distinguish two rival conceptions of secure enjoyment – promotional and protective security. Roughly speaking, on the promotional conception, the degree to which a good is enjoyed securely is solely the chance of that good obtaining in the future. On the protective conception, the degree to which a good is enjoyed securely is dependent upon the chance of that good obtaining regardless of the future dispositions of powerful actors. Whilst the protective account offers a tempting vision of security, I argue that it is too substantive to serve the ecumenical goals that we established at the start of this thesis. I give an error theory for the intuition which supports the protective conception and go on to suggest that the promotional conception appears to account for this intuition in most cases of secure enjoyment.

4.1 Preliminaries

Before we begin, some brief preliminaries. First, as I discussed in Chapter Two, being secure is dependent on enjoying certain goods securely. I posited a general form of accounts of $\text{Security}_{B(x)}$ which could be applied to a variety of entities: such that a state could be secure_B if certain conditions were secure_R, and an individual could be secure_B if other conditions were secure_R, etc. We can therefore talk of the security_R of a great many things – events, knowledge and goods – and the most agreeable way of capturing this diverse set of referents is to encapsulate them within propositions whose truth values correspond to particular states of affairs. Security_R, I suggested, can therefore be analysed as an alethic modal property which tells us something about

the truth of a proposition, not just in the actual world, but in some set of nearby possible worlds. Whilst the generality of this account is important, I also noted that the most important kind of security$_R$ for moral and political philosophy was the secure$_R$ enjoyment of a good by an entity. Throughout this chapter I will sometimes adopt the propositional form by discussing the security$_R$ of p (where p is any felicitous proposition), other times I will speak plainly of an entity's secure$_R$ enjoyment of a good. I will also, for the most part, drop the subscript "R" from security$_R$: but all instances of "security," "secure" and "securely" should be read as referring to the property of security$_R$ unless otherwise explicitly stated.

Second, I argued in Chapter Three that 'security' can be evaluated from a fact-relative and an evidence-relative standpoint. Recall that the fact-relative perspective is akin to what an impartial, fully informed observer might think of an individual's secure enjoyment of a good whilst the evidence-relative perspective captures what particular agents ought to believe about an individual's secure enjoyment of a good given the evidence which is available to them. In this chapter, I will primarily address the secure enjoyment of goods from a fact-relative perspective. I do so in order to get clear on whether security has distinctive ontic (as opposed to epistemic) features which might ground the value of secure enjoyment in the structure of particular states of affairs and their futures. I will therefore abstract away from epistemic limitations and will assume full knowledge of the way things are presently arranged, the possible ways in which they could be arranged in the future, and the likelihood that each of those future arrangements will characterise the actual world. Obviously this information may be inaccessible for a great many goods – particularly goods which arise from complex social processes. The point, however, is that we should establish what security consists in as a property of states of affairs before investigating how changes in our information about those states of affairs influence our beliefs about the security of goods.

Possible worlds and open futures

To explore fact-relative security I will deploy a variant of the familiar possible worlds architecture which accommodates the idea of an open future. The open future thesis suggests that at any point in time in a particular world, most of the facts about the future of that world are unsettled. For example, given a coin is tossed at the present time, t, the open futures thesis posits that whether or not the coin lands on heads at $t+1$ will be unsettled at t. This is not just the claim that these facts are unsettled for most actual agents, but that they are also unsettled for agents like our impartial

observer who have access to all of the facts about the world as it stands now and also the laws of that world.[193] There are several points to note about how the possible worlds framework interacts with the open future thesis.

First, one can think of possible worlds as devices, designed to capture the different ways in which the world could have been or may be in the future.[194] Each possible world describes an entire universe: including all the facts about the physical laws, the position of all the atoms, who is living and dead, the thoughts and feelings of individuals, etc. Importantly, each world has a fully determinate history, such that all of the facts about the world – from the deepest past to the furthest future – are settled. Any difference in the facts (including the history and future) is a difference in worlds: such that the world, w_1, where I finish this thesis by the end of 2012 should be thought of as an entirely different world to the one, w_2, where I finish this thesis by the middle of 2013.[195]

Second, because each possible world has a fully determinate history, to identify a particular time, t, in a particular world, w, we must talk in terms of a world-time pair, $<w, t>$. These identify particular moments in a particular world and let us distinguish between the world as it is in 1968 (i.e. $<w, t_{1968}>$) and the world as it is in 2011 (i.e. $<w, t_{2011}>$). According to the open futures thesis, for each world-time pair $<w, t>$ there is a set of worlds, historically indistinguishable from w up until t, which represent the ways in which w may turn out to be in the future (see Figure 4).[196] Call this set of possible worlds the *possible futures* of the world-time pair $<w, t>$.

Third, one of these possible futures, the *actual future*, is the way things actually will be – this is the world where, amongst other details, the speed of light is 2.99×10^8 m/s, the Normans won the Battle of Hastings and you are reading this chapter. According to the open future thesis, the facts relating to the future in the actual world

---

[193] See Barnes and Cameron, "The Open Future: Bivalence, Determinism and Ontology," p. 295

[194] This is the ersatz view of possible worlds. For some, of course, possible worlds aren't just a device, but are real existing things. See David K. Lewis, *On the Plurality of Worlds* (Oxford: Blackwell, 1986; reprint, 2001), p. 2.

[195] This is the standard approach to possible worlds, although there are some who take a 'branching' view of worlds. The details don't concern us here, but some discussion of the distinction can be found in ibid., p. 206.

[196] It might be natural to suppose that possible futures should be historically *equivalent* to w at t, in the sense of that all facts about w up to t are also true of each possible future up to t, but this is a mistake. As Fred Feldman explains, however, which future does in fact turn out to be actual will have implications for some facts about the past. Suppose I quit smoking yesterday and it is now open to me to either never smoke again or to take up smoking again: if I continue to abstain then the cigarette I smoked yesterday will be my last, if I return to old habits then it will not be. In this sense, minor facts about the past may be different depending on which futures turns out to be the actual future. See Fred Feldman, *Doing the Best We Can* (Dordrecht: D. Reidel Publishing, 1986), pp. 18-19.

are unsettled, such that although there is only one actual world, at any point in time in that world there may be several possible futures which could be the actual future.

Fourth, possible futures are more or less *distant* from one another depending on their similarity. Thus, the world where I eat a bagel for lunch on 31 January 2012 is closer to the actual world than the world where the Allies lost the Second World War, which is closer again than the world where the speed of light is 4.99×10^8 m/s. The notion of similarity is important for understanding the scale of changes to a world, such that a possible future which is drastically different to the world as it currently is, will be more distant (less similar) than another possible future which does not change nearly so much. We will revisit similarity when we discuss the notion of robustness.

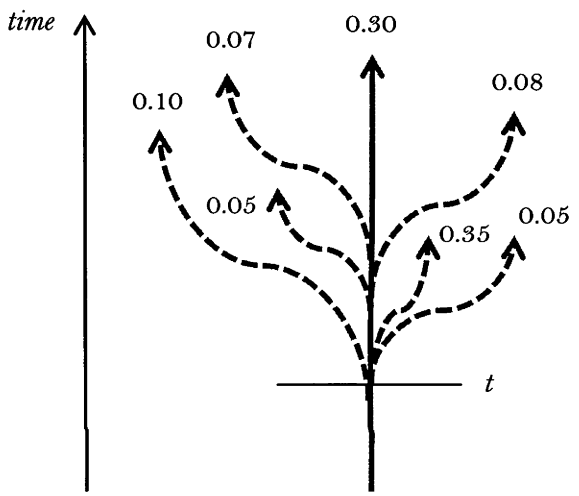


Figure 4: Representation of the open future framework. Each line represents a possible world, with the full line representing world, *w*. At time *t* in world *w* there are several worlds which are possible futures of <*w*, *t*>, represented by dashed lines, which are historically indistinguishable to *w* up to *t* but which diverge from then on (although perhaps not immediately). Numbers represent the chance that each possible future will be the actual future at *t*.

Finally, each possible future of a particular world-time pair may have a greater or lesser chance of being the actual future. The possible world foundations of chance are complex and contested, but the technicalities need not concern us here.[197] The basic idea is that on most accounts of the openness of the future, not all possible futures are equally likely to be the one which actually occurs, and that whatever the more fundamental explanation, an fact-relative perspective will have access to the chance that each possible future will be the actual future. We can model this distribution of chances by positing that there is a chance function which, given all the facts about the

[197] For some proposals, see John C. Bigelow, "Possible Worlds Foundations for Probability," *Journal of Philosophical Logic* 5, no. 3 (1976); Lyon, "Deterministic Probability: Neither Chance nor Credence,".

world at $<w, t>$, assigns to each possible future a chance of being the actual future. These chances will be relative to $<w, t>$, such that at the moment a fair coin is tossed, t_{TOSS}, the possible futures w_H and w_T have an equal chance of being the actual future, but after the coin comes up heads at t_{TOSS+1}, the chance of w_T being the actual world reduces to zero.[198]

Intuitions

To ground our analysis of fact-relative security, I will briefly outline five commonly held intuitions about the nature of secure enjoyment. These intuitions provide a yardstick against which conceptions of secure enjoyment may be judged. Recall, however, that I also set out at the beginning of this thesis to explain security's role in our moral decision-making in a way which is ecumenical towards an array of reasonable conceptions of the Good. The conception of secure enjoyment which best balances these three goals – intuitive fit, utility in moral theorising and moral ecumenicism – will be the best conception of secure enjoyment for my purposes.

First, security comes in degrees. If you buy a single ticket to a 100-ticket lottery then winning the lottery seems less secure than if you buy half of the available tickets.[199] To be sure, sometimes usage of 'security' seems to imply "an absolute condition – something is either secure or insecure – and (this usage) does not lend itself to the idea of a graded spectrum like that which fills the space between hot and cold."[200] Whilst I agree that we might sometimes say that winning the lottery is determinately secure or insecure, I think what we mean is that it is secure to some very high or very low degree. In this sense, we should view the security of p as a scalar property of p.

Second, secure enjoyment is determined not only by what actually occurs, but also by what could occur. If you buy a single ticket to a large lottery, then even if, in actual fact, you win the lottery, it does not seem that winning the lottery was very secure. Whilst in the actual world you may have won, there were very many nearby worlds where you did not win, and this latter fact seems important to how secure winning the lottery is for you. In this sense, the security of p captures a modal fact about p: to be secure with respect to p is to have p, not necessarily in the actual world, but in some class of possible worlds.

---

[198] For a discussion of the temporality of chance, see Lewis, "Humean Supervenience Debugged," p. 475.

[199] I'd also suggest that it seems *50 times* less secure, but for the moment, let us remain agnostic between whether security is an ordinal or a cardinal property.

[200] Buzan, *People, States and Fear: The National Security Problem in International Relations*, p. 24.

Third, security is relative to particular circumstances and times. If you buy a single ticket to a large lottery at time t, then winning the prize at $t+2$ does not seem secure, but if your ticket is declared the winner at $t+1$, winning the prize seems very secure. In this sense, the security of p at time t is determined by facts about p in the possible futures of $<w, t>$. Importantly security does not seem to relate to the counterfactual facts *at* time t. There may be many possible worlds at time $t+1$ where you have a different ticket to the one you actually do – but these count for nothing against the security of your winning the prize at time $t+2$. Whilst you may be *lucky* to have the ticket you do at $t+1$, this does not diminish the security of the prize because it is almost assured you will get it at $t+2$ given you have the ticket that you do.[201]

Fourth, a high degree of secure enjoyment seems to suggest a high chance of enjoying the good. If you buy a single ticket to a large lottery, then winning the prize seems not very secure. It seems intuitively right, however, that you marginally increase the security of winning the prize by buying an additional ticket. The only relevant fact which appears to have changed in this case is that you have increased the chance that the actual future will be one where you win. Thus, *prima facie*, increasing the chance that p seems to increase the security of p.

Fifth, a high degree of secure enjoyment suggests that enjoyment of the good is robust to changes in the circumstances. Suppose that the sweet couple across the street are highly unlikely to ever try to harm you or your family. Whether or not your freedom from violence is secure seems to not only depend upon the chance of being attacked by *those* neighbours, but also by whether or not future sets of neighbours, perhaps totally unreasonable neighbours, would likewise be unlikely to harm you. Security, in this sense, is tied to the idea of a reasonable guarantee: that even if some set of harms is, in the actual circumstances, very unlikely, we are not secure unless we are protected from those harms in a wide range of possible futures. The larger the class of possible futures where p is true, the more of a guarantee one has regarding p, and hence the greater the degree of security over p. Thus, *prima facie*, increasing the range of circumstances across which p is true, increases the security of p.

The first, second and third intuitions are platitudes – they seem to state completely uncontroversial facts about security. The fourth and fifth intuitions give a little more

[201] For a selection of material on the modality of luck, see E. Coffman, "Thinking About Luck," *Synthese* 158, no. 3 (2007); Jennifer Lackey, "What Luck Is Not," *Australasian Journal of Philosophy* 86, no. 2 (2008); Neil Levy, "What, and Where, Luck Is: A Response to Jennifer Lackey," *Australasian Journal of Philosophy* 87, no. 3 (2009).

content to the property of security, but importantly, it is difficult to discern whether chance and robustness are independently important features of security. I think it is obvious that the chance of enjoying a good is constitutive of its security – the higher the chance that you have of being free from violence the more secure that good seems – but I am less confident that robustness is independently important. In what follows I investigate two conceptions of security that explore this tension.

4.2 Promotional and Protective Security

We sometimes take liberty and security to be in opposition to one another,[202] so it may seem odd to refer to debates over the content of freedom in a discussion of secure enjoyment. Nonetheless, understanding the neo-republican account of freedom, and its key points of difference with the standard view of freedom as non-interference, does illuminate two different conceptions of secure enjoyment.

The standard view of liberty, what Isaiah Berlin termed "negative liberty", is that an individual is free in so far as that individual can choose certain courses of action without being interfered with by other agents. So, if A has a choice between x and y, she is free to the extent that B *will not* interfere with A regardless of whether she chooses x or y.[203] This is the theory of freedom as non-interference. Philip Pettit, however, rejects this view of freedom, instead suggesting that A is free to the extent that B does not possess the power to arbitrarily interfere with A's choice, regardless of whether she chooses x or y and regardless of B's disposition towards A. This is the theory of freedom as non-domination.[204]

Consider the situation of a slave, the archetype of the unfree person, whose master is, against all the odds, quite well disposed towards the slave.[205] Despite having absolute power over the slave, the master never interferes in any of the slaves choices and is highly unlikely to ever do so. The theory of freedom as non-interference, counter-intuitively says Pettit, takes a slave in such a position to be free.

---

[202] See, for instance, the critical appraisal of this view in Waldron, "Security and Liberty: The Image of Balance,".

[203] Note that A's freedom is unrelated to how attractive x or y are to her. If B would prevent her from choosing x she is less free, even if she is highly unlikely to choose x. Pace Hobbes, "Leviathan," chap. 21, sect. 2.

[204] Note that only *arbitrary* interferences make an agent less free. On Pettit's account, when Ulysses demands his sailors bind him to the mast, and so interfere with his actions, they do so non-arbitrarily. Philip Pettit, "Freedom and Probability: A Comment on Goodin and Jackson," *Philosophy & Public Affairs* 36, no. 2 (2008): p. 210.

[205] Philip Pettit, *Republicanism: A Theory of Freedom and Government* (Oxford: Oxford University Press, 1999), pp. 22-23.

The theory of freedom as non-domination, however, takes the slave to lack freedom precisely because the master is able to interfere with impunity.

Pettit briefly suggests that the distinction between freedom as non-domination and freedom as non-interference can be understood, to a certain extent, as different ways that non-interference can be secured.

> To try to secure non-interference in the protection sense is to try and reduce interference in those possible worlds where other people take against you or you are not so cunning or whatever; and to do this regardless of the probability of those worlds. To try to secure non-interference in the promotional sense is to try and reduce interference in various possible worlds, but in a way that takes account of how probable it is that those worlds are ways the actual world may be.[206]

Securing non-interference in the promotional sense, Pettit suggests, does not adequately capture what it means to be free, since it can be accomplished by ingratiating oneself with the powerful or duping them into believing you are choosing according to their preferences, and thus "may involve leaving control of the good in the hands of another."[207] This kind of obsequiousness may lower the likelihood of interference, but, in an important sense, it does not seem to protect the agent's freedom. What is required is to secure non-interference in the protective sense: to reduce the probability of interference regardless of the different dispositions which powerful agents may have towards you.[208] If we think properly securing non-interference requires protecting it from changes in the dispositions of powerful actors, Pettit argues, then a concern for non-interference should motivate a concern for non-domination.

Whilst Pettit is concerned only with the security of non-interference, the two conceptions of security he describes may be generalizable to the enjoyment of other goods. Roughly speaking, on the promotional conception, the degree to which a good is enjoyed securely is solely the chance of that good obtaining. On the protective conception, the degree to which a good is enjoyed securely is dependent upon the chance of that good obtaining regardless of the dispositions of powerful actors. In what follows I explore both the promotional and protective conceptions of security and suggest that in terms of intuitive fit, normative ecumenicism and utility in moral

---

[206] ibid., p. 74 fn. 7.

[207] ibid., p. 74.

[208] In later work, Pettit has clarified that this only renders certain domains non-probabilistically relevant to the analysis of freedom: the domain in which "the endangered agent's choice-dispositions vary" and the domain in which "there is variation in the interference-dispositions of endangering agents." It is, in this sense, not equivalent to securing non-interference across *all* possibilities. See Pettit, "Freedom and Probability: A Comment on Goodin and Jackson," p. 218.

theorising, the promotional conception appears to be the most promising account of security.

Lilliput and Blefescu

To give a more concrete analysis of the two conceptions of security, consider a circumstance with only six possible futures. Lilliput and Blefescu are sovereign states that share a river as their primary water source. Lilliput (in the delta) is downstream of Blefescu (in the mountains), who has built a number of large dams upstream of the border. There is no institutional agreement between the two states regarding the shared use of water from the river system, and Blefescu therefore has a large degree of arbitrary power over Lilliput's access to water depending on the prevailing weather pattern. In a La Nina year, Blefescu must release water from its dams (lest there are floods in Blefescu's major cities), and this always provides sufficient water to Lilliput. In an El Nino year, however, Blefescu may either share its water with Lilliput or retain it all for itself. Finally, in the unlikely event that catastrophic climate change creates a Super-El Nino, Lilliput and Blefescu will go thirsty together.

Lilliput's Access to Water at $t+1$		Blefescu's Disposition	
		Friendly	Unfriendly
Weather	La Nina	0.40	0.10
	El Nino	0.36	0.09
	Super El Nino	0.04	0.01

Table 1: Each cell is a distinct possible world, and contains the chance of that world being the actual world. Shaded cells are those were Lilliput retains access to sufficient water at $t+1$. At time t in the world of evaluation, w, it is an El Nino year but Blefescu is friendly, and this remains true in the actual world at t_{+1}.

In the actual world Lilliput might continue to enjoy access to sufficient water at $t+1$, yet that access to water is clearly insecure to a certain degree, since there are nearby worlds where Lilliput goes thirsty. The question, of course, is which conception provides the right analysis of Lilliput's security at time t.

4.3 Promotional Security: Chance

Let us first consider the promotional conception of security. On this account, the degree to which a particular proposition is secure is determined solely by the chance of that proposition being true in the actual future.

> **Promotional**: A proposition, p, is secure to degree x, at a particular world-time pair, $<w, t>$ iff the chance that the actual future of $<w, t>$ is a p-world is x.

This is a natural and intuitive way of capturing the degree to which a good is secure at a world-time pair. In our example, the security of Lilliput's water supply is just the chance that the actual future will be one where Lilliput enjoys access to water (i.e. $p = 0.86$).

That *Promotional* reduces the security of p down to the chance that-p has several benefits. First, very small changes in the chance that-p will be reflected directly in the degree of security of p. If we think that our account of security should be able to straightforwardly distinguish between fine adjustments in the secure enjoyment of a good, as suggested by the intuition that security is a scalar property, then *Promotional* provides just such an account. Second, each possible future contributes to the security of p only in proportion to its likelihood. *Promotional* thus captures the plausible intuition that, all other things being equal, very unlikely futures do not diminish security to the same degree as very likely futures. Third, *Promotional* appears to be compatible with a wide range of conceptions of the Good. Finally, the connection between *Promotional* and moral decision-making is relatively straightforward. If, in a particular state of affairs, the value of p is settled and the chance that p is settled, then the security of p contributes to the value of that state of affairs through a function of the chance that p and the value of p.[209] *Promotional* therefore directly integrates the notion of secure enjoyment into at least one standard framework for moral decision-making: maximising expected moral value.

Chance and robustness

We might be tempted to say, however, that *Promotional* misses an important point about security: which is that it matters not just how likely it is that we will continue to enjoy a particular good in the future, but that a securely enjoyed good should continue to be enjoyed even if some unfortunate change does occur. The fifth intuition about secure enjoyment suggested that a high likelihood of enjoying a good was important in both the actual circumstances and a range of reasonable circumstances. We don't just care that Blefescu and the weather are likely to cooperate so that Lilliput will enjoy access to water, it also seems important that even if they do not that Lilliput will continue to enjoy access to water. Making sure that there is a reasonable guarantee that a good is resilient to, even unlikely, harms seems like an important part of

[209] The most obvious function would be $\mathrm{ch}(p) \times v(p)$. Minor modifications may be made to allow for a risk-averse orientation, such that the value of a 0.5 chance of p will be less than half the value of p. This will be discussed in greater detail in Chapter Six, p. 132.

security that may not be captured by *Promotional*. Let us call this a worry about robustness.

Like security, the precise characterisation of robustness is contested. There is a wealth of work on robustness in the context of scientific explanation.[210] It is also the foundation of solutions to the problem of closure and knowledge in epistemology.[211] Sometimes robustness is taken to mean the sensitivity of a fact to present or past subjunctives – such that Aasim's love for his partner Jeff is said to be robust if Aasim and Jeff would still be lovers even if we made a number of small changes to the past or present (i.e. in the way they had met, or their present taste in movies etc.). Security, however, relates explicitly to future subjunctives – given the way things are and have been, will Aasim and Jeff continue to be lovers across a wide range of ways in which the future could turn out for them. For the purposes of this thesis I therefore understand robustness in the following way.

> **Future Robustness**: A proposition, p, is more robust than another proposition, q, across the possible futures of a world-time pair, $<w, t>$ iff p is true in a wider class of possible futures of $<w, t>$ than q (regardless of the likelihood that any of those futures will be the actual world).

Propositions which remain true despite more radical changes – i.e. in a wider class of worlds that includes more dissimilar futures – will be more robust. The second law of thermodynamics might be true across all of the futures compatible with the physical laws as they are in our world, and it would therefore be more robust than the putative law of supply and demand, which seems only to be true across futures compatible with certain assumptions about resource scarcity and human nature.[212] Likewise, that the US President is decided by an electoral college might be true across all the worlds compatible with the current US constitution, but the two-party system might only be compatible with a much narrower range of future political and social arrangements. The general idea is that the size of the class of futures were p remains true acts as a way of measuring the degree to which p is robust.

Treating robustness as independently important to security is initially appealing because it directly captures the intuition that the more dissimilar the futures where p continues to be true, the more secure the proposition in question. In a sense,

---

[210] Calcott, "Wimsatt and the Robustness Family: Review of Wimsatt's Re-Engineering Philosophy for Limited Beings,"; Robert Hudson, "The Methodological Strategy of Robustness In the context of Experimental Wimp Research," *Foundations of Physics* 39, no. 2 (2009); Staley, "Robust Evidence and Secure Evidence Claims,".

[211] See Duncan Pritchard, *Epistemic Luck* (Oxford: Clarendon Press, 2005).

[212] For related comments, see Christian List, "Republican Freedom and the Rule of Law," *Politics, Philosophy & Economics* 5, no. 2 (2006).

however, the robustness of a good is already captured by *Promotional*, since the conditional chance of enjoying the good even in very remote circumstances diminishes the overall chance. Consider a flood which is likely to occur every 1-in-100 years. If, because you have neglected to adopt appropriate protections (e.g. levies, stilt foundations), damage to your home given that the flood occurs is almost certain, then the unconditional chance of damage to your home that year is almost 1/100. However, if you have prudently adopted protections and the conditional chance of damage given the flood is only 1/100, then the unconditional chance of damage that year is 1/10000. The difference between the unconditional chance of enjoying the good in the two situations is almost 2 orders of magnitude. The unconditional chance of enjoying a good thus appears to capture part of the intuitive appeal of robustness across circumstances.

Even if unconditional chance captures much of the intuitive appeal of robustness, one might still wonder whether, *all other things being equal*, the robustness of p across possible futures contributes to the security of p. To test this, think of two different situations similar to Lilliput and Blefescu's (see Table 2).

A

Lilliput's Access to Water at t+1		Blefescu's Disposition	
		Friendly	Unfriendly
Weather	La Nina	0.9	0.02
	El Nino	0.02	0.02
	Super El Nino	0.02	0.02

B

Lilliput's Access to Water at t+1		Blefescu's Disposition	
		Friendly	Unfriendly
Weather	La Nina	0.18	0.18
	El Nino	0.18	0.18
	Super El Nino	0.18	0.1

Table 2: Each cell is a distinct possible future, and contains the chance of that future being the actual future. Shaded cells are those where Lilliput retains access to sufficient water at t+1. At time t in the world of evaluation, w, it is a La Nina year and Blefescu is friendly, and this remains true in the actual world at t+1.

In circumstance A, the only instance in which Lilliput will enjoy access to water in a year's time is the one which obtains now: i.e. if Blefescu remains friendly and a La Nina pattern continues to prevail. Nonetheless, though there is only one possible future where p will be true, such a world is very likely to be the actual world. In circumstance B, owing to Lilliput's prudent investment in water recycling, a super El

Nino is the only instance where Blefescu could deny Lilliput access to sufficient water. There are therefore five possible futures where Lilliput would continue to enjoy access to sufficient water, and one where it does not, and each of these worlds has a roughly equal chance of being the actual world. Although Lilliput's chance of enjoying access to water is the same in A as it is in B, Lilliput's access to water is more robust in B than in A, for the simple reason that there are a larger class of possible futures (a larger range of circumstances) where it continues to enjoy access to water.

My strong intuition is that Lilliput's access to water is as secure in A as it is in B – suggesting that the robustness of p is not, even all else being equal, constitutive of the fact-relative security of p. Perhaps though, you do have the intuition that there is a difference between A and B, and that future robustness is constitutive of security. This intuition may be explained by some common psychological biases. Consider that individuals tend to prefer lotteries with many outcomes which produce a payoff to lotteries with equivalent odds but a much smaller number of outcomes which produce a payoff.[213] Thus, a game where a player wins \$100 if a red card is drawn from a full deck tends to be preferred to a game where a player wins \$100 if a coin comes up heads.[214] Winning the prize in the card game is more robust than the prize in the coin toss, despite the fact that they have equivalent probabilities. Likewise, the intuition that robustness matters might be related to a tendency to prefer multi-stage lotteries to equivalent single-stage lotteries. Consider that a nuclear reactor with a 1/1000 chance of front-line failure coupled with a 1/1000 chance of a second-line failure, may intuitively appear to be preferable to a reactor with a 1×10^{-6} chance of front-line failure with no available second-line.[215] The multi-stage reactor is more robust since there are three out of four possible futures where the reactor remains safe (i.e. w_{AB}, $w_{\sim AB}$, $w_{A\sim B}$) whereas the single stage reactor remains safe in only one out of two possible futures. If you are compelled, as you should be, to think such preferences are irrational, then it seems safe to disregard the *ceteris paribus* contribution of robustness to the security of p.

---

[213] See Paul S Carlin, "Violations of the Reduction and Independence Axioms in Allais-Type and Common-Ratio Effect Experiments," *Journal of Economic Behavior & Organization* 19, no. 2 (1992).

[214] See Lee A. Kirkpatrick and Seymour Epstein, "Cognitive-Experiential Self-Theory and Subjective Probability: Further Evidence for Two Conceptual Systems," *Journal of Personality and Social Psychology* 63, no. 4 (1992); Paul D. Windschitl and Gary L. Wells, "Measuring Psychological Uncertainty: Verbal Versus Numeric Methods," *Journal of Experimental Psychology: Applied* 2, no. 4 (1996).

[215] The effect is known as 'negative recency', where the occurrence of the initial instance of bad luck is thought to diminish the chance of a further instance of bad luck. See Matthew Rabin, "Inference by Believers in the Law of Small Numbers," *The Quarterly Journal of Economics* 117, no. 3 (2002).

I think we therefore have good reason to reject robustness as an independently important component of the fact-relative security of a proposition. There is, however, a slightly modified version of the robustness intuition which may be more plausible. This is that the robustness of the *chance* of enjoying a good across *particular kinds* of circumstances may have a special impact on the security of a good, regardless of the probability of those circumstances obtaining. Indeed, we might be tempted to say that *Promotional* misses an important point about security, which is that it matters *why* a harm is unlikely.

4.4 Protective Security: Robustness across Dispositions

For political philosophers, there seems to be something especially worrisome about other agents possessing the capacity to arbitrarily end your enjoyment of a good, even if the chance that these powerful agents will exercise this capacity is similar to the chance that a "natural" impediment to enjoyment may arise. The same concern seems to motivate Pettit's theory of freedom as non-domination: such that individuals are free just in so far as powerful agents are not *able* to arbitrarily interfere with their choices, regardless of the chance that such agents will interfere. In this sense, the degree to which non-interference is secure is determined, not by how little interference actually occurs, nor by how likely that interference is, but by the robust chance of non-interference across the different dispositions that powerful agents might adopt. This special concern for the dispositions of powerful agents can be generalised such that, whilst we remain sensitive to the chance that each possible future will be the actual future, we treat variations in the dispositions of powerful agents as non-probabilistically relevant.[216] This provides the following analysis of protective security.

> **Protective**$_1$: A proposition, p, is secure to degree x, at a particular world-time pair, $<w, t>$, iff, regardless of which disposition, d, powerful agents adopt, the chance of the actual future of $<w, t>$ being a p-world is <u>at least</u> x.

On this account, variations in the dispositions of powerful actors are non-probabilistically privileged, such that to determine the degree of security of p according to *Protective*$_1$ find the chance that p conditional on d_1, and the chance that p conditional on d_2, etc. (repeating for all of the possible dispositions that powerful agents might adopt). The degree to which p is secure is equal to the smallest of these

[216] Pettit, "Freedom and Probability: A Comment on Goodin and Jackson,".

conditional chances. Thus, the security of Lilliput's access to water is given by the *minimum* chance that it will retain access conditional on any of the possible dispositions that Blefescu might adopt.[217] In our case, the chance that Lilliput will retain access to water were Blefescu to develop an unfriendly disposition is ½, and this is therefore the degree of security of Lilliput's access to water. Obviously, this is a far cry from the result of *Promotional*, but it captures what seems like an important point – that Lilliput's access to water is at the mercy of changes in Blefescu's disposition and that this, in an important sense, renders Lilliput's independence less secure.

The arbitrary privileging of dispositions

Whilst *Protective*$_1$ is initially compelling, one might object that it arbitrarily privileges one particular set of changes over other, equally relevant changes. It is not clear why we should think that the dispositions of powerful agents matter regardless of their likelihood, whereas other, potentially as harmful, changes matter only in proportion to their likelihood. Consider a farmer whose crop of opium poppies is at risk both from seizure by a local rebel group and heavy rains in September. Though both variations are equally devastating to her crop, only one, on *Protective*$_1$, is privileged in determining the security of her crop. In a sense, *Protective*$_1$ seems to smuggle in a substantive normative commitment: which is that agent-caused harms deserve greater attention (are more relevant) than natural harms.

To be sure, the privileging of dispositions seems justifiable when we are discussing the security of non-interference. Recall the slave with the beneficient master: the fact that the slave is not free seems to stem from the fact that the slave's freedom from interference is not robust across changes in the disposition of his master. If the security of the slave's non-interference was governed solely by the chance that he would not be interfered with (as it is on the promotional conception), then he could increase his security by ingratiating himself with his master, or cunningly concealing his true courses of action. This seems like the wrong analysis of the security of a slave's freedom, since the concept of having a prerogative is undermined if our enjoyment of that prerogative is under the arbitrary control of someone else (no matter how likely we can make it that they will refrain from interfering).[218] Freedom through ever more subordination is a contradiction.

---

[217] See ibid.

[218] Pettit is of course careful to distinguish between the arbitrary interference of the slave owner as opposed to the non-arbitrary interference of the state – since the state's interference is designed to maintain non-dominative relations between persons. Ibid., p. 210.

Pettit's justification for the privileging of dispositions in the slave case does not, however, seem to uncontroversially apply to other goods. One might think that secure enjoyment of freedom from violence – a paradigmatically important component of security$_B$ – would be a case, like non-interference, where the possible dispositions of other agents should be non-probabilistically relevant. Consider a situation where your neighbour is unlikely to beat you simply because she has always had a kindly disposition towards you, as opposed to a situation where there is a legal apparatus which sanctions those who attack their neighbour. Whilst both situations may result in the same chance of violence, the former relies upon on the high likelihood of your neighbour remaining beneficently disposed, whereas the latter ensures a high chance regardless of whether your neighbour's disposition sours. Importantly, however, the enjoyment of bodily integrity does not seem to be incompatible, in the way that being free does, with the arbitrary control of others. It seems perfectly reasonable to say that you may increase your security from violence by ingratiating yourself with your neighbour – by making yourself loved rather than preparing your defences – even as we recognise that you may have traded one value (freedom) for another (security of body).

This case against *Protective*$_i$ may become stronger when we consider the security of your access to goods without an interpersonal element, such as water, food or shelter. The enjoyment of these goods is certainly compatible with the arbitrary control of another: each litre of water is the same whether our access to it is controlled by a well-keeper or by the weather. Whilst we might once again think it is an undesirable trade-off, it therefore does not seem like a contradiction for you to make your access to water more secure by ingratiating oneself with the well-keeper. Whilst I do not think our intuitions are particularly strong in these cases, leaving these goods under the arbitrary control of others does not obviously invite contradiction in the same sense as can occur with freedom.

We can perhaps explain the difference in our intuitions by recalling that Pettit's appeal to protective security is supposed to motivate us to adopt his vision of the content of liberty: it is not an argument for thinking about security *per se.* Most of the pull of the slave case rests on the fact that being under the direct control of another seems incompatible with being free. In this sense, the attractiveness of *Protective*$_i$, over the more natural *Promotional,* seems dependent on whether some strategies which increase the promotional security of non-interference seem to diminish the agent's freedom. I think that, once we try and generalise beyond goods whose enjoyment is diminished by ceding control to others, we are less motivated to

reject the promotional conception. That the capriciousness of actors is treated as more important than the capriciousness of the environment seems like an arbitrary distinction once we extend protective security to apply to goods other than freedom.

A defender of *Protective*$_I$ might reply that security, being an important political concept, should be primarily concerned with the actions, intentions and dispositions of agents – and not the natural impediments to the enjoyment of goods.[219] Indeed, agent-causation is sometimes held to be the defining feature of security:

> Those threats which are grounded in the purposive behaviour of other actors – social threats – are distinguishable in terms of the policy required to address them from those which arise from the chance occurrences of the natural order and require different measures...Natural risks are, by nature and definition, unintended....(and) in so far as they arise from natural causes...do not fall under the category of security.[220]

The privileging of the dispositions of powerful agents, so goes the reply, is justified on the grounds that the goal of a political philosophy is to temper the social (rather than natural) threats we face as citizens bound to live together. This may be a coherent position if we were adjudicating between which goods should be included in a conception of Security$_B$, but I think it requires far too many prior normative commitments to operate as a constraint on the concept of Security$_R$. There are several political philosophies, even philosophies of freedom,[221] which pay heed to the need to secure citizens from 'natural' threats to their wellbeing. Moreover, the distinction between 'natural' and 'social' threats is difficult to identify for a wide swathe of deficiencies – including famines, [222] ill-health and harms arising from natural disasters.[223] In this sense, I think the privileging of the dispositions of powerful agents, as an artefact of Pettit's discussion of freedom, cannot be justified in an account of protective security which aspires to be applied generally.

Security and salient properties

Whilst robustness across the dispositions of powerful agents may appear to be arbitrary, there may yet be a way of rehabilitating the protective conception so that it

---

[219] See, for instance, Pettit, *Republicanism: A Theory of Freedom and Government*, pp. 52-53.

[220] McSweeney, *Security, Identity and Interests: A Sociology of International Relations*, p. 89.

[221] See Sen and Nussbaum's conception of freedom as a set of basic capabilities to be and do. M.C. Nussbaum, *Creating Capabilities: The Human Development Approach* (Belknap Press of Harvard University Press, 2011); Amartya Sen, *Development as Freedom* (Oxford: Oxford University Press, 1999).

[222] Amartya Sen, *Poverty and Famines: An Essay on Entitlement and Deprivation* (Oxford: Oxford University Press, 1983).

[223] See, for instance, the review of the literature on the social determinants of health and critique of the idea that health is a natural good in Sridhar Venkatapuram, *Health Justice: An Argument from the Capabilities Approach* (London: Polity, 2011).

is generalizable to other goods. Those who are attracted strongly to the robustness intuition might suggest that robustness across different classes of changes may be important for the security of different goods, such that the precise changes which must be non-probabilistically privileged would depend on the proposition being considered. Consider that the security of a friendship might be characterised by a high chance of remaining friends regardless of how much wealth either of you possessed,[224] or that the security of a scientific claim might be characterised by a high chance of it being true regardless of which more fundamental theory is true.[225] Likewise, having a chance of being able to speak freely, regardless of the content of your speech, seems important to the secure enjoyment of the right to political expression. On this more convoluted account, the security of a proposition is dependent on its chance of being true across a salient class of properties.

> **Protective$_2$**: A proposition, p, is secure to degree x at a particular world-time pair, $<w, t>$, iff, regardless of which property, v, obtains, the chance of the actual future of $<w, t>$ being a p-world is at least x. Where v is a member of a *salient* class of properties, V (ie. where $v \in V$).

To determine the degree of security of p according to *Protective*$_2$: find the chance that p conditional on v_1, and the chance that p conditional on v_2, etc. repeating for all members of the salient class of properties, V. The degree to which p is secure with respect to V is equal to the smallest of these conditional chances. Thus, if we take Blefescu's dispositions to be the salient class of properties then Lilliput's access to water is secure to degree .5 (see Table 1).[226] On the other hand, if we take the weather to be the salient class of properties then Lilliput's access to water is totally insecure. *Protective*$_2$ thus gives different answers to the question – "how secure is Lilliput's access to water?" – depending on what class of properties is salient.

Whilst this may provide a more ecumenical vision of security than *Protective*$_1$, it introduces its own problems. In particular, for most propositions it is not clear that there is a class of properties which should be non-probabilistically privileged. Certainly there appears to be little reason to non-probabilistically privilege either changes in the weather or fluctuations in Blefescu's dispositions when contemplating the security of Lilliput's access to water. To be sure, it appears that the security of

---

[224] See Pettit, "Making Good: The Challenge of Robustly Demanding Values,".

[225] On this last point, see Staley, "Robust Evidence and Secure Evidence Claims,".

[226] It is important to note that the security of a good in respect to all of the possible changes will always be either one (just in case every world is a p-world) or zero (just in case there is some world which is not a p-world). Thus, in the case as we have described it, Lilliput's access to water would be maximally insecure in respect to changes in Blefescu's dispositions and the weather.

some goods, such as friendship, might be constituted by robustness across some salient class of properties (e.g. your wealth). Like Pettit's argument for the importance of robustness over the dispositions of powerful agents, however, cases which appear to support *Protective*$_2$ should be read as arguments for the content of particular goods, not for a protective conception of security generally. If a good has modally demanding features then those features should constitute what it means to enjoy that good and not what it means to enjoy it securely.

Conclusion

Given these two conceptions of security, is one clearly superior? I take it that the promotional conception appears to be the most natural and intuitive way of discussing the degree to which a good is secure from a fact-relative perspective. It seems right that a circumstance should contribute to or diminish the security of a good in direct proportion to how likely it is to occur, and this accords with standard accounts of moral decision-making. Moreover, the promotional conception is compatible with a wider range of moral theories. The protective account sought to non-probabilistically privilege certain changes, and so treat some (potentially unlikely) changes as more salient to the security of a good than other (potentially more likely) changes. Whilst this appears to be justifiable for certain goods, there is no principled way of deciding which variation to privilege generally. Once we abstract away from securing freedom, we lose much of the rationale for privileging variations in the dispositions of actors, and don't gain any reasons to privilege other variations. Thus, for all the appeal of the protective conception when applied to the security of freedom, it appears to have serious difficulties as a more general account of fact-relative security.

If, from an fact-relative perspective, the security of a proposition at a world-time pair is merely the chance that the possible futures of the world-time pair will be p-worlds, then it is tempting to think that we ought to abandon talk of security altogether. Whilst having a high chance of enjoying some goods (perhaps those that normally constitute conceptions of Security$_B$) is important, it is already a well characterised component of value. In a certain sense, talk of security – as opposed to the chance that p – does not so far appear to add anything to our vocabulary of moral decision-making.

Notice, however, that our dismissal of robustness and embrace of *Promotional* required that we have precise knowledge of all of the chance information, including exactly how conditional chances (e.g. the chance of saving ones house conditional on the occurrence of a flood) contributed to the overall chance that p. Whilst we may be

able to access the probability of some future events – such as the chance that any given Carbon-14 isotope will decay over 6000 years (roughly ½) or the chance that a typical Australian male born in 1980 might live to 90 years of age – we simply do not have the resources to know the precise chance of events which are the result of complex social processes (what, for instance, is the precise chance that China will invade Taiwan before 2025?). Epistemic access to chances almost always relies on extremely well-evidenced background theories and the law of large numbers – things which we cannot avail ourselves of in the context of most goods. There may therefore be space for security as an evidence-relative category – one which illuminates an important component of our moral decision-making in instances where we are uncertain as to the precise chance that we will enjoy a good in the future.

Chapter Five
Evidence and Security

It is rare that we have direct evidence about the chance that an event will occur – rarer still when it comes to the goods whose security we typically care about. As J.M. Keynes long ago noted, for many important propositions – "the price of copper and the rate of interest twenty years hence, or the obsolescence of a new invention, or the (future) position of private wealth-owners in the social system" – our evidence may not warrant a definitive and precise picture of our future prospects.[227] Nonetheless, whilst the fact-relative security of most important political and social goods may be well beyond our epistemic reach, we must nonetheless make decisions which impact upon the secure enjoyment of those goods and we ought do so with due regard to the available evidence. We therefore require an account of the structure, and justification for, evidence-relative judgements about the security of a proposition.

Recall that in Chapter Three I argued that the distinction between objective and subjective security obscures a number of distinct, though inter-related concepts. The most important distinction, I surmised, was that between a fact-relative and an evidence-relative perspective from which the security of a proposition could be assessed. In Chapter Four I suggested that the fact-relative security of a proposition at a particular world-time pair was determined by the content of the possible futures of that world-time pair and the chance that they would be the actual future. In what follows I provide an account of evidence-relative security which seeks to connect it to fact-relative security whilst incorporating some importance considerations regarding the limits of the available evidence.

I begin by suggesting that an evidence-relative perspective is captured by considering the beliefs which a perfectly rational agent would be justified in holding given the available evidence. The evidence-relative security of a proposition can

[227] J. M. Keynes, "The General Theory of Employment," *The Quarterly Journal of Economics* 51, no. 2 (1937): p. 214.

therefore be analysed, broadly speaking, as the fact-relative security of p were those relevant beliefs justified by the evidence true. I then go on to identify those beliefs which are relevant, suggesting that though beliefs about the chance that-p seem intuitively relevant, we would be best served by considering the credence in p which is justified by the evidence. Next I investigate the problem of *imprecise* evidence: i.e., those circumstances where the evidence does not uniquely determine the credence in p which a rational agent ought to hold, but is compatible with a variety of ways of distributing credence amongst the relevant possibilities. Ultimately, I argue that the evidence-relative security of a proposition p is determined by the minimum degree to which p is fact-relatively secure were any of the distributions of credence compatible with the evidence the objective distribution of chance.

5.1 Evidence, Beliefs and Justification

Before I begin, let me clarify some of the concepts which I will use throughout this chapter.

First, in order to fix the evidence-relative perspective, I will discuss the beliefs which rational agents would be justified in holding given the available evidence. Talk of such idealized agents is designed to bracket away the kinds of cognitive biases and heuristics which we identified with belief-relative security in Chapter Three. These agents are hypothetical – there need not actually be an agent who holds that set of beliefs nor has access to the particular set of evidence which we are discussing. In this sense, such agents are merely constructs designed to identify the set of beliefs which the available evidence justifies.

Second, throughout this chapter I will make a distinction between *chance* and *credence*.[228] The chance that a proposition is true is the objective probability that it will be true, which is determined by the physical laws and facts at the actual world-time pair. It should be thought of as a fact-relative property which is independent of both the beliefs of agents and the available evidence, such that a proposition may have a high chance of being true despite being highly improbable on the presently available evidence. An agent's credence in a proposition, on the other hand, is the agent's *degree of belief* in that proposition, which is best understood as a subjective measure of the probability that the proposition will be true. An agent's distribution of credence amongst various propositions is a measure of an agent's confidence in the truth of

[228] An introduction to this relatively standard distinction between chance and credence can be found in Mellor, *Probability: A Philosophical Introduction*, pp. 8-13.

those proposition, such that a credence of 1 in p entails absolute certainty that p is true, whilst a credence of 0 in p entails absolute certainty that p is false, and a credence of 0.5 entails being precisely divided between p and $\sim p$. Whilst actual agents might adopt a particular credence in a proposition for a variety of reasons, in this chapter I will be primarily concerned with the distributions of credence which a rational agent would be justified in adopting given the available evidence.[229] In this sense, I will seek to identify the set of beliefs, and distributions of credence amongst those beliefs, which are justified by the evidence.

Third, whilst there are a number of ways of defining the *available evidence*, I will treat it as merely the set of propositions, E, about which a rational agent ought to be certain of given their circumstances.[230] Exactly what propositions an agent ought to be certain of given a set of circumstances is an open question,[231] but it is important to note that the relevant sense of ought being used here is epistemic, rather than moral. Moreover, whilst we sometimes act as if particular pieces of evidence justify a set of beliefs, what we are interested in is what an agent's *total* evidence – i.e., the set of all propositions about which an agent ought be certain – justifies.[232]

Fourth, and perhaps most importantly, let me explain what I mean by the notion that the evidence *justifies* certain beliefs. My account of evidence-relative security is based upon a broadly Bayesian epistemology. The hallmarks of Bayesianism are the view that belief comes in degrees measured by the distribution of credence amongst propositions, that those distributions conform to the probability axioms, and that agents update their distributions of credence by conditionalizing upon the total evidence.[233] Whilst Bayesianism is compatible with a range of theories of evidential

---

[229] We should be careful to distinguish this notion of justified credence from "evidential" or "epistemic" probability. These kinds of probability seem to play a critical role in determining the degree to which a piece of evidence confirms a particular hypothesis, but for the purposes of this thesis I place them to one side. For a discussion, see ibid., pp 80-90; Williamson, *Knowledge and Its Limits*, p. 209.

[230] Recall the discussion in Chapter Three, p. 180.

[231] Feldman, "Having Evidence," p. 226; Mark Schroeder, "Having Reasons," *Philosophical Studies* 139(2008).

[232] Feldman, "Having Evidence," p. 226.

[233] The conditionalization procedure typically employs Bayes Rule, such that an agent who has an initial credence in p, $c_i(p)$, and comes to possess evidence E, should end up with a posterior credence in p, $c_p(p)$, such that $c_p(p) = c_i(p|E) = c_i(E|p) \times \frac{c_i(p)}{c_i(E)}$. Since,

$$c_i(E) = c_i(E|p).c_i(p) + c_i(E|\sim p).c_i(\sim p)$$

then the full formulation of Bayes Rule is:

justification,[234] the standard approach is a form of coherentism, where a particular distribution of credence is justified by the available evidence if that distribution has been arrived at through updating prior beliefs by conditionalizing upon the available evidence.[235]

To make this approach more concrete, consider an example. If I am unsure whether a coin is fair (F) or double-headed ($\sim F$), and lack any relevant evidence one way or the other, then I ought not to become certain either way until such time as I have relevant evidence.[236] Moreover, I am justified in adopting any degree of belief within the interval $0 < \mathrm{cr}(F) < 1$.[237] Once I come to have some relevant evidence, say that I tossed the coin and it came up heads, then no matter which initial credence between 0 and 1 I adopted,[238] my distribution will update (using the standard conditionalization principles) such that I adopt a posterior credence in F which is less than my initial credence in F.[239] Given more evidence, perhaps a string of twenty

---

$$c_p(p) = \frac{c_i(E|p).c_i(p)}{c_i(E|p).c_i(p) + c_i(E|\sim p).c_i(\sim p)}$$

[234] See James M. Joyce, "How Probabilities Reflect Evidence," *Philosophical Perspectives* 19, no. 1 (2005): pp. 157-58.

[235] For defences of this view, see Paul Horwich, *Probability and Evidence* (Cambridge: Cambridge University Press, 1982); Colin Howson and Peter Urbach, *Scientific Reasoning: The Bayesian Approach* (La Salle: Open Court, 1989).

[236] We might think of this as a further constraint on distributions of credence which seems to accompany Bayesianism – which is that agents ought to remain open-minded about (i.e. not assign credence 0 to) propositions which are not ruled out by their evidence. I believe that such a constraint is a plausible epistemic norm *in general*, but even if open-mindedness is not ultimately defensible as a general epistemic norm, however, it will remain plausible as a constraint on the formation of security judgements. Alan Hájek and Michael Smithson, "Rationality and Indeterminate Probabilities," *Synthese* 187, no. 1 (2012)

[237] In the view of some Imprecise Bayesians, you ought to have a total credal state which is "spread" over the entire interval. See ibid; James M. Joyce, "A Defense of Imprecise Credences in Inference and Decision-Making," *Philosophical Perspectives* 24, no. 1 (2010); Isaac Levi, "The Paradoxes of Allais and Ellsberg," *Economics and Philosophy* 2, no. 01 (1986), and criticism from Adam Elga, "Subjective Probabilities Should Be Sharp," *Philosophers' Imprint* 10, no. 05 (2010).

[238] It is not quite true that *any* initial distribution of credence would update in this way. Some very extreme distributions of credence – i.e., those which are infinitesimally close to either 0 or 1 – will fail to update to within such a range. I place such formal problems to one side, suggesting merely that there may be some additional "reasonableness" constraints on initial distributions of credence. See Joyce, "A Defense of Imprecise Credences in Inference and Decision-Making," pp. 290-91.

[239] Say that my initial credence in F, $c_i(F) = 0.9$. My initial credence in a single head conditional on F, $c_i(H|F) = 0.5$, and my initial credence in a single head conditional on not-F, $c_i(H|\sim F) = 1$. Following Bayes Rule, I should end up with a posterior credence in p, $c_p(p)$, through the following procedure.

$$c_p(F) = \frac{c_i(H|F).c_i(F)}{c_i(H|F).c_i(F) + c_i(H|\sim F).c_i(\sim F)}$$

$$c_p(F) = \frac{0.5 \times 0.9}{(0.5 \times 0.9) + (1 \times 0.1)}$$

$$c_p(F) = \frac{0.45}{0.55} \approx 0.82$$

heads in a row, the set of posterior credences in F compatible with the evidence ought to converge still further towards zero.[240] In the limit, the distribution of credence which an agent would be justified in possessing (regardless of their initial distribution of credence) may converge upon a precise degree of credence which is uniquely justified.[241]

The basic idea is that given a lack of evidence relevant to p there will be a wide range of justifiable degrees of credence in p. The distributions, c, which an agent would be justified in holding given the available evidence constitute a set, C, of credence distributions which are compatible with the available evidence. As more and more relevant evidence becomes available, the set of justifiable distributions of credence will begin to converge according to the rules of conditionalization. In this sense, the specificity of a Bayesian approach (i.e. what it counts as justified and what it does not) is reliant on the specificity of the available evidence.

5.2 First Moves

With these preliminaries out of the way, I will start by suggesting that the distinction between fact-relative and evidence-relative security is similar to the distinction between objective and subjective reasons. A subjective reason for an agent to act in a particular way is typically analysed as something that would be an objective reason for the agent to act in that way if the agent's beliefs were, in actual fact, true.[242] Thus, an agent has a subjective reason to drink a glass which they believe to contain gin, even if that glass actually contains petrol (such that they do not have an objective reason to drink it).[243] We might say something similar about security.

For all $c_i(F)$ greater than zero and less than one, $c_p(F) < c_i(F)$.

[240] Even given our initial, quite strong, prior of $c_i(F) = 0.9$, twenty heads in a row gives us strong reason to converge upon a very small credence in F. Since $c_i(20H|F) \approx 9.54 \times 10^{-7}$ and $c_i(20H|{\sim}F) = 1$, then $c_p(F) \approx \frac{8.58\times10^{-7}}{8.58\times10^{-7}+0.1} \approx 8.58 \times 10^{-6}$

[241] Of course, some Objective Bayesians believe that there is almost always a uniquely justified credence in p, since in instances of limited evidence, agents ought to adopt a principle, such as the Principle of Insufficient Reason to provide them with an "objectively determined" credence. I will discuss this objection to my view at the end of this chapter.

[242] There is debate over exactly how the relation between subjective and objective reasons should be analysed with a counterfactual or material conditional: i.e. for some authors, agents have a subjective reason to Φ when they would have an objective reason to Φ if their beliefs were true, whereas for other authors an agent has a subjective reason to Φ when their beliefs *entail* that they have an objective reason to Φ. See Schroeder, "Having Reasons," p. 67; Eric Vogelstein, "Subjective Reasons," *Ethical Theory and Moral Practice* 15, no. 2 (2012): p. 243.

[243] This example originally comes from Bernard Williams, *Moral Luck* (Oxford: Oxford University Press, 1981).

> **ERS:** The evidence-relative security of a proposition p is x just in case were those beliefs justified by the available evidence true, then the fact-relative security of p would be x

Essentially, ERS aligns the evidence-relative security of p with the fact-relative security of p in the possible world nearest to the actual world in which S's justified beliefs are true. Consider the following case:

> COIN TOSS: A coin is about to be tossed – if it comes up heads then you will win a prize and if it comes up tails then you will be asked to pay a small fine. You are certain, having been told by a reliable source, that the coin is fair. Unbeknownst to you, the coin is unfairly weighted towards tails such that the chance of it coming up heads on any given toss is 1/100.

If what you believed about the coin was true (i.e. that the coin was fair), then the fact-relative security of you winning the prize would be ½. No matter that the prize is actually much less fact-relatively secure, in the world nearest to the actual world where all of your beliefs are true, the fact-relative security of the prize is precisely ½.

Whilst ERS captures our intuition that evidence-relative security is, in some sense, determined by the beliefs which are compatible with the available evidence, it is not clear that evidence-relative security should be determined by any and all of those beliefs which are justifiable given the available evidence. Consider a modified coin toss:

> COIN TOSS II: *Unbeknownst to you* a coin is about to be tossed; if it comes up heads you will win a prize, if comes up tails you will be asked to pay a small fine. The coin is unfairly weighted towards tails such that the chance of it coming up heads on any given toss is 1/100. You have a large set of justified beliefs, including that Essendon won the premiership last year. All of you beliefs are true, except for the fact that Collingwood won the premiership.

In the world nearest to the actual world where the agent's beliefs are true, the fact-relative security of the prize is 1/100 and thus, according to ERS, the evidence-relative security of the prize is also 1/100. This seems absolutely incorrect, however, since if you have no beliefs about the coin – you are not even aware of it being tossed – how can you have an opinion about the security of the prize? The problem is that nothing within ERS requires that an agent's evidence have a connection with the fact-relative security of the relevant proposition. This problem suggests that only some features of an agent's doxastic state, those related in some way to fact-relative security, should

matter in determining evidence-relative security, and that our account should identify those beliefs precisely.[244]

Chance that-*p*

To this end, perhaps the most obvious thought is that the evidence-relative security of a proposition should be based on justified beliefs about the fact-relative security of that proposition. Given that fact-relative security is determined by the chance that-*p*, let us begin by positing that the only relevant beliefs are those beliefs about the chance that-*p*.

> **ERS-1:** The evidence-relative security of a proposition p for an agent, S, is x just in case were those beliefs about the chance that-p justified by the available evidence true, then the fact-relative security of p would be x.

If the evidence justifies certainty about the chance that p – e.g. given the evidence, you ought to be certain that $ch(p) = 0.5$ – we can deduce that if the chance that-p really were 0.5 then the security of p would be 0.5.[245] Evidence about chances thus seems to straightforwardly warrant judgements about evidence-relative security.

This approach gives the intuitively correct answer in both COIN TOSS and COIN TOSS II. In the former case, it correctly identifies that the evidence-relative security of the prize as an even split, since the agent knows the relevant chance information. In the latter case, it correctly posits that there is no evidence-relative sense of the security of the prize, since the agent has no beliefs about a coin being tossed. Moreover, it directly connects evidence-relative security with fact-relative security – since the only relevant beliefs are those which are justified by evidence about chances. ERS-1 therefore seems like a powerful modification of ERS and is an initially appealing account of evidence-relative security.

A problem looms for ERS-1, however, which is that whilst agents may not be justifiably certain about the chance that-*p*, they may nonetheless possess a degree of credence in p which is justified by the available evidence. Recall that the chance that-p is the probability of p being true that is enshrined in the laws of the universe: the objective fact of the matter as to which possible futures of a world-time pair are more likely than others to be the actual future. Asked whether you are certain of the chance that it will rain today, most individuals should defer their judgement – how can one be

---

[244] A similar problem for subjunctive analyses of subjective reasons is identified by Vogelstein, "Subjective Reasons,".

[245] We would analyse the counterfactual claim in the standard way: that if in the nearest world where the justified beliefs about the chance that p are true the fact-relative security of p is x, then the evidence-relative security of p is x.

justifiably certain of a property which seemingly requires knowledge of the processes which govern the weather, and all of the contingent facts relevant to those processes? In this sense, it should be rare that agents are justifiably certain of the chance that-p.

Nonetheless, whilst it is rare for an agent to be certain about the chance that-p, agents may nonetheless have indirect evidence about the probability that p will be true. Thus, although I may not be certain of the objective chance of rain is today, or have any beliefs at all about the chance of rain, I may nonetheless possess a justified degree of credence in the proposition that it will rain. Consider the following case, due to James Joyce:[246]

> URN I: You confront an urn containing coins of various biases, β, where β is equal to the chance of that coin landing heads.[247] One of the coins will be drawn at random and subsequently tossed. If the randomly drawn coin comes up heads you will win a prize, if it comes up tails you will be asked to pay a small fine. You are told nothing about the composition of the urn, except that for each coin in the urn of bias β there is a corresponding coin of bias $1 - \beta$. A coin is drawn from the urn and is about to be tossed.

In this case, you have no direct evidence about the bias of the coin which is about to be tossed and so would not be justified in holding beliefs about the chance of tossing a head (except that the chance is within the range $0 \leq \mathrm{ch}(h) \leq 1$). According to ERS-1, therefore, there is no sense in which winning the prize is evidence-relative secure (or not). This seems, however, to underplay the distributions of credence which are justified by the evidence. Indeed, given your evidence that for each coin of bias β within the urn there is a corresponding coin of $1 - \beta$, you ought to be precisely as confident that you will win the prize as not.[248] The insight is that whilst the evidence may not justify beliefs about the chance that a proposition will be true, it may nonetheless justify a determinate degree of belief in a proposition.

---

[246] Joyce, "A Defense of Imprecise Credences in Inference and Decision-Making," p. 289.

[247] For simplicity sake, posit that the machine which produces such coins does so perfectly, and that the bias of each coin is measurable in increments of $\frac{1}{100}$.

[248] To see why, consider that the urn may contain only four coins, with β values of $\frac{2}{10}, \frac{8}{10}, \frac{4}{10}$ and $\frac{6}{10}$ respectively. Because all our evidence about the draw suggests it is random, the probability of drawing a certain coin, conditional on this evidence, is ¼ for all coins. Thus, according to standard conditionalization rules, we can determine the credence which an agent ought to assign to winning the prize. $\mathrm{cr}(h) = \left(\frac{2}{10} \times \frac{1}{4}\right) + \left(\frac{8}{10} \times \frac{1}{4}\right) + \left(\frac{4}{10} \times \frac{1}{4}\right) + \left(\frac{6}{10} \times \frac{1}{4}\right) = \frac{20}{40} = \frac{1}{2}$. This will generalise to any n number of coins within the urn and so, despite the fact that you have no information about the actual chance functions your credence should be constrained to ½.

Credence in *p*

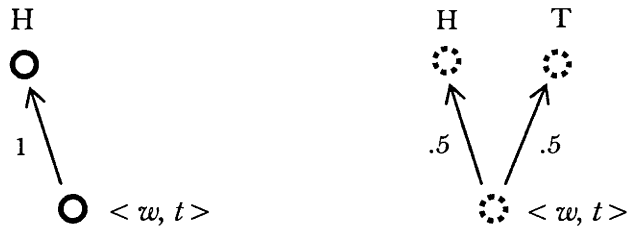
The fact that the evidence may justify a particular credence in a proposition, despite not justifying beliefs about the chance that-*p*, suggests that the account of evidence-relative security ought to be broadened. I suggest the following revision to ERS-1.

> **ERS-2:** The evidence-relative security of a proposition *p* is *x* just in case were the credence in *p* justified by the available evidence true, then the fact-relative security of *p* would be *x*.

Given the credence in *p* which is justified by the evidence available in URN I, ERS-2 suggests that the evidence-relative security of *p* should be ½. This provides a far more accurate picture than ERS-1 of the evidence-relative security of winning the prize.

In order to make sense of ERS-2, however, I must explain how the credence in *p* which is justified by the evidence could be true. To begin with, think of the distribution of credence justified by the evidence as being determined by a particular function, *c*, which assigns to each proposition, *p*, a degree of belief (i.e., $c: p \rightarrow \mathrm{cr}(p)$). Next, think of this credence function as a kind of evidential model of the objective chance function at a particular world-time pair, such that the propositions which an agent believes to some degree (i.e. those compatible with the evidence) model the actual set of possibilities, whilst the agent's credence in those propositions acts as a model of the chance that those possibilities will be actualised. To illustrate how this operates, suppose that the coin which is drawn from the urn in URN II is fully biased towards heads, such that the objective chance that the future will be one where the coin lands heads is 1 (see below). Given the evidence available, however, a rational agent should possess a model of the future which includes both heads-worlds and tails-worlds – each of which are equally likely to be the actual future.[249]

[249] Importantly, though we are taking an agent's credence in *p* to be an evidential model of the chance that-*p*, a credence function may assign a credence *x* to *p* which the agent is certain is not the objective chance that-*p*. Thus, I may know that the coin is either a full tails coin or a fair coin, such that the probability that a coin coming up heads tomorrow is 0.25. We might still say that I possess a evidential model of a set of possible futures where the chance of heads is 0.25, even though I know it to be the case that the actual chance of heads is not 0.25 (i.e. it is either 0 or 0.5).



Fact-relative future Evidence-relative model

Figure 5: The fact-relative and evidence-relative futures of a world w at time t, ($<w, t>$). The probability that the actual future will be a possible future where the coin comes up heads (H) is 1 in the fact-relative case, and 0.5 in the evidence-relative representation. The probability that the actual future will be a possible future where the coin comes up tails (T) is 0 in the fact-relative case and 0.5 in the evidence-relative representation.

It should now be clear what it means for the credence in p justified by the evidence to be true: i.e., that the credence function justified by the evidence coincides with the chance function at a particular world time pair. We can therefore restate ERS_2 in the following way:

> **ERS-2*:** The evidence-relative security of a proposition p is x just in case were the credence function justified by the available evidence the objective chance function, then the fact-relative security of p would be x.

Thus, according to *ERS-2**, the evidence-relative security of p is just the fact-relative security of p in the nearest world in which the evidential model of the future which is justified by the evidence is equivalent to the fact-relative set of possible futures and the distribution of chance amongst those futures.

5.3 Limited and Imprecise Evidence

ERS-2* seems to be a plausible account of evidence-relative security in those circumstances where the evidence is compatible with only one credence function. It is less clear, however, how evidence-relative security should be characterised in situations where an agent's evidence may justify a range of different credence functions. Recall that in instances where there is no evidence relevant to the truth or falsity of a proposition, situations where an agent is ignorant about p, the set of credence distributions compatible with the evidence will include the full interval $0 \leq \mathrm{cr}(p) \leq 1$. Where an agent has a limited amount of relevant evidence, perhaps because the relevant propositions is dependent on rare or unknown events, then the set of credence distributions compatible with the evidence may converge only slightly (say to some interval $a \leq \mathrm{cr}(p) \leq b$). Likewise, where the available evidence is ambiguous, perhaps because its interpretation relies upon complex and contingent social theories, then the set of credence distributions compatible with the evidence may converge upon

many different regions of the full interval $\{0,1\}$.[250] Situations where the available evidence is limited or imprecise are common for most social and political goods, in large part because the probability of retaining such goods relies upon complex and contingent theories. The basic premise is that in many instances, the set of credence distributions compatible with the evidence, C, may include many different distributions of credence. To simplify this discussion, consider a modified version of URN I:

> URN II: You confront an urn containing coins of various biases, β, where β is equal to the chance of that coin landing heads. One of the coins will be drawn at random and subsequently tossed. If the randomly drawn coin comes up heads you will win a prize, if it comes up tails you will be asked to pay a small fine. You are told nothing about the composition of the urn, except that each coin is biased such that $\frac{1}{4} \leq \beta \leq \frac{3}{4}$. A coin is drawn from the urn and is tossed.

URN II presents agents with a very different set of evidence than URN I. It is not just that there is little direct evidence of the precise chance of the coin coming up heads (h), we also lack the kind of symmetry information that was available in the case of URN I. Whilst in URN I the evidence warranted the straightforward adoption of a particular credence in h, our evidence about the composition of URN II leaves open the possibility that it could be filled entirely with coins of bias ¼, entirely with coins of bias ¾ or with potentially any composition of coins of any bias in between.

Given what is known about URN II, there are therefore a very large number of credence functions which are compatible with the available evidence.[251]

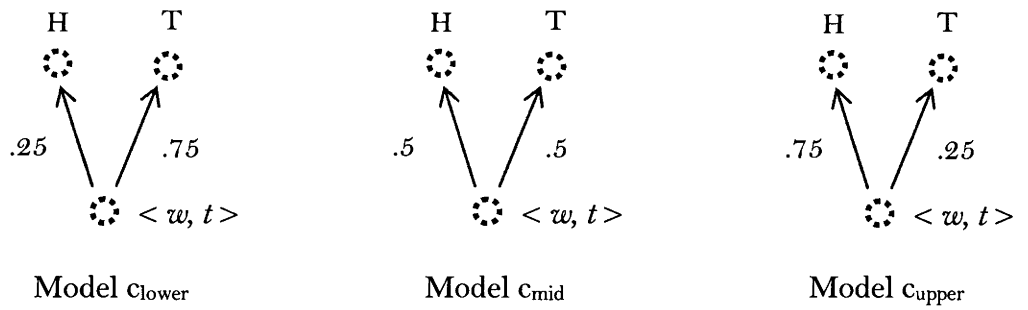
	$\mathrm{cr}(h)$	$\mathrm{cr}(t)$
c_{lower}	0.25	0.75
	...	
c_{mid}	0.50	0.50
	...	
c_{upper}	0.75	0.25

The credence functions, c, which an agent would be justified in holding given the available evidence constitute a set, C, of credence functions which are compatible with the evidence. Where C includes more than one credence function, then there are multiple evidential models of the future which are compatible with the evidence, each of which represents a different objective chance function (see below).

---

[250] It might be the case, for instance, that conditional on the truth of some theory A, the evidence justifies distributions of credence such that $j \leq \mathrm{cr}(p) \leq k$. Whilst, conditional on the truth of some other theory B, the evidence justifies distributions of credence such that $m \leq \mathrm{cr}(p) \leq n$. See Staley, "Robust Evidence and Secure Evidence Claims,".

[251] Some argue that there is only one credence function which is justified by the evidence. I will discuss this position, and its unsuitability for the purposes of assessing evidence-relative security, below.



Model c_{lower} Model c_{mid} Model c_{upper}

Figure 6: A selection of evidential models of the future compatible with the evidence available in URN II. Each credence function justified by the evidence forms part of a evidential model of the future, which can then be used to identify an appropriate world for the counterfactual analysis in ERS_2*.

The credence function c_{lower}, for instance, models a world-time pair where the chance that the actual future is an h-world is 0.25. The credence function c_{upper}, on the other hand, models a world-time pair where the chance that the actual future is an h-world is 0.75. We can think of these models as different world-time pairs upon which the counterfactual analysis in ERS-2* may rest.

In situations where the evidence appears to be compatible with multiple credence functions, and hence multiple evidential models, ERS-2* seems like an inadequate account of evidence-relative security. In the panoply of situations where our evidence is ambiguous, vague or imprecise, ERS-2* does not give any guidance on which of the compatible credence functions ought to ground the counterfactual analysis, and we are therefore left to suppose that the evidence-relative security of a proposition is indeterminate. On a straightforward reading of ERS-2*, it suggests that the evidence-relative security of winning the prize in URN II is somewhere within the interval $\{0.25, 0.75\}$. ERS-2* therefore seems to suggest that there is *no fact of the matter* as to whether the prize is more secure in URN I or URN II.[252] Intuitively, however, there is something determinately different about the evidence-relative security of winning the prize in URN I and URN II. Indeed, since agents are justified in possessing a credence in h of at least ½ in URN I, but may adopt a credence as low as ¼ in URN II, it seems as though the evidence suggests that winning the prize in URN I is more reliable than in URN II. In the next section I will propose a conception of evidence-relative security which seeks to capture this notion, by re-introducing the notion of robustness.

[252] Since it is indeterminate whether $S(h) = ½$ is greater than, less than or equal to the interval $¼ \leq S(h) \leq ¾$.

5.4 Robust Credence

Consider a further account of evidence-relative security: one which determinately differentiates between URN I and URN II by emphasising the intuition that evidence-relative security is identified with a cautious appraisal of the evidence. This proposal suggests that the evidence-relative security of p is determined by the minimum credence in p assigned by all of the credence functions justified by the available evidence. The evidence-relative security of enjoying a good could thus be seen as the safe or reliable probability of enjoying that good, regardless of which credence function justified by the evidence was the most accurate evidential model of the objective chance function. We might therefore adopt the following modification to ERS-2*.

> **ERS-3:** The evidence-relative security of a proposition p is x just in case were any of the credence functions justified by the available evidence equivalent to the objective chance function, then the fact-relative security of p would be <u>at least</u> x.

Essentially, ERS-3 re-introduces the notion of robustness which I discussed in Chapter Four. Recall that in my discussion of Pettit's two conceptions of fact-relative security, I identified that a concern for robustness might motivate the adoption of a protective conception of security. *Protective*$_2$ posited that the security of p was determined by the chance that p regardless of which property, amongst a salient class of properties, obtained. The failing of *Protective*$_2$ was that, for most propositions, the class of properties which should be non-probabilistically privileged was either opaque, or rested on a set of substantive normative commitments incompatible with our ecumenical purpose. There was, therefore, no pervasively salient class of properties which could be non-probabilistically privileged in determining fact-relative security. In the evidence-relative case, however, there appears to be a class of properties that was not present for our impartial observer – variations in the composition of the future and the chance that each possible future might be the actual future. If the available evidence is compatible with more than one credence function, then there are multiple models of the future which are compatible with the future. In this sense, we might think of the evidence-relative security of p as the chance that-p which is robust across the evidential models of the future which are compatible with the evidence.

Unlike the various versions of the protective conception of fact-relative security, there is no mystery as to why particular variations should be non-probabilistically privileged on the account offered by ERS-3. The variations which may be treated as non-probabilistically relevant are those where we lack the information to eliminate all but one credence function. Where we have some information, but not enough to

uniquely identify the credence in p which is justified by the evidence, we privilege all those credence functions which are compatible with the evidence. In this sense, the choice of which variations to privilege non-probabilistically is determined by whether or not our evidence warrants us assigning precise degrees of credence to those variations to begin with.

Objections

There seem to be three major of objections to ERS-3 as an account of evidence-relative security. Let me deal with each in turn.

First, on some views of the relationship between evidence and credence, there is always a uniquely justified credence function for any given set of evidence. Proponents of this view suggest that in circumstances where the evidence is incomplete or imprecise, rational agents should adopt a rule – sometimes called the "Principle of Insufficient Reason" (PIR) – which distributes their credence uniformly across all of the propositions which are equally supported by the evidence.[253] For URN II, one way of applying this thought is that a rational agent ought to distribute their credence equally amongst all of the possible distributions of coins within the urn, such that the uniquely justified credence in h would be ½.[254] Unfortunately, PIR has a dubious reputation. In particular, the use of PIR is famously prone to paradox, since the selection of the appropriate set of hypotheses over which to uniformly distribute your credence may be made on a number of incompatible grounds.[255] This formal problem is normally taken to be devastating, but for those of us concerned with

---

[253] Or, as Roger White puts it, "propositions p and q are evidentially symmetrical for a subject if his evidence no more supports one that the other". Roger White, "Evidential Symmetry and Mushy Credence," *Oxford Studies in Epistemology* 3(2010): p. 161.

[254] To see why, suppose that we simply consider the two most extreme hypotheses about the distribution of coins within the urn, d_1 and d_2, such that $\Pr(h|d_1) = \frac{1}{4}$ and $\Pr(h|d_2) = \frac{3}{4}$. The available evidence does not give agents a reason to be any more confident in d_1 than in d_2, such that agents should assign $\mathrm{cr}(d_n) = \frac{1}{2}$ to each hypothesis. Using standard conditionalization rules, we can determine the credence which an agent ought to assign to winning the prize: $\mathrm{cr}(h) = \mathrm{cr}(d_1).\Pr(h|d_1) + \mathrm{cr}(d_2).\Pr(h|d_2) = \left(\frac{1}{2} \times \frac{1}{4}\right) + \left(\frac{1}{2} \times \frac{3}{4}\right) = \frac{1}{2}$. This will generalise to any countable n number of hypotheses about the contents of the urn and so, despite the fact that you have no evidence about distribution of coins within the urn, your credence should be constrained to ½.

[255] Consider the famous example of a cube factory: You possess a cube which has been manufactured at a factory which produces cubes of widths of less than 2 cm. Apart from this piece of evidence you have no more information about the cube. What ought be your credence in the cube being less than 1 cm wide? The PIR supposes that you ought distribute your credence over the set of hypotheses which are evidentially symmetrical. But which set of hypotheses to choose? One may partition the possibility space according to hypotheses about the width of the cube (0 cm < width < 2 cm), the area of the cube (0 cm² < area < 4 cm²) or the volume of the cube (0 cm³ < area < 8 cm³). Depending on which set of hypotheses you choose to uniformly distribute your credence across, you will either believe that Pr (width > 1 cm) = 0.5, 0.25 or 0.125. See White, "Evidential Symmetry and Mushy Credence," pp. 164-65.

security there is another, less formal reason to reject the use of PIR: which is that it seems to justify epistemic leaps of faith. Recall that evidence-relative security is supposed to reflect features of the future which an agent can rely upon given the available evidence. In URN II, the evidence is compatible with many different distributions of coins within the urn and many different ways in which a particular distribution could be more or less likely.[256] The insistence on a uniform distribution of credence amongst the various distributions of coins seems to ignore the fact that it is entirely compatible with the evidence that an agent may distribute almost all their credence to just one distribution. If all that a belief (or degree of belief) must do in order to be justified is to be coherent with the evidence, then it appears that we should count as justified *any* credence function which can form a coherent set with the evidence.

Second, one might object that ERS-3 allows for propositions and their negations to have degrees of security which violate the probability axioms. According to ERS-3, for instance, the security of winning the prize in URN II is ¼ and the security of not winning is also ¼, and were we to treat these degrees of security as probabilities they would indeed violate the probability axioms. Indeed, evidence-relative security is not a probability measure, since it may violate the complementarity norm, which requires that given a probability x in the truth of A, one ought also believe that the probability that not-A is $1 - x$.[257] In this sense, using the degree to which a proposition is secure as a basis for action may result in inconsistent, or practically irrational actions.[258] To be sure this would be troublesome if one where to treat the evidence-relative security

---

[256] In particular, for every distribution of coins, d_n, the evidence is compatible with $0 \leq \Pr(d_n) \leq 1$.

[257] The probability axioms are generally taken as requiring that for a sample space, X, that includes only two incompatible propositions A and B: (1) Non-negativity: $0 \leq \Pr(A) \leq 1$; (2) Normalization: $\Pr(X) = 1$; and (3) Additivity: $\Pr(A \vee B) = \Pr(A) + \Pr(B)$, such that it directly follows that $\Pr(A) = 1 - \Pr(\neg A)$. See Alan Hájek, "Scotching Dutch Books?," *Philosophical Perspectives* 19, no. 1 (2005): p. 140.

[258] In particular that you are at risk of making a "Dutch Book". Suppose that the prize on offer in URN II was $100, and that you were a fairground worker concerned to make a profit of willing punters. Unfortunately, your employer neglected to tell you the actual composition of the urn, and so your guess as to the probability of winning is as good as any other persons. Suppose Jon came along and, being a terrible gambler, offered to give you $30 to play the game. Given how secure winning the prize appears to be, it seems that the odds are in your favour and you should take this sucker's money and let him play the game. Rosa, a more cunning gambler, simultaneously sidles up to you and offers you a side-bet: "If Jon wins", she says, "I will give you $30, but if Jon loses then you must give me half the prize-money." Happy days! Given the security of the proposition that Jon will *not* win (i.e ¼) then the odds on this bet are also (very much) in your favour and you should accept. Despite the initially appealing reasoning which brought you to the conclusion that you should accept both bets simultaneously, you now have a payoff situation that guarantees a loss. See the explanation, and criticism of, the Dutch Book argument in Alan Hájek, "Arguments for–or against–Probabilism?," *The British Journal for the Philosophy of Science* 59, no. 4 (2008); Hájek, "Scotching Dutch Books?,".

of a proposition as the odds at which one should determinately accept a wager. Yet nothing we have discussed so far suggests that the security of propositions should be a guide to practical action. On the contrary, in Chapter Six I will argue that the evidence-relative security of a proposition should never be the sole consideration in decision-making problems, even if it may form part of a pluralistic decision framework. Unless we believe that the degree to which an agent believes a proposition to be secure should be his sole guide to action, then ERS-3 escapes the Dutch Book argument.

Third, one might believe that ERS-3 does not adequately represent where the weight of the evidence lies, since very extreme initial credences in a proposition (i.e. those infinitesimally close to 0 or 1) may result in a posterior credence which is well below or well above that assigned by the bulk of credence distributions within the set compatible with the evidence *C*. If evidence-relative security is based upon the minimum degree of credence which is assigned by *any* function within *C*, then those extreme distributions might seem to have a disproportionate influence on the evidence-relative security of a good. There appear to be two solutions to this problem. The first is to include some constraints on which initial distributions of credence are *reasonable*, such that those extreme distributions which would be insensitive to the evidence would be deemed unreasonable and excluded from the beginning.[259] On the other hand, one might wish to suggest that the reasonability criterion should be built into the account of evidence-relative security, such that the credence in p which determined the evidence-relative security of p need only be robust across a reasonable sub-set of the credence functions compatible with the evidence.[260] I am sympathetic to both these views, but wish to note that what counts as reasonable is likely to be highly dependent on moral considerations. Consider that we might wish to include some very extreme priors in determining the evidence-relative security of the Large Hadron Collider producing a world-devouring black-hole,[261] but that excluding extreme priors seems reasonable in determining whether or not it will rain today. ERS-3 is designed to capture the credence in p – reasonable or otherwise – which can be relied upon given the evidence, and abstracts away from such moral considerations. In Chapter Six I will illustrate how the weight one accords to evidence-relative security in a

---

[259] For a proposal along these lines, see Luis Raúl Pericchi and Peter Walley, "Robust Bayesian Credible Intervals and Prior Ignorance," *International Statistical Review / Revue Internationale de Statistique* 59, no. 1 (1991): pp. 3-5.

[260] i.e., **ERS-4**: The evidence-relative security of a proposition p is x just in case were any of a reasonable sub-set of the credence functions justified by the available evidence equivalent to the objective chance function, then the fact-relative security of p would be at least x.

[261] Max Tegmark and Nick Bostrom, "Astrophysics: Is a Doomsday Catastrophe Likely?," *Nature* 438, no. 7069 (2005).

pluralistic decision-making process might be a better space for capturing such reasonableness considerations.

Conclusion

I have argued that ERS-3 is a plausible account of the evidence-relative security of a proposition. That account was developed by considering a series of competing accounts alongside various cases where agents had access to different evidence about the fact-relative security of a proposition. Whilst ERS-1 provides a straightforward link between fact-relative and evidence-relative security, it underplayed the extent to which an agent's evidence could constrain the probability of a proposition being true. ERS-2, whilst providing a straightforward account of the link between an agent's evidence and evidence-relative security, violated the intuition that evidence-relative security should be based on a cautious reading of the evidence available to an agent. ERS-3, on the other hand, has three major virtues.

First, ERS-3 captures the intuition that security is associated with a cautious disposition towards the world. ERS-3 aligns the evidence-relative security of p with the fact-relative security of p in the possible world nearest to the actual world in which the "most cautious" of the evidential models is true.[262] An agent who contemplates the security of a good according to ERS-3 is concerned to identify only those evidential models of the future which are incompatible with the evidence they have before them. Whilst they may use heuristics, rules-of-thumb or hunches to settle upon a particular credence function within the set of those compatible with the evidence, an agent has reason to take into consideration the degree of confidence that the agent may determinately possess given the evidence.

Second, ERS-3 is generalizable to both instances where there is direct evidence about chances and situations where agents lack any relevant evidence. In situations where the evidence includes the objective chance that-p, then agents ought to possess a degree of credence in p which matches the chance that-p.[263] In such situations, ERS-3 correctly identifies that the evidence-relative security of p ought to be equivalent to the fact-relative security of p. Likewise, in situations where an agent's evidence does not contain any information relevant to the truth of p, ERS-3 gives the intuitively

---

[262] When I say "cautious" I don't mean that if that evidential model were true it would necessarily involve the least desirable outcome. The idea is that for any proposition, p, the most cautious evidential model will be the model which gives p the lowest chance of being true of any of the models compatible with the evidence.

[263] This dictum is similar to David Lewis' Principal Principle. See Lewis, "A Subjectivist's Guide to Objective Chance," ; Lewis, "Humean Supervenience Debugged,".

plausible suggestion that p is, from an evidence-relative perspective, completely insecure. For these reasons, ERS-3 suggests an intuitively plausible ordering of the evidence-relative security of winning the prize in COIN-TOSS (i.e., ½), URN I (i.e., ½) and URN II (i.e., ¼).

Finally, evidence-relative security, as captured by ERS_3, seems to capture a distinctive, and potentially important, aspect of moral decision-making. For reasons explored by many others,[264] the standard decision theoretic rule to perform the act which maximises expected utility does not allow agents to perform what appear to be rational decisions in circumstances where the evidence is limited or ambiguous. Whilst fact-relative security is straightforwardly compatible with standard accounts of expected utility theory, the putative importance of evidence-relative security might offer a way of capturing the intuitive importance of avoiding such epistemic risks. I will explore these issues further in Chapter Six.

[264] Peter Gärdenfors and Nils-Eric Sahlin, "Unreliable Probabilities, Risk Taking, and Decision Making," *Synthese* 53, no. 3 (1982); Joyce, "A Defense of Imprecise Credences in Inference and Decision-Making,"; Isaac Levi, "Ignorance, Probability and Rational Choice," *Synthese* 53, no. 3 (1982); Hájek and Smithson, "Rationality and Indeterminate Probabilities,".

Chapter Six
Security and Goodness

In this chapter I consider whether there are reasons to regard security as morally valuable; either as a feature of states of affairs or as a consideration in how agents ought to act. Intuitively, security seems valuable in two ways. First, security might be a feature of a state of affairs which is (at least partly) constitutive of the moral goodness of that state of affairs. Consider that, *ceteris paribus*, a state of affairs where important goods are enjoyed securely seems intuitively more valuable than a state of affairs where those goods are not so secure. This gives us a *prima facie* reason to suppose that security might itself be a property which adds to the final value of a state of affairs. Second, even if security is not constitutive of the goodness of a state of affairs, it may be constitutive of the goodness of the acts which an agent could perform. Consider that, *ceteris paribus*, an act which increases the degree to which a good is securely enjoyed is intuitively better than an act which does not. As I will explore, this intuition has particular importance in the context of evidence-relative moral decision-making.

In what follows I argue that the moral importance of security is located in its role as a consideration in how agents ought to act. In section 6.1, I introduce what I call the standard teleological model, which suggests that the goodness of an act is reducible to the goodness of the possible futures which that act may bring about. This model suggests that secure enjoyment of a good may only be valuable in so far as it promotes the actual enjoyment of a good. In section 6.2, I explore security as a component of the goodness of a world, noting that the value of security is dependent on the value of the good being secured. I go on to argue that though it might seem plausible that the security of finally valuable goods contributes to the goodness of a world, it is difficult to adequately explain this intuition. In section 6.3, I argue that, *ceteris paribus*, the fact-relative security of goodness is the primary consideration in determining the fact-

relative goodness of acts. Moreover, drawing on utility theory, I suggest that the importance which agents place on the fact-relative security of each degree of goodness models their attitude towards risks posed by an indeterminate future. In section 6.4, I turn to consider the role that evidence-relative security plays in modelling an agent's attitude towards the risks posed by the limits of their evidence. In particular, I argue that the importance placed on the degree of expected utility which is evidence-relative secure, as opposed to the optimal degree of expected utility, might plausibly form part of a pluralistic framework for decision-making under conditions of limited evidence.

6.1 The Standard Teleological Model

In order to illustrate security's distinctive role in moral decision-making, I will initially adopt what I call the standard teleological model.[265] This model is defined by three theses about the nature of moral decision-making.

First, the model holds that the deontic status of an act – i.e., whether it is permissible, obligatory, right etc. – is a function of the goodness of that act.[266] Broadly speaking, a teleological ethics holds that acts can be ordered according to their goodness, and that the deontic status of an act is determined by their place in that ordering. The standard model is a maximising teleology, such that an agent is morally obliged to perform that act which is best, but we can imagine a satisficing teleology which suggests that any sufficiently good act is permissible.[267]

Second, the model holds that the goodness of an act is a function of the goodness of the possible futures where that act is performed.[268] There are many competing theories of what exactly constitutes the goodness of a world, but all agree that in

---

[265] Whilst teleological ethics is a broad church, the basic precepts enumerated here are broadly agreed upon by both adherents and critics. See John Broome, *Weighing Goods: Equality, Uncertainty and Time* (Oxford: Blackwell, 1991), pp. 10-16; Philip Pettit, "Consequentialism," in *A Companion to Ethics*, ed. Peter Singer (Oxford: Blackwell Publishers, 1991), pp. 18-19; Scanlon, *What We Owe to Each Other*, pp. 79-81.

[266] This view of the connection between the deontic and evaluative status of an act is at the crux of the debate between teleologists and non-teleologists. See Broome, *Weighing Goods: Equality, Uncertainty and Time*, pp. 11-16; W.D. Ross, *The Right and the Good* (Oxford: Oxford University Press, 1930), pp. 8-11; Scanlon, *What We Owe to Each Other*, pp. 95-99; Michael Stocker, "Rightness and Goodness: Is There a Difference?," *American Philosophical Quarterly* 10, no. 2 (1973); Michael J. Zimmerman, "The Good and the Right," *Utilitas* 19, no. 3 (2007).

[267] See Jean-Paul Vessel, "Supererogation for Utilitarianism," *American Philosophical Quarterly* 47, no. 4 (2010); Michael J. Zimmerman, *The Concept of Moral Obligation* (Cambridge University Press, 1996), p. Ch. 8.

[268] This should not be confused with the consequentialist thesis that the goodness of an act is determined only by its consequences independent of the value of performing that act. A teleologist can hold, for instance, that the intrinsic badness of the act of torturing the innocent diminishes the goodness of a world, even if the consequences of torturing the innocent would bring about some stupendous state of affairs. See Broome, *Weighing Goods: Equality, Uncertainty and Time*, p. 4; Scanlon, *What We Owe to Each Other*, p. 80.

principle every possible world has a degree of goodness determined by the finally valuable components of that world.[269] In this sense, a teleological ethics is concerned to *promote* the realisation of those states of affairs which are valuable as ends-in-themselves. This may be contrasted with a non-teleological ethics, where the right act must *respect* or honour a particular value, such as autonomy, even where by so doing it may bring about a world where others fail to honour that value.[270]

Third, where an act does not fully determine which future will arise, the standard teleological model holds that the goodness of an act is determined by the goodness of each possible future where that act is performed, weighted by the probability that that future will be the actual future. Importantly, the standard model suggests that this procedure for handling indeterminacy does not itself instantiate particular value commitments.[271] In this sense, the standard model holds that all of the relevant *values* to consider in determining the goodness of acts are features of the goodness of possible worlds. Whilst there may be other considerations, such as the probability of those futures being realised, they are considerations which are merely instrumental to the promotion of what is actually valuable.

According to the standard teleological model, security might play a role in moral decision-making in one of two ways: (i) as a component of the goodness of a world, or (ii) as a component in the neutral decision procedure which one may adopt in order to determine the best act. If security is best conceived as a component of the former, then the secure enjoyment of a good may have some final value beyond the mere promotion of the enjoyment of that good. If, however, security is a component of the latter (neutral) decision procedure, then on the standard model the secure enjoyment of a good will only be valuable in so far as it promotes the enjoyment of that good. As we shall see, contemplating the role of security in moral decision-making calls into question the view that the goodness of an act is reducible to the goodness of the possible worlds where that act is performed.

---

[269] See Fred Feldman, "On the Consistency of Act- and Motive-Utilitarianism: A Reply to Robert Adams," in *Utilitarianism, Hedonism, and Desert*, ed. Fred Feldman (Cambridge: Cambridge University Press, 1997), p. 72.

[270] See Pettit, "Consequentialism," p. 17, and the critique of promotional accounts of value in Scanlon, *What We Owe to Each Other*, pp. 83-95.

[271] see Mark Colyvan, Damian Cox, and Katie Steele, "Modelling the Moral Dimension of Decisions," *Nous* 44, no. 3 (2010), and a critique of the insistence on value-neutral decision rules in Sven Hansson, "The Harmful Influence of Decision Theory on Ethics," *Ethical Theory and Moral Practice* 13, no. 5 (2010).

6.2 Security and the Best Possible Future

Before considering whether security contributes to the goodness of a possible world, let me establish some preliminary observations about the nature of the value of security.

First, we are concerned in this section with whether security is finally, rather than merely instrumentally, valuable. Roughly speaking, finally valuable goods are those goods which are valuable as "ends-in-themselves", whilst instrumentally valuable goods are those goods which are only valuable in so far as they promote or maintain a finally valuable good.[272] If security is to contribute to the goodness of a state of affairs then it must be valuable as an end in itself and not merely valuable in so far as it promotes the enjoyment of a good.

Second, whilst security may be finally valuable, it is not *intrinsically* valuable. Whilst it is a potentially controversial claim, I take it that something may be an end-in-itself by virtue of its intrinsic properties or by virtue of a relation between that thing and some other thing.[273] The property of security is certainly not intrinsically valuable, since as a measure of the truth of a proposition in the future, it has little meaning without reference to that proposition. In this sense, we are not so much investigating the value of security *per se* as the value of the security of p – where p is some morally relevant proposition.

Third, that p is itself finally valuable is a necessary condition for the security of p to be finally valuable. Consider that the security of the proposition "all babies born tomorrow will be tortured horribly" seems very dis-valuable by virtue of the dis-value of the proposition being secured. Moreover, consider the secure enjoyment of an instrumental good such as money. If enjoying money is merely instrumental to the realisation of other goods (even if it is necessary), then it is hard to see how the *secure* enjoyment of that money could be finally valuable. Suppose that $10 will straightforwardly give you ten units of happiness, then having a high chance of $10 will give you a high chance of ten units of happiness. The secure enjoyment of an

---

[272] Of course, there is much more nuance to these categories than these simple definitions suggest. For instance, we ought to recognise a distinction between wholly instrumental, wholly final (ultimate) goods and mixed goods. A great work of art, in so far as it makes people happier, has some instrumental value, but, in so far as its value is not entirely explainable by gains in people's happiness, it might also have some final value.

[273] For instance, historical artefacts, the protection of which many hold to be an end-in-itself, are only finally valuable by virtue of their relation to an historic event. See Fred Feldman, "Hyperventilating About Intrinsic Value," *The Journal of Ethics* 2, no. 4 (1998); Shelly Kagan, "Rethinking Intrinsic Value," *Journal of Ethics* 2, no. 4 (1998).

instrumental good therefore seems to be merely instrumental to the secure enjoyment of some final good.

Finally, to consider whether the security of a final good is constitutive of the goodness of a state of affairs I will consider whether security contributes to the goodness of particular possible worlds. This is both because it accords with the framework I developed in Chapter Four and Five, and because it will be difficult to determine the independent value of security without controlling for its role in the realisation of future goods within a world. Consider a proposition – Xoli will enjoy pleasure at time t – which identifies a finally valuable state of affairs.[274] Clearly, the security of this proposition at time t-1 is *instrumentally* valuable to the realisation of the finally valuable state of affairs at time t, but it is difficult to see whether or not the security of the proposition is itself finally valuable. Considering whether the security of finally valuable goods contributes the goodness of whole possible worlds, on the other hand, eliminates the instrumental role which security plays in the realisation of finally valuable states of affairs. As wholly temporally extended entities, all of the instrumental linkages within possible worlds have been realised, and in this sense, differences in the goodness of possible worlds will be differences in final value.

With these preliminary observations out of the way, I will now consider whether there is reason to suppose that security is finally valuable. To do so, I will employ a simple contrast strategy, fixing all the facts about two possible worlds expect for the degree to which the finally valuable goods within those worlds are secure. Consider the following:[275]

> **World I** is a small planet inhabited by moral agents. Life on this planet is idyllic, with no war, hunger or inequality, and much beauty, pleasure and wisdom besides. Life on World I persists in its idyllic state for five years before it is instantaneously obliterated by the explosion of a nearby star.

[274] I am assuming, in this instance, that this state of affairs forms part of a Moorean "organic unity." We needn't contemplate whether the value of the state of affairs "Xoli is experiencing pleasure at midday" might be different depending on whether Xoli's pleasure is derived from a job well done or some despicable act. See G. E. Moore, *Principia Ethica* (Cambridge: Cambridge University Press, 1903), pp. 27-28.

[275] A similar example appears in Michael J. Selgelid, "The Value of Security: A Moderate Pluralist Perspective," in *Ethics and Security Aspects of Infectious Disease Control: Interdisciplinary Perspectives*, ed. Christian Enemark and Michael Selgelid (Burlington: Ashgate, 2012).

> **World II** is almost identical to World I. Also inhabited by a population of moral agents, life on this planet is idyllic in just the same way as it is on World I. Unbeknownst to anyone on World II, however, it is surrounded by a large number of asteroids. During any given year, there is a very high objective chance (i.e., 9/10) that the orbit of one of these asteroids will degrade, colliding with World II and permanently destroying it. Luckily, despite the high chance of catastrophe, no asteroid ever collides with World II. Life on World II persists in its idyllic state for five years before it is instantaneously obliterated by the explosion of a nearby star.

The difference between World I and World II is a difference in their fact-relative security. The history of these two planets is indistinguishable: at each and every moment they are filled with exactly the same amount of beauty, pleasure, wisdom and inequality. The only thing which differentiates each moment on World I from each moment on World II is the set of possible futures at that moment. In this sense, they form a good test case of whether we think security is finally valuable.

I take it that the most common intuitive response to this case will be that there is no discernible difference between the goodness of World I and the goodness of World II. The basic idea behind this intuition is that given the knowledge that lives lived on either planet will, in actual fact, be indistinguishable from one another, many people will be indifferent between living a life on World I and living a life on World II. The intuition is clearest if one harbours a subjectivist view of final value (i.e., that the goodness of a possible world is reducible to the subjective experiences of the agents within that possible world),[276] but it also seems compatible with most person-affecting views of final value (i.e., that the goodness of a possible world is determined by whether or not it is good for the agents within that world).[277] Indeed, if one holds the view that the goodness of a possible world is determined only by what actually occurs to agents in that world, then living a life on World I would be equivalent to living a life on World II.[278] Whilst I share this intuition, some might nonetheless believe that World I is preferable to World II. This may be for two reasons.

---

[276] This is most obvious on a straight-forward mental-states view of goodness, but it is also compatible with actual and informed desire accounts satisfaction accounts. For an overview of many different subjectivist accounts of goodness, see Griffin, *Well-Being: Its Meaning, Measurement and Moral Importance*, pp. 1-39.

[277] The most well-known version of this view is that given by Derek Parfit, who ultimately rejects it. Derek Parfit, *Reasons and Persons* (Oxford: Clarendon Press, 1984), p. 363.

[278] Of course, more sophisticated person-affecting accounts might hold that some goods are modally demanding, in the sense that their existence in the actual world relies upon facts about what occurs in some relevant set of merely possible worlds. In a similar vein one might hold that risks to a person are bad for that person, regardless of whether the risk materialises. These views are discussed below.

Expected value as good-in-itself

First, one might hold the view that the *expected* value of the future of each moment of time within a world contributes, in some way, to the goodness of that world. For each moment of time, *t*, the future of *t* will have some fact-relative expected value, i.e, the chance-weighted sum of the value of the possible futures of *t*.[279] At any moment in World I (except the last) it is objectively certain that World I will continue to be idyllic, but at any moment in World II (expect the last) there is a very high chance that World II will suffer a catastrophe. The expected value of the future of any moment within World I will therefore be greater than the expected value of any moment within World II. In this respect, at any moment before the end of each world, living on World I would appear to be preferable to living on World II.

This helps explain the intuition that World I is better than World II in the following way. Suppose that we divide on World I and World II coarsely, such that each moment of time amounts to precisely one year on the planet. Suppose further that each year is indistinguishable from any other, such that a year on each world generates precisely v units of what we might call underlying final value.[280] Given these assumptions, we can calculate the expected value of the future at the conclusion of each moment in time on each world (see table below).[281] If the thesis that the expected value of each moment contributes to the goodness of that world is true, then it supports the intuition that World I is more finally valuable than World II.[282]

WORLD I	t_1	t_2	t_3	t_4	t_5
Final Value	v	v	v	v	v
Expected Value	$4v$	$3v$	$2v$	v	0

[279] The expected value of the future of a world-time pair <*w*, *t*> is given, in the traditional way, by the chance-weighted sum of the value, *v*, of each possible future, *f*, of that world-time pair: $EV(w,t) = \sum \mathrm{ch}(f).v(f)$.

[280] The underlying value of a moment is simply the final value of that moment independent of the expected value of the future of that moment. The underlying value of a world-time pair will derive from things which are *true* at that world-time pair – i.e. particular experiences of moral agents, such as instances of pleasure, or the existence of certain finally valuable objects. Importantly, however, those things which are true at that world-time pair need not be *intrinsic* to that world-time pair. For instance, the value of a particular instance of pleasure might be different depending on whether it forms part of an extended experience of pleasure which spans across multiple world-time pairs within the actual world. Moreover, the value of making a promise that you will in fact fulfil might contribute to the underlying value of t_1, even if that value is by virtue of facts about t_2. The idea is that whatever the underlying value of each moment, it represents the final value of that moment which is not attributable to the expected value of the future of that moment.

[281] For World I, the expected value of the future at the conclusion of the *n*th year is given by $EV(n) = 1 \times (5-n)v$.

[282] Assuming that the goodness of each world *in toto* is simply the sum of the final value and the expected value of each moment, then the goodness of World I is $15v$, whereas the goodness of World II is $5.4321v$.

WORLD II	t_1	t_2	t_3	t_4	t_5
Final Value	v	v	v	v	v
Expected Value	$0.1111v$	$0.111v$	$0.11v$	$0.1v$	0

Unfortunately, this explanation of the initial intuition leads to two unappealing upshots.

First, we can increase the total goodness of the world merely by more finely dividing the moments of time within that world.[283] Suppose that we arbitrarily carve up World I into ten parts (instead of 5), and distribute the value accrued such that each moment in time has actual value $v/2$ (see table below).

	t_1	t_2	t_3	t_4	t_5	SUM
Final Value	v	v	v	v	v	5v
Ex. Value	$4v$	$3v$	$2v$	v	0	$10v$

	t_1	t_2	t_3	t_4	t_5	t_6	t_7	t_8	t_9	t_{10}	SUM
Final Value	$v/2$	$v/2$	$v/2$	$v/2$	$v/2$	$v/2$	$v/2$	$v/2$	$v/2$	$v/2$	5v
Ex. Value	$4.5v$	$4v$	$3.5v$	$3v$	$2.5v$	$2v$	$1.5v$	v	$v/2$	0	$22.5v$

On any view where the expected value of each moment in time contributes proportionately to the goodness of a world, then the goodness of a world will increase as the number of temporal parts increases.[284]

Second, if the expected goodness of each moment in time contributes to the goodness of the world, then the *order* in which those moments occur will alter the goodness of the world. Consider two further deterministic worlds, III and IV:

WORLD III	t_1	t_2	t_3	t_4	t_5	SUM
Final Value	v	v	v	$2v$	$5v$	$10v$
Expected Value	$9v$	$8v$	$7v$	$5v$	0	$29v$

WORLD IV	t_1	t_2	t_3	t_4	t_5	SUM
Final Value	$5v$	$2v$	v	v	v	$10v$
Expected Value	$5v$	$3v$	$2v$	v	0	$11v$

In World III, the underlying value of each moment increases over time, whereas in World IV, the underlying value of each moment decreases over time. On the view that expected value contributes to the goodness of a world, World III is determinately

[283] I was assisted in seeing this by Robert Kirby.

[284] One might suppose that there is some "basic" number of temporal parts to a world. If this is the case then this problem may lose much of its force.

better than World IV, despite the fact that they contain (and are determined to contain) equivalent amounts of underlying value. That we can alter the goodness of a world merely be re-ordering its parts seems intuitively wrong, since two deterministic worlds with equivalent amounts of underlying final value ought to be equivalently good.

There therefore seem to be serious problems with explaining the intuition that World I is better than World II with an appeal to the view that the expected value of each moment in time within a world contributes to the goodness of that world. Whilst these problems may not be insurmountable, the onus is upon those seeking to explain their intuition in this way to provide a detailed axiology which deals sensibly with these problems.

Modally-demanding futures

Another way in which one might seek to explain the intuition that World I is determinately better than World II is by appeal to the idea that the goodness of a world is not determined solely by what actually occurs in that world, but also by what occurs in other merely possible circumstances. Recall Pettit's claim that in order to enjoy the good of freedom, for instance, one must not only enjoy the non-interference of the powerful in the actual world, but also in those possible worlds where they cease to be favourably disposed towards you. Whilst there are many goods which are not modally demanding in this sense, there are a large number of goods which are: e.g. friendship, dignity, and the rule of law.[285] To count someone as a friend, for instance, it must be the case that this friend exhibits the modally undemanding good of concern for your interests, not just in the actual world, but in some salient class of possible worlds. Thus, the content of a world, and hence its goodness, might be influenced by what would counterfactually occur in worlds nearby to that world.

Is the security of *p* a good of this kind? Well, it is not straightforwardly so, since I argued in Chapter Four that the secure enjoyment of a good is simply the chance that the actual future will be one where you enjoy that good. Pettit's claim, however, is that in order to enjoy a modally-demanding good, like freedom, it must be the case that one enjoys its modally-undemanding counterpart (in this case non-interference) in a range of counterfactual circumstances regardless of their likelihood of being realised. One's chance of being free from interference in the future is not what is at issue – it is whether or not you would presently be free from interference were some relevant facts

[285] See List, "Republican Freedom and the Rule of Law,"; Pettit, "Making Good: The Challenge of Robustly Demanding Values,".

altered (and regardless of whether they were counterfactually likely). One can think of this as a difference in the perspective from which the analysis occurs: to examine the modal robustness of a good, examine the way the world could have been right now; to examine the security of a good, examine the way the world might be in the future. To be sure, the security of the modally undemanding counterpart (e.g. non-interference) of a modally demanding good (e.g. freedom) may give us an insight into whether or not the modally demanding good is actually being enjoyed. This, however, does not argue for the contribution of security *per se* to the goodness of a world, merely security's usefulness as a guide to the goodness of that world.

In this respect, I think it is unlikely that holders of the intuition that World I is better than World II will be able to appeal to considerations of modal robustness in defence of their view. Perhaps there is some insight to be gleaned from the notion of modally-demanding goods, which could then be applied in a more thorough defence, but it is not straightforwardly applicable. Given my belief that most individuals will hold the intuition that World II is just as good as World I, and the difficulty in justifying the intuition that it is not, I will proceed as if the fact-relative security of goods does not contribute to the goodness of possible worlds.

Of course, regardless of whether the secure enjoyment of finally valuable goods contributes to the goodness of a possible world, the secure enjoyment of finally valuable goods might yet be a consideration in determining which acts an agent ought to perform. Recall that the standard teleological model suggests that in situations where it is not clear which future an act will bring about, an agent ought to use a neutral decision procedure to choose which act is best. In what follows I challenge the notion that the goodness of an act is reducible to the goodness of possible worlds. Indeed, as we shall see, contemplating the importance of securing the good requires considering what the best attitude towards two types of risk.

Importantly, just as there is a distinction between fact-relative and evidence-relative measurement of security, there is a distinction between fact-relative and evidence-relative oughts.[286] For the purposes of the next two sections I will

[286] Many argue that there is a third sense of ought, which is what an agent ought do given what they do, in fact, believe (and regardless of whether those beliefs are justified by the evidence). It is plausible that such belief-relative oughts bear on whether we ought blame or praise an agent for their acts, but it seems to me that they do not bear on what constitutes right action. For reasons why, see Zimmerman, *Living with Uncertainty: The Moral Significance of Ignorance*, pp. 13-14. For more on the general distinction see Feldman, *Doing the Best We Can*, p. 46; Derek Parfit, *On What Matters*, vol. I (Oxford: Oxford University Press, 2011), pp. 150-53.

understand the former as the goodness of an act given the actual facts and the latter as the goodness of an act given those beliefs justified by our evidence.[287] If we assume this familiar position, then fact-relative security would seem to be a consideration in fact-relative oughts, whereas evidence-relative security would appear to be a consideration in evidence-relative oughts. In what follows I interrogate security's role in determining the goodness of acts from both these perspectives.

6.3 Fact-relative Security and the Goodness of Acts

Let us begin by interrogating the role of fact-relative security as a consideration in what agents ought to do from the point of view of an impartial observer. Consider the following case:

> **TOXIN:** A patient arrives in the emergency room suffering from a deadly neurotoxin. His doctor has three kinds of antidote available to her. Antidote X is the antidote specific to this neurotoxin and will cure the patient. Antidote Y is an all-purpose but slow-acting antidote, which though it will save the patient's life will result in permanent facial paralysis. Antidote Z is an antidote to an unrelated toxin and will have no effect such that the patient will succumb to the toxin. If no compound is administered the patient will die.
>
> Unfortunately, when the doctor visits the dispensary she finds that the vials of antidote are not properly separated, and are instead stored in two urns, A and B. Urn A contains only Antidote Y. Urn B contains Antidote X and Antidote Z in precisely the ratio of 4:1. The doctor must choose from which urn to draw a vial.

For simplicity sake, let us suppose that there are only three relevant outcomes of TOXIN each of which corresponds to a particular possible future, f_i: (i) the patient dies, f_D, (ii) the patient makes a partial recovery, f_P, or (iii) the patient makes a full recovery, f_R. The chance that each possible future will be the actual future, conditional on taking a vial from a particular urn is given by the following table.[288]

	$ch(f_D \mid \phi)$	$ch(f_P \mid \phi)$	$ch(f_R \mid \phi)$
A	0	1	0
B	$1/5$	0	$4/5$

[287] See Parfit, *On What Matters*, pp. 150-51.

[288] To be perfectly accurate, these are actually *sets* of possible futures – i.e. those futures where, regardless of what else happens, the patient dies, etc. For simplicity sake, I treat these sets as if they are single futures.

Moreover, supposing that the only propositions relevant to the goodness of the actual future are those which describe the outcome for the patient, the security of goodness *itself* will be different depending on the urn from which the doctor draws a vial. To make the situation more concrete, let us assume that the goodness of the possible futures of TOXIN has the following cardinal structure:[289]

	Death (f_D)	Partial Recovery (f_P)	Full Recovery (f_R)
$g(f)$	0	15	20

Recalling that the fact-relative security of a proposition is simply the chance that the actual future will be one where the proposition is true, we can therefore calculate the security of certain degrees of goodness. For instance, if the doctor draws a vial from urn A then fifteen units of goodness will be totally secure. If, on the other hand, the doctor draws a vial from urn B then twenty units of goodness will be secure to degree 4/5. The doctor's choice, therefore, is between drawing from an urn which will fully secure a less desirable outcome, or drawing from an urn with the potential for a more desirable, but less secure, outcome.

Expected goodness and security

On the standard teleological model, the importance of security is incorporated into decisions by comparing the *expected* goodness of each act. Roughly speaking, the fact-relative expected goodness of an act φ is the sum of the chance-weighted goodness of the possible futures, f, conditional on φ.[290] This can be formalised as follows:

$$EG(\phi) = \sum \mathrm{ch}(f|\phi).g(f)$$

In this sense, the expected goodness of choosing compound A is given by:

$$\begin{aligned} EG(A) &= \mathrm{ch}(f_D|A).g(f_D) + \mathrm{ch}(f_P|A).g(f_P) + \mathrm{ch}(f_R|B).g(f_C) \\ &= (0 \times 0) + (1 \times 15) + (0 \times 20) \\ &= 15 \end{aligned}$$

[289] One might doubt that goodness has a cardinal structure at all, but I place this issue to one side for now. See Broome, *Weighing Goods: Equality, Uncertainty and Time*, pp. 91, 121-50.

[290] Those familiar with decision theory will note that this is similar to Jeffrey's conditional decision theory. As Lewis and others have pointed out, this conditional formula is susceptible to Newcomb's Paradox in evidence-relative cases. I will not pursue the matter here, but the various formal solutions to Newcomb's Paradox, including causal decision theory ought to be compatible with my account of the importance of fact-relative and evidence-relative security. See Richard C. Jeffrey, *The Logic of Decision* (Chicago: University of Chicago Press, 1990); David Lewis, "Causal Decision Theory," *Australasian Journal of Philosophy* 59, no. 1 (1981).

Likewise, the expected goodness of administering compound B is given by:

$$\begin{aligned} EG(B) &= \mathrm{ch}(f_D|B).g(f_D) + \mathrm{ch}(f_P|A).g(f_P) + \mathrm{ch}(f_R|B).g(f_C) \\ &= \left(^1/_5 \times 0\right) + (0 \times 15) + \left(^4/_5 \times 20\right) \\ &= 16 \end{aligned}$$

The standard teleological decision model enjoins agents to *maximise* expected goodness.[291] On this account, an option φ is best if, and only if, its expected goodness is at least as great as any other option φ available to the agent. Given the expected goodness of drawing a vial from urn B is greater than the expected goodness of drawing a vial from urn A, the standard teleological model suggests that the doctor ought to administer a vial from urn B.

Importantly, whilst the standard teleological model therefore suggests that the doctor ought to draw from the urn which is less secure, it nonetheless takes into account the security of each degree of goodness within its decision-criteria. If we hold that the fact-relative security of p is simply the chance that-p and that the expected goodness of an act is simply the chance-weighted sum of its possible futures, then the standard teleological model for decision-making regards the fact-relative security of each degree of goodness as a consideration in what an agent ought to do.

Security of goodness and utility

We might wonder, however, whether the security of each degree of goodness is *equally important* to our all-things-considered judgement regarding what an agent ought to do. The idea that agents ought to maximise expected goodness treats the security of each degree of goodness as equally weighty. It is not clear, however, that agents ought to weight the security of each degree of goodness equally.[292] In TOXIN, for instance, it is at least *prima facie* plausible that the doctor ought to place special weight on securing a minimal amount of goodness (i.e., the partial cure), even if that means performing an act which does not maximise expected goodness. Alternatively, we can imagine a happy circumstance where our choices provide a surfeit of good options and little risk of serious harm, where it is plausible that we might place less weight on securing a minimally decent future and place greater weight on securing futures which offer magnificent delights. The importance we accord to the security of each degree of goodness will reflect the degree to which we are risk-averse (or risk-seeking) with

[291] See Frank Jackson, "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection," *Ethics* (1991): p. 463; Parfit, *Reasons and Persons*, p. 25; Pettit, "Consequentialism," pp. 232-33.

[292] See, for instance, the discussion on the view that the expected goodness of the outcomes of an act determines the goodness of an act in Broome, *Weighing Goods: Equality, Uncertainty and Time*, p. 124.

respect to that degree of goodness. We can call this an agent's attitude towards *outcome risks* – the risk that some unfavourable outcome might occur.[293]

The standard way of accommodating attitudes towards outcome risks amongst decision theorists is to suggest that, rather than maximising expected goodness, agents ought to perform that act which maximises expected *utility*.[294] The formalisation of the expected utility of an act is much the same as the formalisation of the expected goodness of an act; substituting the goodness of a future for it utility, *u*.

$$EU(\phi) = \sum \mathrm{ch}(f|\phi).u(f)$$

On this, slightly modified account of the standard teleological model, an option φ is best if, and only if, its expected utility is at least as great as any other option φ available to the agent.

The term "utility" is sometimes used by philosophers and economists to indicate a measure of the goodness of particular futures for someone,[295] or alternatively as the strength of an agent's subjective preference for an option, but I have in mind a usage which is closest to that in decision theory.[296] For our purposes, we should understand the utility of a future as a measure which takes into account *both* (i) the goodness of the future, and (ii) the appropriate attitude towards outcome risks to that degree of goodness. Translated into the language of security, the utility of a possible future is the goodness of that future weighted by the importance of the security of that future. The weighting of the security of each unit of goodness can be straightforwardly represented by what is known as a *utility function*, which assigns to each degree of goodness a degree to which it contributes to the goodness of an act (see Figure 7).

We can imagine a large array of potential utility functions, representing a large array of ways of weighting the security of each degree of goodness in TOXIN. A *linear* utility function, for instance, represents a situation where the security of each degree of goodness is equally weighted (see function A in Figure 7). A *concave* utility function, on the other hand, represents a situation where less weight is placed on the security of each additional degrees of goodness (see function B in Figure 7). A *convex*

---

[293] The terminology of "outcome risk" comes from Nils-Eric Sahlin, "Unreliable Probabilities, Paradoxes, and Epistemic Risks," in *Handbook of Risk Theory: Epistemology, Decision Theory, Ethics, and Social Implications of Risk*, ed. Sabine Roeser, et al. (Dordrecht: Springer, 2012), p. 490.

[294] Karsten Klint Jensen, "A Philosophical Assessment of Decision Theory," in *Handbook of Risk Theory: Epistemology, Decision Theory, Ethics, and Social Implications of Risk*, ed. Sabine Roeser, et al. (Dordrecht: Springer, 2012), p. 423.

[295] See, for instance, Derek Parfit, "Equality and Priority," *Ratio* 10, no. 3 (1997): p. 205.

[296] See the discussion of this confusion in John Broome, ""Utility"," *Economics and Philosophy* 7, no. 01 (1991).

utility function will place greater importance on the security of each additional degree of goodness. We might also posit a range of mixed functions (such as function C in Figure 7), which are risk seeking with respect to initial degrees of goodness and risk averse with respect to degrees of goodness beyond some reference point.[297]

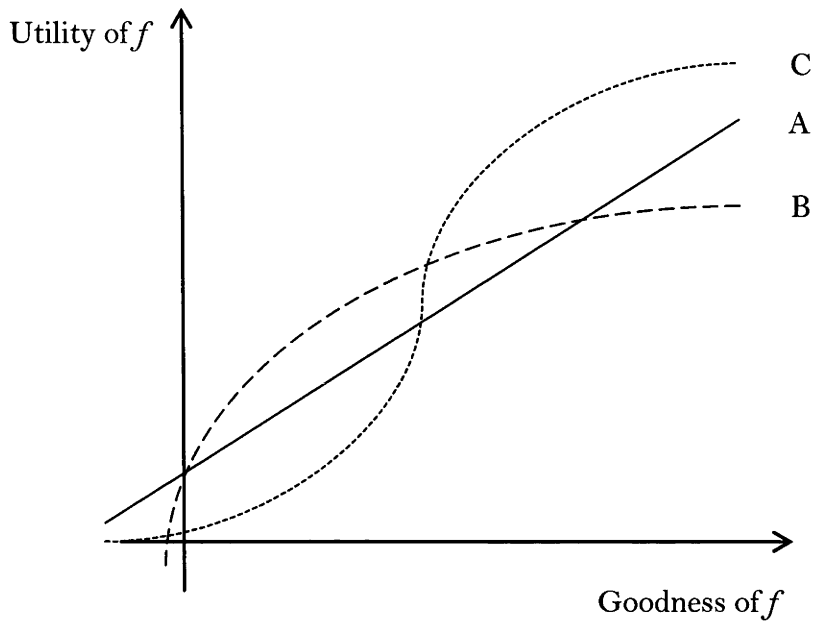


Figure 7: Possible utility-functions with respect to the goodness of possible futures. (A) Linear utility function representing risk-neutrality (B) Concave utility function representing standard aversion to risk (C) S-shaped utility function representing risk-seeking up to a threshold degree of goodness and risk-aversion beyond that threshold.

Thus, given a particular utility function, u, the utility of the various possible futures of TOXIN will be given by:

	Death (f_D)	Partial Recovery (f_P)	Full Recovery (f_R)
$u(f)$	$u(0)$	$u(15)$	$u(20)$

In this sense, given a heavy enough weighting towards the security of initial degrees of goodness, we can posit that the difference between the utility of a partial cure and the utility of a cure might be sufficiently narrowed such that the expected utility of administering a vial from urn A is greater than the expected utility of administering a vial from urn B. This would explain the intuition that administering a vial from B may not be the best act, despite the fact that it would maximise expected goodness.

[297] See Daniel Kahneman and Amos Tversky, "Prospect Theory: An Analysis of Decision under Risk," *Econometrica* 47, no. 2 (1979); A. Tversky and D. Kahneman, "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of risk and uncertainty* 5, no. 4 (1992).

One might object that this distinction between goodness and utility is artificial, since there is no way of *cardinally* measuring the goodness of futures independent of their utility.[298] Indeed, the most plausible ways of providing a cardinal measure of an agent's subjective preference for a future, the so-called representation theorems of Ramsey and Savage, compare an agent's subjective preferences for sets of uncertain futures such that they already incorporate the agent's attitudes towards risk.[299] Whilst I do not have the space to do complete justice to this objection, I will make a brief observation in response: which is that it seems to confuse the measurement of the utility of a future with the reasons for possessing that particular utility function.[300] The representation theorems are purely formal devices, designed in the first place to axiomatise the rationality of an agent's subjective preferences, and they have nothing to say about the reasons one might have for possessing a particular utility function. My account of the importance of security, on the other hand, identifies a property whose relative importance, coupled with the goodness of futures, provides a reason for adopting a particular utility function. If the only way to construct a cardinal measure of the goodness of a future is to mix the goodness of that future up with our attitudes towards risking the good, then so be it – but we should be mindful that this is precisely what we are doing.

Fact-relative security and outcome risks

Even if we believe that there is a valid distinction between the goodness and utility of possible futures, it does not follow that we ought to weight the security of each degree of goodness differently in our all-things considered judgements on what the doctor ought to do in TOXIN. Indeed, it might be the case that the most morally appropriate utility function is one which is linear, and therefore instantiates a neutral attitude towards outcome risks.

Amongst the class of all different utility functions, there will be some which are more morally appropriate than others. For instance, an aggressively risk-seeking utility function, applied to cases like TOXIN, will suggest that administering a vial from urn B maximises expected utility even in circumstances where the chance of

---

[298] Broome, *Weighing Goods: Equality, Uncertainty and Time*, p. 146; Jensen, "A Philosophical Assessment of Decision Theory," p. 413.

[299] See Frank P. Ramsey, "Truth and Probability," in *Philosophical Papers*, ed. Frank P. Ramsey and D.H. Mellor (Cambridge: Cambridge University Press, 1990 [1926]); J. von Neumann and O Morgenstern, *Theory of Games and Economic Behaviour* (Princeton, N.J.: Princeton University Press, 2004 [1944]).

[300] See Jensen, "A Philosophical Assessment of Decision Theory," p. 424; Philip Pettit, "Decision Theory and Folk Psychology," in *Foundations of Decision Theory: Issues and Advances*, ed. Michael Bacharach and Susan Hurley (Oxford: Blackwells, 1991), pp. 166-71.

failure was very high. Likewise, an ultra-conservative utility function would recommend administering a vial from A, regardless of whether urn B was filled almost entirely with vials of antidote X. Abiding by either of these kinds of utility function seems morally worse than abiding by milder utility functions which have attitude towards risk that accord with our intuitions. Given the rejection of these extreme attitudes towards outcome risk, one might think that the best attitude is the mildest of the utility functions – one which is simply a linear transformation of the goodness of each possible future.[301] This view, however, seems to become less attractive when we consider cases where there is a risk of catastrophe. Whilst it may be the case that the best attitude towards risk, *in this case*, is to adopt a linear utility function, it is by no means clear that this will always be the case. Consider a modified version of TOXIN, where the doctor must choose between administering a vial from urn A or urn B to one hundred patients all at once. Whilst the chance of each possible future occurring remains the same, the stakes are radically increased.[302]

	Death (f_D)	Partial Recovery (f_P)	Full Recovery (f_R)
$u(f)$	$u(0)$	$u(1500)$	$u(2000)$

Given a risk-neutral utility function, the expected utility of A will amount to 1500, whereas the expected utility of administering a vial from B is 1600. Given the dictum to maximise expected utility, it seems as though the doctor should administer a vial from urn B. Intuitively, however, the risk that all 100 patients might die seems to weigh very heavily in favour of administering a vial from urn A. The risk involved in drawing from urn B appears too great, since being sure that you save the lives of every patient seems more morally weighty than possessing a chance of curing every patient.

Ultimately, I cannot adjudicate on whether the most morally appropriate utility function reflects mild risk-aversion, risk-neutrality, or mild risk-seeking with respect to goodness. Whilst I am sympathetic to the view that it may be morally best to be mildly risk-averse with respect to the choice presented to the doctor in TOXIN, this is a position which reflects a substantive value commitment. Others may have the opposite intuition, that the doctor ought to be risk-neutral with respect to the good of her patient such that she administer compound B. Importantly, whatever one's

[301] Broome, *Weighing Goods: Equality, Uncertainty and Time*, p. 148.

[302] Let us assume that each patient's outcome contributes to the goodness of a future to precisely the same degree as in TOXIN. It is conceivable, however, that the marginal addition of equivalently good lives to a world does not increase the goodness of that world linearly. There may be diminishing marginal returns on the value of a life.

intuition regarding the best action, it should be clear that the attitude towards risk, and hence the weight accorded to the security of each unit of goodness, is a choice which reflects a particular value commitment.[303] Moreover, which utility function is most morally appropriate will not be dictated by what constitutes the *goodness* of a possible world. Given any account of the goodness of a possible future, that account will be compatible with the full range of ways of weighting the security of each degree of goodness. The utility of a given future will require consideration of the morally appropriate attitude towards outcome risks: a substantive consideration in determining the goodness of act which is not itself determined by the goodness of a world. In this sense, the fact-relative security of goodness might not merely be a consideration in moral decision-making, but one which requires agents to reflect upon values which are independent of the goodness of possible futures.

So far we have presented an account of right action which, with some modifications, accords with the standard teleological account. Almost all should accept that the fact-relative security of the good is an important consideration in determining what agents ought to do relative to the facts. Moreover, the prospect that we ought to weight the security of each unit of goodness differentially can be accounted for by the dictum to maximise expected utility. In this sense, we seem to have reason to admit that there are substantive value considerations which do not reduce down to the goodness of the possible futures which acts may bring about. Nonetheless, in the sense that decision-theoretic accounts of right action are already cognizant of the distinction between utility and goodness, this modification is broadly in line with the standard teleological model.

In the next section I do away with the assumption that agents know the relevant chance information, and ask what agents ought to do relative to their evidence. In so doing, I will forward a potentially controversial view on the role of evidence-relative security as a consideration in a pluralistic account of moral decision-making. This view, which suggests that evidence-relative security plays an important role in decision-making where the evidence is compatible with multiple credence functions,

[303] The problem of how to weigh the security of each unit of value becomes even more acute if we think that some forms of final value are incommensurable with one another, such that they cannot be straightforwardly conglomerated into cardinal rankings of the goodness of possible futures. For instance, we might think that the satisfaction of preferences and the autonomous formation of preferences are both finally valuable, yet have no clear way of determining their overall contribution to a cardinal ordering of possible futures. If this is the case, our situation might be further complicated by the fact that it may be appropriate to be risk neutral with respect to the satisfaction of an agent's preferences, but very risk averse with respect to an agent's autonomy.

potentially illuminates why the security of the good appears to be *especially* valuable in contexts where we have limited evidence.

6.4 Evidence-relative Security and Evidence-relative Oughts

Suppose that we modify TOXIN such that the available evidence regarding the composition of one of the urns is vague. How does this affect the role of security in deliberations on what ought to be done?[304]

> **TOXIN II:** Just like before, a patient arrives in the emergency room suffering from a potent neurotoxin. This time, however, the only available antidotes are X (which will cure the patient) and Z (which will have no effect, such that the patient dies). Unfortunately, when the doctor visits the dispensary she finds that the vials of antidote are not properly separated, and are instead stored in two urns, A and B. Urn A is labelled "1 X: 1 Z" whilst Urn B is simply labelled "X and Z". Given the dispensary's idiosyncratic labelling practices, the doctor can be certain that Urn A contains precisely a ratio of one vial of antidote X to every one vial of antidote Z. Likewise, whilst she can be certain that Urn B contains only vials of antidote X and/or antidote Z, she has no reason to suppose they are in any particular ratio. The doctor must choose from which urn to draw a vial.

In TOXIN two, we can assume that there are only two relevant outcomes: either the patient dies or the patient is cured. In order to illustrate the importance of evidence-relative security, I will assume both that the goodness of the futures in TOXIN II is identical to the goodness of the futures in TOXIN and that the most morally appropriate utility function is risk-neutral.

	Death (f_D)	Recovery (f_R)
$u(f)$	0	20

Given the limited information available to the doctor, there are several distributions of credence compatible with her evidence. In particular there are a range of ways in which the doctor would be justified in assigning her credence to each future conditional on administering a vial from either A or B.

	$cr(f_D\|A)$	$cr(f_R\|A)$	$cr(f_D\|B)$	$cr(f_R\|B)$
c_{lower}	0.5	0.5	1	0
		...		
c_{mid}	0.5	0.5	0.5	0.5
		...		
c_{upper}	0.5	0.5	0	1

[304] This is simply a moralized version of a case provided by Daniel Ellsberg in his landmark study of ambiguity aversion. See Daniel Ellsberg, "Risk, Ambiguity, and the Savage Axioms," *The Quarterly Journal of Economics* 75, no. 4 (1961).

The doctor's choice, therefore, is between drawing from an urn which contains an even mixture of X and Z or drawing from an urn which contains an unknown mixture of X and Z. In this sense, she is faced with options which provide different degrees of evidence-relative security.

Evidence-relative security and expected goodness

One might begin by having an erroneous thought: that we should seek to apply evidence-relative security to TOXIN II in the same way that we applied fact-relative security to TOXIN. Recall that the fact-relative expected goodness of an act φ can be characterised as the sum of the security-weighted goodness of the possible futures, f, conditional on φ. In this sense, we might think that the evidence-relative expected goodness of an act φ can be characterised in just the same way. This, however, is a mistake. In order for an expectation (of goodness, utility or any other value) to be coherent the weightings of each state of affairs must obey the probability axioms. The reason the fact-relative security of p could form part of an expectation of the goodness of an act is that it does obey the probability axioms; indeed, it is simply the objective chance that-p. Recall, however, the account of evidence relative security which I proposed in Chapter Five:

> **ERS-3:** The evidence-relative security of a proposition p is x just in case were any of the credence functions justified by the available evidence equivalent to the objective chance function, then the fact-relative security of p would be <u>at least</u> x.

On this account, the evidence-relative security of f is not a probability. Indeed, as we showed in Chapter Five, the evidence relative security of p and $\sim p$ will collectively violate the complementarily axiom of probability theory. Whilst the evidence-relative security of p takes the same *value* as the probability that-p in a particular model of the future which is compatible with the evidence, it is not itself a probability. Were an agent to attempt to calculate the expected goodness of an act from the evidence-relative security of a state of affairs, that agent would potentially be open to accepting a so-called Dutch Book (a sequence of gambles which collectively lead to a sure loss).[305]

Expected utility and evidential models

A more coherent approach is to calculate the expected utility of an act using each credence distribution which is compatible with the evidence. Recall that the evidence-relative security of p is determined by the credence in p that is robust across the

[305] See Chapter Five, p. 144.

various credence functions, c, which are members of the set, C, of all the models compatible with the available evidence. Each credence function acts as a model of the structure of the future of the world, where each possible future, f, has a particular probability of being the actual future. We can therefore think of a credence function which assigns to a proposition p a credence of x, as a kind of evidential model of a set of circumstances where the fact-relative security of p is x. Within such evidential models, we are therefore able to calculate the expected goodness of each act.

Likewise, we may identify the most morally appropriate utility function, u, to apply to the circumstances modelled in that representation.[306] In this sense, for each representation of the future we can also calculate the expected utility of each act φ. This will not be the actual expected utility of an act, but the expected utility of an act were it the case that the evidential model of the future was equivalent to the actual structure of the future. In this sense, the expected utility of an act φ from within one of these evidential representations, e, is the sum of the probability-weighted value of those possible futures, f, conditional on φ. This can be formalised as follows:

$$EU_c(\phi) = \sum \mathrm{cr}_c(f|\phi).u_c(f)$$

Let us turn now to the decision facing the doctor in TOXIN II. In order to make things simple, lets us suppose that the morally appropriate attitude towards outcome risks is to be risk-neutral, such that that the utility function remains constant throughout all of the evidential models. Because all of the credence functions compatible with the evidence assign the same credence to the possible futures conditional on drawing a vial from urn A, the expected utility of administering a vial from urn A within any evidential model of the future will be:

$$\begin{aligned} \mathrm{EU}_c(A) &= \mathrm{cr}_c(f_D|\phi).u_c(f_D) + \mathrm{cr}_c(f_R|\phi).u_c(f_R) \\ &= (0.5 \times 0) + (0.5 \times 20) \\ &= 10 \end{aligned}$$

On the other hand, since there multiple credence functions compatible with the evidence about urn B, such that the evidence-relative expected utility calculations for compound B are more complex. In particular:

---

[306] It is likely to be the case that an agent's attitudes towards outcome risks, and hence their utility function, ought to be invariant across the different models of the future. Some behavioural economists note, however, that it appears that agent's actual attitudes towards outcome risk are relative to the decision circumstances. See Amos Tversky and Daniel Kahneman, "The Framing of Decisions and the Psychology of Choice," *Science* 211, no. 4481 (1981).

$$\begin{aligned}
\mathrm{EU}_{c(\text{lower})}(B) &= \mathrm{cr}_c(f_D|\phi).u_c(f_D) + \mathrm{cr}_c(f_R|\phi).u_c(f_R) \\
&= (1 \times 0) + (0 \times 20) \\
&= 0
\end{aligned}$$

$$\begin{aligned}
\mathrm{EU}_{c(\text{mid})}(B) &= \mathrm{cr}_c(f_D|\phi).u_c(f_D) + \mathrm{cr}_c(f_R|\phi).u_c(f_R) \\
&= (0.5 \times 0) + (0.5 \times 20) \\
&= 10
\end{aligned}$$

$$\begin{aligned}
\mathrm{EU}_{c(\text{upper})}(B) &= \mathrm{cr}_c(f_D|\phi).u_c(f_D) + \mathrm{cr}_c(f_R|\phi).u_c(f_R) \\
&= (0 \times 0) + (1 \times 20) \\
&= 20
\end{aligned}$$

Thus given the range of credence functions compatible with the doctor's evidence assigned to futures conditional on B, the expected utility of administering a vial from urn B, according to the full set of models within *C*, is best represented by the interval $0 \leq EU(B) \leq 20$.

Given that the evidence does not uniquely determine the expected utility of drawing a vial from urn B, it seems that the dictum to maximise expected utility does not determine what the doctor ought to do. On some distributions of credence the dictum will recommend choosing urn A, on others it will recommend choosing urn B and on (at least) one it will recommend that we should be indifferent between choosing compound A or B. It therefore appears that, given the evidence available to the doctor, there is no fact of the matter as to whether compound A or compound B yields the optimal expectation.[307]

That in instances of limited evidence there appears to be no fact of matter regarding the act which maximises expected utility might motivate two different revisions to the standard maximisation principle. First, one might suggest that an agent ought to perform an act only if it maximises expected utility across *all* the credence functions compatible with the evidence. Second, some authors, notably I.J. Good,[308] suggest that an agent ought to perform *any* act which maximises expected utility according to some credence function within the set compatible with the evidence.[309] Of these two rules, the most plausible is Good's rule. Consider that in

---

[307] Note that this is not the claim that the doctor ought to be indifferent between either act, rather it is the much more radical claim that there is no fact of the matter about the goodness of the available acts.

[308] I. J. Good, "Rational Decisions," *Journal of the Royal Statistical Society. Series B (Methodological)* 14, no. 1 (1952): p. 114.

[309] Hájek and Smithson discuss a different version, where an agent is *required* to φ only if φ maximises expected value according to all probability functions compatible with their evidence. I treat this as equivalent to Good's rule for our purposes. See Hájek and Smithson, "Rationality and Indeterminate Probabilities," p. 44.

situations like TOXIN II, where there is no option which maximises expected utility across all the credence functions compatible with the evidence open to you, then the first rule would posit that whatever you do is morally impermissible! Since we are frequently in epistemic situations similar to TOXIN II, where many different acts maximise expected utility given the available evidence, the first rule would suggest that we frequently act impermissibly. Good's Rule, on the other hand, says that when our evidence is incomplete we need only exclude those options which determinately do not maximise expected utility – all others are permissible.

Whilst Good's rule provides a determinate analysis of what the doctor ought to do in TOXIN II, it does seem to miss something morally important. It is plausible, for instance, that the *best* thing to do might be to administer a vial from urn A, since the available evidence about that urn is clear. Indeed, when individuals are asked to make a choice between gambles whose odds are based on precise evidence and equivalent gambles whose odds are based on vague or imprecise evidence, they almost always choose the gamble with the certain odds.[310] Administering a vial from compound B is akin to accepting a gamble which is based on vague evidence, and in this sense seems to require taking an *epistemic* risk. Intuitively, at least, avoiding the imposition of an epistemic risk, on yourself or others, seems morally valuable. It is for this reason that I suspect that administering a vial from urn A might be better than administering a vial from urn B.

Evidence-relative security and evidential risks

It is at this point that evidence-relative security might make its distinctive contribution to moral decision-making. Indeed, it has been suggested that in instances where two acts maximise expected utility according to at least one representation of the future compatible with your evidence, that a secondary decision-criteria, what is sometimes called a "security level," should be used to determine which act is best.[311] The evidence-relative security of two different propositions might play just such a role as a secondary decision-criteria.

First, agents might seek to maximise the amount of *goodness* that is evidence-relative secure.[312] For each act, φ, there will be some degree of goodness which will be totally evidence-relative secure – i.e., were *any* of the credence functions within C the

[310] See Ellsberg, "Risk, Ambiguity, and the Savage Axioms,"and recent empirical confirmation by Yoram Halevy, "Ellsberg Revisited: An Experimental Study," *Econometrica* 75, no. 2 (2007).

[311] See Isaac Levi, *Hard Choices: Decision Making under Unresolved Conflict* (Cambridge: Cambridge University Press, 1990), pp. 122-47.

[312] This is the approach favoured by Levi, see Levi, "Ignorance, Probability and Rational Choice,"

actual chance function, the chance that the actual future will include at least that degree of goodness would be one. This will correspond to the goodness of the worst possible future where φ is realised. Given a set of options, this criterion, known amongst decision-theorists as "minimax", suggests that agents perform that act whose worst outcome is at least as valuable as any other acts worst outcome. Minimax is intuitively appealing in situations, like the Rawlsian original position, where an agent is completely ignorant of the distribution of probabilities amongst alternatives.[313] In TOXIN II, however, the degree of goodness that is totally secure in either case would not differentiate between the two urns, since the worst outcome in either situation is the same: death. In this sense, a concern for the evidence-relative security of final value does not seem to explain the intuition that urn A is preferable to urn B.

Second, agents might seek to maximise the amount of *expected utility* that is evidence-relative secure. For each act, φ, there will be some level of expected utility which will be totally evidence-relative secure – i.e., were *any* of the credence functions within *C* the actual chance function, the chance that the future would be one were the expected utility of performing φ would be at least that great would be one.[314] For simplicity, I shall call this the secure expected utility (SEU) of an act. Given a set, *C*, of credence functions, *c*, compatible with the evidence, the secure expected utility of an act, φ can be formalised as:

$$\mathrm{SEU}(\phi) = \min_{c \in C} \sum \mathrm{cr}_c(f|\phi).u_c(f)$$

In contexts where the evidence is compatible with multiple ways of distributing ones credence, considering the SEU of an act amounts to considering the degree of expected utility of that act which an agent can rely upon given her evidence.[315] Recalling that each credence function acts as a model of the future, the SEU of an act can be thought of as the minimum amount of utility which an agent could expect, regardless of which model of the future is closest to the way in which the future is actually structured. Whilst it may be the case that the future is structured in more favourable ways,

---

[313] See Rawls, *A Theory of Justice*, pp. 154-56 and the famous critique of his reasoning from John C. Harsanyi, "Can the Maximin Principle Serve as a Basis for Morality? A Critique of John Rawls's Theory," *The American Political Science Review* 69, no. 2 (1975).

[314] It may seem slightly baroque to insist that the expected utility of an act can be *secure* in the sense of having some probability of being true in the future, but if we consider that from within each representation, *e*, the expected utility of φ-ing at *t* will remain the same at times after *t* (even if φ is indeed the act which is performed), then it is perfectly correct to say this proposition can be evidence-relative secure to a certain degree.

[315] See Chapter Five.

relative to the available evidence an agent can be sure that the expected utility of an act is no less than the SEU of that act.

Importantly, the SEU of an act can be contrasted with the Maximum Expected Utility (MEU) of an act. Recall that the evidence-relative security of a proposition, p, identifies the minimum degree of credence in p which any credence function compatible with the evidence would be willing to assign to p. This can be contrasted with the *optimistic* probability of p: the maximum credence in p which any credence function compatible with the evidence would be willing to assign to p. Just as there is a level of secure expected utility for each act, there will be some level of maximum expected utility – the maximum degree of expected utility that would be assigned to an act by some credence function compatible with the evidence.[316]

In TOXIN II, the amount of secure expected utility differentiates between the two urns: since the secure expected utility of administering a vial from urn A is 10, whilst the secure expected utility of administering a vial from urn B is precisely zero. The importance of secure expected utility might therefore explain the intuition that administering a vial from urn A is preferable to administering a vial from urn B. In what follows I discuss the secure expected utility of an act as a consideration in determining the goodness of that act.

Maximising secure expected utility

Given a set of options, agents who are maximally averse to epistemic risk would seek to maximise secure expected utility (MSEU). According to MSEU, an agent ought to perform the act whose minimum expected utility is at least as great as any other acts minimum expected utility.[317] Indeed, one might be tempted to think that MSEU should serve as the *sole* consideration in what an agent ought do relative to their evidence, such that one might permissibly perform only those acts whose secure expected utility is at least as great as any other option. Certainly, variants of MSEU have been advocated by decision-theorists seeking to provide a criterion of decision which is generally applicable.[318] Moreover, MSEU is equivalent to maximising expected value in situations where only one credence function is compatible with the evidence, such that it comports with the standard decision-theory in situations where the evidence is good.

---

[316] We can formalise this as: $\mathrm{MEU}(\phi) = \max_{c \in C} \sum \mathrm{cr}_c(f|\phi).u_c(f)$.

[317] This is known to decision theorists as the maxi-min expected utility rule. See Gärdenfors and Sahlin, "Unreliable Probabilities, Risk Taking, and Decision Making,".

[318] See ibid; L Hurwicz, "Some Specification Problems and Applications to Econometric Models," *Econometrica* 19(1951); A Wald, *Statistical Decision Functions* (New York: Wiley & Sons, 1950).

Unfortunately, MSEU leads to some intuitively implausible restrictions on what one might *permissibly* do. Suppose that evidence suggests that there are 100 vials in urn B, at least 49 of which contain compound X. Given this evidence, the doctor's credence that the patient will fully recover ought to be within the interval $0.49 \leq \mathrm{cr}(f_R) \leq 1$. MSEU suggests that, given this modified evidence about urn B, it still be *impermissible* to administer a vial from urn B, since the secure expected utility of doing so would be less than the secure expected utility of choosing a vial from urn A.[319] It seems intuitively implausible, however, that it is morally impermissible to choose urn B in such a situations, since whilst SEU(B) is only slightly less than SEU(A), the maximum expected utility of administering a vial from urn B is much greater than that offered by administering a vial from urn A. In this sense, MSEU seems to imply a maximally conservative strategy, one which becomes implausible once we consider situations where the evidence, whilst compatible with a slightly more pessimistic credence function, is compatible with many more optimistic distributions of credence.

More problematically, MSEU leads to highly counterintuitive proscriptions in some complex decision situations. Considering the following set of interlocking games:[320]

> **Stone Game:** Suppose I am holding a stone in either my left or my right pocket, but that you are completely ignorant as to which pocket the stone rests in. If you guess correctly, then you win \$100, if you guess wrongly I will spot you \$10.

The SEU of guessing either "Right" or "Left" is therefore \$10, such that the overall SEU of playing the Stone Game is \$10.[321] Suppose further that I offer you a choice between two gambles which incorporate Stone Game:

> **Gamble 1:** A fair coin is tossed. If it comes up heads you can either play the Stone Game with a bet on "Left" or take \$9. If it comes up tails you can either bet on "Right" or take \$9. If you choose to play the Stone Game with the choices determined by the coin, then the payoffs for a correct or incorrect 'guess' remain \$100 and \$10 respectively.
>
> **Gamble 2:** A fair coin is tossed. If it comes up heads you can either play the Stone Game with a bet on "Left" or take \$11. If it comes up tails you can either bet on "Right" or take \$11.

---

[319] In particular, $\mathrm{SEU}(B) = (0.49 \times 20) + (0.51 \times 0) = 9.8$, whilst $\mathrm{SEU}(A) = (0.5 \times 20) + (0.5 \times 0) = 10$.

[320] The following is Sahlin's own explanation of a problem Ted Seidenfeld raised in response to Gärdenfors and Sahlin's version of MSEU. Sahlin, "Unreliable Probabilities, Paradoxes, and Epistemic Risks," p. 488; Teddy Seidenfeld, "Decision Theory without "Independence" or without "Ordering"," *Economics and Philosophy* 4, no. 02 (1988).

[321] The SEU of a choosing "Right" in the Stone Game is given by $\mathrm{cr}_{\min}(R) \times 100 + \mathrm{cr}_{\max}(L) \times 10$. Given the credence function which assigns $\mathrm{cr}(R) = 0$ and $\mathrm{cr}(L) = 1$, the SEU of "Right" is \$10 (and likewise for "Left"). The SEU of playing the Stone Game is therefore \$10.

If you opted to play Gamble 1, using MSEU as your only decision criteria, you ought to always opt to play the Stone Game.[322] In this respect, Gamble 1's overall SEU is straightforwardly computable – it is \$27.50.[323] If you opted to play Gamble 2, however, following MSEU requires that you ought to always opt to take the \$11, such that the overall SEU of Gamble 2 is only \$11. Given a choice between Gamble 1 and Gamble 2, MSEU suggests that you ought to take Gamble 1.

This, however, seems like a mistake. Gamble 2 looks like it straightforwardly dominates Gamble 1, since nothing about the odds, the payoffs or the SEU of the Stone Game changes, and in Gamble 2 you possess an alternative to playing the Stone Game which is more attractive (i.e. \$11 as opposed to \$9)! The problem, it seems, is your slavish application of MSEU across a sequential series of bets. In this sense, it appears unreasonable to employ MSEU as the *sole* consideration in decision-making.

Secure expected utility and maximum expected utility

Whilst maximising secure expected utility should not, therefore, be the sole consideration in moral decision-making, we might nonetheless consider the secure expected utility of an act along with other criteria. Recall that just as there is a level of secure expected utility for each act, there will be some level of maximum expected utility (MEU) – the maximum degree of expected utility that would be assigned to an act by some model of the future compatible with the evidence.[324] Acting on the MEU of an act involves taking a risk – an epistemic risk – that more precise evidence about the future would not lead us to reject those models of the future which lead to the maximal expectation. Acting on the secure expected utility of an act, on the other hand, does not carry such a risk (even if it leads to practically irrational acts). In this sense, the degree to which an agent places importance on avoiding an epistemic risk can be modelled by considering the weight one might accord the secure expected utility of an act as opposed to the maximum expected utility of an act.[325]

---

[322] Since the SEU of playing the Stone Game is \$10, whereas the SEU of taking \$9 is exactly that.

[323] The SEU of Gamble 1 given that you ought to play the Stone Game is given by:

$$\begin{aligned}\text{SEU(Gamble 1)} &= \big(\text{cr}(H) \times \text{cr}_{\text{mid}}(L) \times u(L)\big) + \big(\text{cr}(T) \times \text{cr}_{\text{mid}}(R) \times u(R)\big) \\ &= (0.5 \times 0.5 \times 100) + (0.5 \times 0.5 \times 10) \\ &= 27.5\end{aligned}$$

[324] See p. 149.

[325] One might formalise this as an all-thing considered decision-consideration: "Security-Weighted Expected Utility" Where s is a normalised weighting which reflects the importance one places on the amount of expected utility which is secure, as opposed to the amount of expected utility which is evidentially possible.

$$\text{SWEU}(\phi) = s\left(\min_{c \in C} \sum \text{cr}_c(f|\phi).u_c(f)\right) + (1-s).\left(\max_{c \in C} \sum \text{cr}_c(f|\phi).u_c(f)\right)$$

Consider TOXIN II once more. Given the stakes at play and the evidence available, it seems that the doctor ought to weight the SEU of the available options more heavily than the MEU of each option in determining the goodness of each act, such that it would be best to administer a vial from compound A. However, if one lowers the stakes, such that little harm will occur if the wrong medication is administered, then taking a vial from the more epistemically risky urn might be the best option. Likewise, in circumstances where a patient has "run out of options", such as might be the case with a previously incurable illness, the best option might be to take an epistemic risk on a drug about which there is limited evidence.

The exact weight that one might place on either the secure or maximum expected utility of an act in an agent's all-things-considered judgements, will be dependent on the overall attitude towards epistemic risk which an agent adopts.[326] Like the moral appropriateness of attitudes towards outcome risk, this is not a discussion which I wish to settle here. Some attitudes towards epistemic risk, such as a maximally risk-seeking approach, seem obviously impermissible in cases where the stakes are high and the present is tolerable. Likewise, a maximally risk-averse strategy seems permissible in some cases – e.g. where the stakes are very high and there is no relevant evidence – but cannot be generally recommend. There may be no definitive answer as to which attitude towards epistemic risk is best.

Importantly, whatever the most morally appropriate attitude towards epistemic risk, it should be clear that this attitude is not dictated by what constitutes the *goodness* of a possible world. Given any account of the goodness of a possible future, that account will be compatible with the full range of ways of weighting the secure expected utility of an act against the maximum expected utility of that act. In this sense, the secure expected utility of an act might not merely be a consideration in determining the goodness of an act, but one which requires agents to reflect upon values which are not reducible to the goodness of worlds.

Conclusion

In this chapter I have argued that both fact-relative and evidence-relative security appear to be important considerations in determining our fact-relative and evidence-

---

SWEU models an agent's attitudes towards epistemic risk, such that if $s = 1$ then the agent is maximally averse to epistemic risks, and if $s = 0$ then the agent is maximally seeking towards epistemic risks, and if $s = 0.5$ then the agent is neutral towards epistemic risks.

[326] This accords well with Isaac Levi's approach, which posits that the choice of secondary decision-criterion is a value consideration. Levi, *Hard Choices: Decision Making under Unresolved Conflict*, p. 127.

relative moral obligations. I began by exploring the possibility that the security of finally valuable goods is itself constitutive of the goodness of a possible world. Whilst some may have the intuition that security does contribute to goodness in such a way, the potential explanations for such a contribution implied a number of controversial axiological commitments.

In considering the fact-relative goodness of acts, the fact-relative security of goodness plays an important role in distinguishing what we ought to do given indeterminacy over the outcome of our actions. Moreover, the weight which ought to be accorded to the security of each degree of goodness models a substantive value consideration: the attitude towards risk which should be adopted in assessing the goodness of acts. Likewise, in considering our evidence-relative obligations, the importance of the secure expected value of an act seemed to explain a common intuition regarding the preferability of evidentially precise, rather than evidentially vague, situations. In situations where there are multiple credence functions compatible with the evidence, placing special weight on the secure expected utility of acts in determining their all-things-considered goodness, models the importance of avoiding an epistemic risk. Whilst I did not give a conclusive argument for aversion with respect to either outcome risks or epistemic risks, I suggested that it may be morally preferable to give some weight to security, such that our actions ensure the achievement of a minimally decent future.

Conclusion

At the beginning of this thesis I suggested that security was like justice: a powerful but opaque term. As the argument in this thesis has progressed I have sought to clarify the meaning of security and identify its distinctive contribution to moral decision-making. By identifying its role in helping agents to navigate the future, even when the available evidence is limited, I hope to have shed light on why security is an important moral and political concept. In particular, there are four major conclusions which I hope this thesis enables us to draw.

First, whilst the meaning of 'security' is complex, it is by no means un-analysable. Scholars in the disciplines of Security Studies, Criminology and Political Theory, whilst they have sometimes engaged in conceptual analysis of security, have tended to view it as an essentially contested concept.[327] As this thesis shows, however, there are a number of important, and hitherto unrecognised, distinctions which bear on the meaning of 'security'. In particular, 'security' is a polyseme which can refer to the distinct concepts of security practices ($\text{Security}_{\text{P}}$), the state of being secure ($\text{Security}_{\text{B}}$) and the notion of security as a quality of reliability ($\text{Security}_{\text{R}}$). Likewise, the common distinction between subjective and objective security obscures a more complex set of distinctions between the kinds of goods (i.e., concrete or affective) which can constitute conceptions of being secure, and the different perspectives (i.e., fact-relative, evidence-relative and belief-relative) from which the degree of security be analysed. Identifying these distinctions allows philosophers and citizens to demand that the invocation of security is accompanied by an explanation of exactly what set of factual and evaluative claims are being made. In this sense, the conceptual analysis of security is not only possible, but helps to clarify the kinds of claims which agents are making when they invoke security.

Second, I argued that the concept of being secure has a particular structure and that at the core of this structure is the concept of security as a quality of reliability. In

[327] Recall the discussion in Chapter Two, p. 43.

particular, I argued that to be secure_B as a type of entity, a set of relevant propositions, $(p_1, p_2, \ldots p_n)$ must be secure_R to a sufficient degree. This account of the structure of the concept of Security_B identified Security_R as an uncontested core, around which different conceptions of the good substitute different conditions for an entity's security. In this sense, I have provided a structured way of identifying exactly which parts of the concept of being secure are contested. Whilst the concepts of human security and national security are likely to remain essentially contested, this can now be understood as a contest over the importance of the secure enjoyment of particular goods for particular referents, rather than as a conceptual contest over the meaning of 'security' itself.

Third, the secure enjoyment of finally valuable goods seems to be an important consideration in moral decision-making. Where the available evidence is limited, there may be multiple distributions of credence which are compatible with the evidence. In such situations, placing special weight on the secure expected utility of acts in determining their all-things-considered goodness, appears to model the importance of avoiding epistemic risks. The precise weight which ought to be accorded to the secure expected utility of an act, as opposed to the maximum expected utility of an act will model the morally appropriate attitude one ought to adopt towards epistemic risks.

Finally, given these observations about the role which secure enjoyment plays in deliberations on the goodness of acts, it seems as though the standard teleological model, whereby the goodness of an act is based solely on the goodness of the worlds that it may bring about, should be modified to account for the value of securing oneself from both outcome risks and epistemic risks. Whilst there are many sophisticated accounts of the contribution of finally valuable goods – such as equality, liberty, wellbeing, promise-keeping, beauty and knowledge – to the goodness of worlds, the standard model holds that the goodness of acts is determined simply by the goodness of worlds and the use of a value-neutral decision rule. The idea that the goodness of an act might be determined by a pluralistic, value-laden decision procedure is a potential reason to reject the standard teleological model.

Future Directions

Whilst this thesis takes significant steps towards clarifying the concept of security, the value of secure enjoyment and its role in moral decision-making, there is a wealth of further work which ought to be performed.

First, this thesis is by no means a complete account of the role that secure enjoyment plays in moral decision-making. Further work must be done to identify the

terms upon which the secure enjoyment of a good ought to be balanced against maximising the expectation of a good. Identifying the circumstances in which the evidence-relative security of a good ought to be accorded priority, and where its importance ought to be diminished in favour of maximising the expectation of a good, are critical questions for decision theorists and moral philosophers who aspire to provide practical guidance.

Second, the appropriate connection between evidence-relative security and the affect of security, seems to be an important site for further inquiry. It is tempting to think that moral agents would be better off if they were fully informed of the risks and dangers they faced, so that they may take appropriate and reasonable measures to protect themselves. On the other hand, it seems reasonable to prefer a degree of quietude at the price of a little ignorance, especially if one can delegate the promotion of security to a competent political authority. The appropriate balance between a rational awareness of the risks and the minimization of anxiety, and the extent to which those aware of the risks ought to make them known, seem to me to be fascinating and important questions.

Third, those philosophers keen to establish the boundaries of human rights and the obligations of political institutions (particularly at the global level), ought to contemplate the appropriate content of conceptions of being secure. A pluralistic notion of human security – which identified a set of core goods which each individual ought to enjoy securely – might, for instance, serve to limit the "duties of assistance" which the global rich owe to the global poor. On the other hand, attention by political philosophers to the notion of state security, so important to debates in international relations, may serve to illuminate debates on humanitarian intervention and the role of global political institutions.

Finally, we might bring the concept of secure enjoyment (and decision-making under limited evidence) to bear on a range of key practical questions. The global proliferation of nuclear weapons, the emergence of new pandemic diseases and the safety of emerging technologies are all issues which pose very grave dangers but about which our evidence is poor. In another vein, the feasibility of programs of social justice and schemes for the elimination of poverty and disease are likewise afflicted by the blight of inadequate evidence. Faced with problems of such a kind, the prudent course of action might be to seek to increase the security of those goods which are vitally important, even if that means delaying or abandoning our greatest expectations.

Bibliography

Alkire, Sabina. *A Conceptual Framework for Human Security: Working Paper 2*, Crise Working Paper. Oxford: Centre for Research on Inequality, Human Security and Ethnicity (CRISE), 2003.

Arends, J. Frederik M. "From Homer to Hobbes and Beyond - Aspects of 'Security' in the European Tradition." In *Globalization and Environmental Challenges: Reconceptualizing Security in the 21st Century*, edited by Hans Günter Brauch, Úrsula Oswald Spring, Czeslaw Mesjasz, John Grin, Pál Dunay, Navnita Chadha Behera, Béchir Chourou, Patricia Kameri-Mbote and P. H. Liotta. New York: Springer, 2008.

Austin, J. L. *How to Do Things with Words.* Edited by J. O. Urmson and Marina Sbisà. 2nd ed. Oxford: Oxford University Press, 1980.

Bajpai, Kanti. "Human Security: Concept and Measurement." In *Kroc Institute Occasional Paper.* Notre Dame: Joan B. Kroc Institute for International Peace Studies, 2000.

Baldwin, David A. "The Concept of Security." *Review of International Studies* 23, no. 1 (1997): 5-26.

Barnes, Elizabeth, and Ross Cameron. "The Open Future: Bivalence, Determinism and Ontology." *Philosophical Studies* 146, no. 2 (2009): 291 - 309.

Basic Law for the Federal Republic of Germany.

Bentham, Jeremy. "Principles of the Civil Code." edited by John Bowring, 1843.

Bigelow, John C. "Possible Worlds Foundations for Probability." *Journal of Philosophical Logic* 5, no. 3 (1976): 299–320.

Bitar, D., A. Goubar, and J.C. Desenclos. "International Travels and Fever Screening During Epidemics: A Literature Review on the Effectiveness and Potential Use of Non-Contact Infrared Thermometers." *Eurosurveillance* 14, no. 6 (2009): 1-5.

Booth, Ken. "Security and Emancipation." *Review of International Studies* 17, no. 4 (1991): 313-26.

Booth, Ken. *Theory of World Security.* New York: Cambridge University Press, 2007.

Broome, John. ""Utility"." *Economics and Philosophy* 7, no. 01 (1991): 1-12.

Broome, John. *Weighing Goods: Equality, Uncertainty and Time.* Oxford: Blackwell, 1991.

Bull, Hedley. *The Anarchical Society: A Study of Order in World Politics.* London: Macmillan, 1977.

Burke, Anthony. *Beyond Security, Ethics and Violence: War against the Other.* New York: Routledge, 2007.

Buzan, Barry. *People, States and Fear: The National Security Problem in International Relations.* Sussex: Wheatsheaf Books, 1983.

Buzan, Barry, and Lene Hansen. *The Evolution of International Security Studies.* Cambridge: Cambridge University Press, 2009.

Buzan, Barry, and Ole Wæver. *Regions and Powers: The Structure of International Security.* Cambridge: Cambridge University Press, 2003.

Buzan, Barry, Ole Wæver, and Jaap De Wilde. *Security: A New Framework for Analysis.* London: Lynne Rienner Publishers, 1998.

Calcott, Brett. "Wimsatt and the Robustness Family: Review of Wimsatt's Re-Engineering Philosophy for Limited Beings." *Biology and Philosophy* 26, no. 2 (2011): 281-93.

Carlin, Paul S. "Violations of the Reduction and Independence Axioms in Allais-Type and Common-Ratio Effect Experiments." *Journal of Economic Behavior & Organization* 19, no. 2 (1992): 213-35.

Clarke, Barry. "Eccentrically Contested Concepts." *British Journal of Political Science* 9, no. 1 (1979): 122-26.

Coffman, E. "Thinking About Luck." *Synthese* 158, no. 3 (2007): 385-98.

Colyvan, Mark, Damian Cox, and Katie Steele. "Modelling the Moral Dimension of Decisions." *Nous* 44, no. 3 (2010): 503-29.

Constitution of Ireland.

Déclaration Des Droits De L'homme Et Du Citoyen De 1789. Assemblée nationale, 1790.

Deudney, D. "The Case against Linking Environmental Degradation and National Security." *Millenium: journal of international studies* 19, no. 3 (1990): 461-76.

Dillon, Michael. *Politics of Security: Towards a Political Philosophy of Continental Thought.* New York: Routledge, 1996.

Earman, J. *A Primer on Determinism.* Dordrecht: D. Reidel Publishing, 1986.

Elga, Adam. "Subjective Probabilities Should Be Sharp." *Philosophers' Imprint* 10, no. 05 (2010).

Elliott, L. "Imaginative Adaptations: A Possible Environmental Role for the UN Security Council." *Contemporary Security Policy* 24, no. 2 (2003): 47-68.

Ellsberg, Daniel. "Risk, Ambiguity, and the Savage Axioms." *The Quarterly Journal of Economics* 75, no. 4 (1961): 643-69.

Enemark, Christian. *Disease and Security: Natural Plagues and Biological Weapons in East Asia.* New York: Routledge, 2007.

Ex parte Quirin, 317 U.S. 1 (1942).

Feldbaum, Harley. "U.S. Global Health and National Security Policy." Washington: Centre for Strategic and International Studies, 2009.

Feldman, Fred. *Doing the Best We Can.* Dordrecht: D. Reidel Publishing, 1986.

Feldman, Fred. "Hyperventilating About Intrinsic Value." *The Journal of Ethics* 2, no. 4 (1998): 339-54.

Feldman, Fred. "On the Consistency of Act- and Motive-Utilitarianism: A Reply to Robert Adams." In *Utilitarianism, Hedonism, and Desert,* edited by Fred Feldman, 65-75. Cambridge: Cambridge University Press, 1997.

Feldman, Richard. "Having Evidence." In *Evidentialism: Essays in Epistemology,* edited by Earl Conee and Richard Feldman. Oxford: Oxford University Press, 2004.

Gabriel, Ute, and Werner Greve. "The Psychology of Fear of Crime. Conceptual and Methodological Perspectives." *British Journal of Criminology* 43, no. 3 (2003): 600-14.

Gallie, W. B. "Essentially Contested Concepts." *Proceedings of the Aristotelian Society* 56 (1955): 167-98.

Gärdenfors, Peter, and Nils-Eric Sahlin. "Unreliable Probabilities, Risk Taking, and Decision Making." *Synthese* 53, no. 3 (1982): 361-86.

Giddens, Anthony. *Modernity and Self-Identity: Self and Society in the Late Modern Age.* Cambridge: Polity Press, 1991.

Good, I. J. "Rational Decisions." *Journal of the Royal Statistical Society. Series B (Methodological)* 14, no. 1 (1952): 107-14.

Goodin, Robert E. *Political Theory and Public Policy.* Chicago: University of Chicago Press, 1982.

Goodin, Robert E., and Frank Jackson. "Freedom from Fear." *Philosophy and Public Affairs* 35, no. 3 (2007): 249-65.
Gray, Emily, Jonathan Jackson, and Stephen Farrall. "Reassessing the Fear of Crime." *European Journal of Criminology* 5, no. 3 (2008): 363-80.
Gray, John N. "On the Contestability of Social and Political Concepts." *Political Theory* 5, no. 3 (1977): 331-48.
Griffin, James. *On Human Rights.* Oxford: Oxford University Press, 2008.
Griffin, James. *Well-Being: Its Meaning, Measurement and Moral Importance.* Oxford: Clarendon Press, 1986.
Hájek, Alan. "Arguments for–or against–Probabilism?" *The British Journal for the Philosophy of Science* 59, no. 4 (2008): 793-819.
Hájek, Alan. "Scotching Dutch Books?" *Philosophical Perspectives* 19, no. 1 (2005): 139-51.
Hájek, Alan, and Michael Smithson. "Rationality and Indeterminate Probabilities." *Synthese* 187, no. 1 (2012): 33-48.
Hale, C. "Fear of Crime: A Review of the Literature." *International Review of Victimology* 4, no. 2 (1996): 79-150.
Halevy, Yoram. "Ellsberg Revisited: An Experimental Study." *Econometrica* 75, no. 2 (2007): 503-36.
Hamdi v. Rumsfeld, 542 U.S. 507 (2004).
Hammond, Mason. "Res Olim Dissociabiles: Principatus Ac Libertas: Liberty under the Early Roman Empire." *Harvard Studies in Classical Philology* 67 (1963): 93-113.
Hampson, Fen Osler, Jean Daudelin, John B. Hay, Todd Martin, and Holly Reid. *Madness in the Multitude: Human Security and World Disorder.* Ottawa: Oxford University Press, 2002.
Hansen, Lene. "The Little Mermaid's Silent Security Dilemma and the Absence of Gender in the Copenhagen School." *Millenium: journal of international studies* 29, no. 2 (2000): 285-306.
Hansson, Sven. "The Harmful Influence of Decision Theory on Ethics." *Ethical Theory and Moral Practice* 13, no. 5 (2010): 585-93.
Harsanyi, John C. "Can the Maximin Principle Serve as a Basis for Morality? A Critique of John Rawls's Theory." *The American Political Science Review* 69, no. 2 (1975): 594-606.
Haslanger, Sally. "Gender and Race:(What) Are They?(What) Do We Want Them to Be?" *Nous* 34, no. 1 (2000): 31-55.
Haslanger, Sally. "What Are We Talking About? The Semantics and Politics of Social Kinds." *Hypatia* 20, no. 4 (2005): 10-26.
Hobbes, Thomas. "De Corpore Politico." In *Human Nature and De Corpore Politico*, edited by J.C.A Gaskin. Oxford: Oxford University Press, 1994.
Hobbes, Thomas. "Leviathan." edited by J.C.A Gaskin. Oxford: Oxford University Press, 1996.
Hoefer, Carl. "Determinism, Causal." In *Stanford Encyclopedia of Philosophy*, edited by Edward Zalta, 2004.
Horwich, Paul. *Probability and Evidence.* Cambridge: Cambridge University Press, 1982.
Howson, Colin, and Peter Urbach. *Scientific Reasoning: The Bayesian Approach.* La Salle: Open Court, 1989.
Hudson, Robert. "The Methodological Strategy of Robustness In the context of Experimental Wimp Research." *Foundations of Physics* 39, no. 2 (2009): 174-93.
Hurwicz, L. "Some Specification Problems and Applications to Econometric Models." *Econometrica* 19 (1951): 343-44.

Huysmans, J. "Security! What Do You Mean?: From Concept to Thick Signifier." *European Journal of International Relations* 4, no. 2 (1998): 226.
Jackson, Frank. "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection." *Ethics* (1991): 461-82.
Jeffrey, Richard C. *The Logic of Decision.* Chicago: University of Chicago Press, 1990.
Jensen, Karsten Klint. "A Philosophical Assessment of Decision Theory." In *Handbook of Risk Theory: Epistemology, Decision Theory, Ethics, and Social Implications of Risk,* edited by Sabine Roeser, Rafaela Hillerbrand, Per Sandin and Martin Peterson. Dordrecht: Springer, 2012.
Joyce, James M. "A Defense of Imprecise Credences in Inference and Decision-Making." *Philosophical Perspectives* 24, no. 1 (2010): 281-323.
Joyce, James M. "How Probabilities Reflect Evidence." *Philosophical Perspectives* 19, no. 1 (2005): 153-78.
Kagan, Shelly. "Rethinking Intrinsic Value." *Journal of Ethics* 2, no. 4 (1998): 277-97.
Kagan, Shelly. "Well-Being as Enjoying the Good." *Philosophical Perspectives* 23, no. 1 (2009): 253-72.
Kahneman, Daniel, and Amos Tversky. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47, no. 2 (1979): 263-92.
Keynes, J. M. "The General Theory of Employment." *The Quarterly Journal of Economics* 51, no. 2 (1937): 209-23.
King, Gary, and Christopher J.L. Murray. "Rethinking Human Security." *Political Science Quarterly* 116 (2001): 585-610.
Kirkpatrick, Lee A., and Seymour Epstein. "Cognitive-Experiential Self-Theory and Subjective Probability: Further Evidence for Two Conceptual Systems." *Journal of Personality and Social Psychology* 63, no. 4 (1992): 534-44.
Korematsu v. United States, 323 U.S. 214 (1944).
Kornblith, Hilary. "Justified Belief and Epistemically Responsible Action." *Philosophical Review* 92, no. 1 (1983): 33-48.
Krause, Keith. "The Key to a Powerful Agenda, If Properly Delimited." *Security Dialogue* 35, no. 3 (2004): 367-68.
Lackey, Jennifer. "What Luck Is Not." *Australasian Journal of Philosophy* 86, no. 2 (2008): 255-67.
Lasswell, Harold D. *National Security and Individual Freedom.* New York: McGraw-Hill, 1950.
Leaning, Jennifer, and Sam Arie. "Human Security: A Framework for Assessment in Conflict and Transition." Washington: United States Agency for International Development/Complex Emergency Response and Transition Initiative, 2000.
Leibniz, Gottfreid Wilhelm. "Letter of 1705." In *Die Werke von Leibniz,* edited by Onno Klopp. Hannover: Klindworth, 1864-1873.
Levi, Isaac. *Hard Choices: Decision Making under Unresolved Conflict.* Cambridge: Cambridge University Press, 1990.
Levi, Isaac. "Ignorance, Probability and Rational Choice." *Synthese* 53, no. 3 (1982): 387-417.
Levi, Isaac. "The Paradoxes of Allais and Ellsberg." *Economics and Philosophy* 2, no. 01 (1986): 23-53.
Levy, Neil. "What, and Where, Luck Is: A Response to Jennifer Lackey." *Australasian Journal of Philosophy* 87, no. 3 (2009): 489-97.
Lewis, David. "Causal Decision Theory." *Australasian Journal of Philosophy* 59, no. 1 (1981): 5-30.
Lewis, David. "Humean Supervenience Debugged." *Mind* 103, no. 412 (1994): 473-90.

Lewis, David. "A Subjectivist's Guide to Objective Chance." In *Studies in Inductive Logic and Probability*, edited by Richard C. Jeffrey, 83-132: University of California Press, 1980.
Lewis, David K. *On the Plurality of Worlds.* Oxford: Blackwell, 1986. Reprint, 2001.
Liddell, Henry George, and Robert Scott. *A Greek-English Lexicon.* Edited by Henry Stuart Jones and Roderick McKenzie. 9th ed. Oxford: Clarendon Press, 1940.
Linklater, Andrew. "Dialogic Politics and the Civilising Process." *Review of International Studies* 31 (2005): 141-54.
List, Christian. "Republican Freedom and the Rule of Law." *Politics, Philosophy & Economics* 5, no. 2 (2006): 201-20.
Liversidge v. Anderson, AC 206 (1942).
Locke, John. "Two Treatises of Government." London: J. M. Dent & Sons, 1690 [1924]. Reprint, 1990.
Lycan, William G. "Evidence One Does Not Possess." *Australasian Journal of Philosophy* 55, no. 2 (1977): 114 – 26.
Lyon, Aidan. "Deterministic Probability: Neither Chance nor Credence." *Synthese* 182, no. 3 (2011): 413-32.
Maclean, Sandra J. "Microbes, Mad Cows and Militaries: Exploring the Links between Health and Security." *Security Dialogue* 39, no. 5 (2008): 475-94.
Mathew, Gervase. "The Character of the Gallienic Renaissance." *The Journal of Roman Studies* 33 (1943): 65-70.
McDonald, Matt. "Securitization and the Construction of Security." *European Journal of International Relations* 14, no. 4 (2008): 563-87.
McEachrane, Michael. "Emotion, Meaning, and Appraisal Theory." *Theory & Psychology* 19, no. 1 (2009): 33-53.
McSweeney, Bill. *Security, Identity and Interests: A Sociology of International Relations.* Cambridge: Cambridge University Press, 1999.
Mellor, D.H. *Probability: A Philosophical Introduction.* New York: Routledge, 2005.
Mill, John Stuart. "Utilitarianism." In *On Liberty and Other Essays*, edited by John Gray. Oxford: Oxford University Press, 1991.
Mitzen, Jennifer. "Ontological Security in World Politics: State Identity and the Security Dilemma." *European Journal of International Relations* 12, no. 3 (2006): 341-70.
Moore, G. E. *Principia Ethica.* Cambridge: Cambridge University Press, 1903.
Morgan, Glyn. *The Idea of a European Superstate: Public Justification and European Integration.* Princeton: Princeton University Press, 2007.
Morgenthau, Hans. *Politics among Nations: The Struggle for Power and Peace.* 4th ed. New York: Alfred A. Knopf, 1965.
Neocleous, Mark. *Critique of Security.* Montreal: McGill-Queen's University Press, 2008.
Newman, Edward. "Critical Human Security Studies." *Review of International Studies* 36, no. 01 (2010): 77-94.
Nozick, Robert. *Anarchy, State, and Utopia.* Malden: Blackwell Publishing, 1974.
Nussbaum, M.C. *Creating Capabilities: The Human Development Approach*: Belknap Press of Harvard University Press, 2011.
Owen, Taylor. "Human Security-Conflict, Critique and Consensus: Colloquium Remarks and a Proposal for a Threshold-Based Definition." *Security Dialogue* 35, no. 3 (2004): 373.
Oxford Latin Dictionary. Vol. 2. Oxford: Oxford University Press, 1982.
Parfit, Derek. "Equality and Priority." *Ratio* 10, no. 3 (1997): 202-21.
Parfit, Derek. *On What Matters.* Vol. I. Oxford: Oxford University Press, 2011.
Parfit, Derek. *Reasons and Persons.* Oxford: Clarendon Press, 1984.

Paris, Roland. "Human Security: Paradigm Shift or Hot Air?" *International Security* 26, no. 2 (2001): 87-102.
Peoples, Columba. "Security after Emancipation? Critical Theory, Violence and Resistance." *Review of International Studies* 37, no. 3 (2011): 1113-35.
Pericchi, Luis Raúl, and Peter Walley. "Robust Bayesian Credible Intervals and Prior Ignorance." *International Statistical Review / Revue Internationale de Statistique* 59, no. 1 (1991): 1-23.
Peterson, Susan. "Epidemic Disease and National Security." *Security Studies* 12, no. 2 (2002): 43-81.
Pettit, Philip. "Consequentialism." In *A Companion to Ethics*, edited by Peter Singer. Oxford: Blackwell Publishers, 1991.
Pettit, Philip. "Decision Theory and Folk Psychology." In *Foundations of Decision Theory: Issues and Advances*, edited by Michael Bacharach and Susan Hurley, 147-75. Oxford: Blackwells, 1991.
Pettit, Philip. "Freedom and Probability: A Comment on Goodin and Jackson." *Philosophy & Public Affairs* 36, no. 2 (2008): 206-20.
Pettit, Philip. "Making Good: The Challenge of Robustly Demanding Values." *Uehiro Lectures* (2011), http://www.practicalethics.ox.ac.uk/lectures/2011_uehiro_lectures.
Pettit, Philip. *Republicanism: A Theory of Freedom and Government.* Oxford: Oxford University Press, 1999.
Pritchard, Duncan. *Epistemic Luck.* Oxford: Clarendon Press, 2005.
Rabin, Matthew. "Inference by Believers in the Law of Small Numbers." *The Quarterly Journal of Economics* 117, no. 3 (2002): 775-816.
Ramsey, Frank P. "Truth and Probability." In *Philosophical Papers*, edited by Frank P. Ramsey and D.H. Mellor. Cambridge: Cambridge University Press, 1990 [1926].
Rawls, J. *A Theory of Justice.* Rev. ed. Cambridge: Belknap Press, 1999.
Rawls, John. *Justice as Fairness: A Restatement.* Cambridge, MA: Belknap Press, 2001.
Raz, Joseph. *The Morality of Freedom.* Oxford: Oxford University Press, 1986.
Roe, Paul. "The 'Value' of Positive Security." *Review of International Studies* 34 (2008): 777-94.
Romm, Joseph J. *Defining National Security: The Nonmilitary Aspects.* New York: Council on Foreign Relations, 1993.
Roseman, I.J., and C.A. Smith. "Appraisal Theory: Overview, Assumptions, Varieties, Controversies." In *Appraisal Process in Emotion: Theory, Methods, Research*, edited by K.R. Scherer, A. Schorr and T. Johnstone, 3-19. Oxford: Oxford University Press, 2001.
Ross, W.D. *The Right and the Good.* Oxford: Oxford University Press, 1930.
Rothschild, Emma. "What Is Security?" *Daedalus* 124, no. 3 (1995): 53-98.
Sahlin, Nils-Eric. "Unreliable Probabilities, Paradoxes, and Epistemic Risks." In *Handbook of Risk Theory: Epistemology, Decision Theory, Ethics, and Social Implications of Risk*, edited by Sabine Roeser, Rafaela Hillerbrand, Per Sandin and Martin Peterson, 477-98. Dordrecht: Springer, 2012.
Scanlon, T. M. *What We Owe to Each Other.* Cambridge: Belknap Press, 1998.
Schlatter, Richard. "Thomas Hobbes and Thucydides." *Journal of the History of Ideas* 6, no. 3 (1945): 350-62.
Schroeder, Mark. "Having Reasons." *Philosophical Studies* 139 (2008): 57-71.
Seidenfeld, Teddy. "Decision Theory without "Independence" or without "Ordering"." *Economics and Philosophy* 4, no. 02 (1988): 267-90.
Selgelid, Michael J. "The Value of Security: A Moderate Pluralist Perspective." In *Ethics and Security Aspects of Infectious Disease Control: Interdisciplinary*

Perspectives, edited by Christian Enemark and Michael Selgelid. Burlington: Ashgate, 2012.
Sen, Amartya. *Development as Freedom.* Oxford: Oxford University Press, 1999.
Sen, Amartya. *Poverty and Famines: An Essay on Entitlement and Deprivation.* Oxford: Oxford University Press, 1983.
Shakespeare, William. "The Tragedy of Macbeth." In *William Shakespeare: The Complete Works (Compact Edition)*, edited by Stanley Wells and Gary Taylor. Oxford: Clarendon Press, 1988.
Shue, Henry. *Basic Rights: Subsistence, Affluence, and Us Foreign Policy*: Princeton University Press, 1980 [1996].
Smith, C. A., and L. D. Kirby. "Putting Appraisal in Context: Toward a Relational Model of Appraisal and Emotion." *Cognition & Emotion* 23, no. 7 (2009): 1352-72.
Smith, GM. "Into Cerberus' Lair: Bringing the Idea of Security to Light." *British Journal of Politics & International Relations* 7, no. 4 (2005): 485-507.
Smith, Holly. "Culpable Ignorance." *The philosophical review* 92, no. 4 (1983): 543-71.
Smith, Steve. "The Contested Concept of Security." In *Critical Security Studies and World Politics*, edited by Ken Booth. Boulder: Lynne Rienner Publishers, 2005.
Sorrell, Tom. "Hobbes's Moral Philosophy." In *The Cambridge Companion to Hobbes's Leviathan*, edited by Patricia Springborg, 128-56. Cambridge: Cambridge University Press, 2007.
Staley, Kent W. "Robust Evidence and Secure Evidence Claims." *Philosophy of Science* 71, no. 4 (2004): 467-88.
Steele, Brent J. *Ontological Security in International Relations: Self-Identity and the Ir State.* New York: Routledge, 2008.
Stocker, Michael. "Rightness and Goodness: Is There a Difference?" *American Philosophical Quarterly* 10, no. 2 (1973): 87-98.
Striker, Gisela. "Ataraxia: Happiness as Tranquility." *The Monist* 73, no. 1 (1990): 97-110.
Swanton, Christine. "On the "Essential Contestedness" of Political Concepts." *Ethics* 95, no. 4 (1985): 811-27.
Talbot, William. "Bayesian Epistemology." In *The Stanford Encyclopedia of Philosophy*, edited by Edward Zalta, 2008.
Tegmark, Max, and Nick Bostrom. "Astrophysics: Is a Doomsday Catastrophe Likely?" *Nature* 438, no. 7069 (2005): 754-54.
Thomas, Nicholas, and William T. Tow. "The Utility of Human Security: Sovereignty and Humanitarian Intervention." *Security Dialogue* 33, no. 2 (2002): 177-92.
Thucydides. "History of the Peloponnesian War." In *Perseus Digital Library*, ed Thomas Hobbes, 1843.
Thucydides. "The Peloponnesian War."
Tversky, A., and D. Kahneman. "Advances in Prospect Theory: Cumulative Representation of Uncertainty." *Journal of risk and uncertainty* 5, no. 4 (1992): 297-323.
Tversky, Amos, and Daniel Kahneman. "The Framing of Decisions and the Psychology of Choice." *Science* 211, no. 4481 (1981): 453-58.
U.S. Constitution.
Ullman, Richard H. "Redefining Security." *International Security* 8, no. 1 (1983): 129-53.
United Nations Development Program. "Human Development Report 1994." New York: United Nations Development Program, 1994.
Universal Declaration of Human Rights. G.A. res. 217A (III).

Valverde, Mariana. "Governing Security, Governing through Security." In *The Security of Freedom: Essays on Canada's Anti-Terrorism Bill*, edited by Ronald J. Daniels, Patrick Macklem and Kent Roach. Toronto: University of Toronto Press, 2001.

Venkatapuram, Sridhar *Health Justice: An Argument from the Capabilities Approach.* London: Polity, 2011.

Vessel, Jean-Paul. "Supererogation for Utilitarianism." *American Philosophical Quarterly* 47, no. 4 (2010): 299-319.

Vogelstein, Eric. "Subjective Reasons." *Ethical Theory and Moral Practice* 15, no. 2 (2012): 239-57.

von Neumann, J., and O Morgenstern. *Theory of Games and Economic Behaviour.* Princeton, N.J.: Princeton University Press, 2004 [1944].

Wæver, Ole. "Securitization and Desecuritization." In *On Security*, edited by Ronnie D. Lipschutz, 46-86. New York: Columbia University Press, 1995.

Wald, A. *Statistical Decision Functions.* New York: Wiley & Sons, 1950.

Waldron, Jeremy. "Safety and Security." *Nebraska Law Review* 85, no. 2 (2006): 454-507.

Waldron, Jeremy. "Security and Liberty: The Image of Balance." *The Journal of Political Philosophy* 11, no. 2 (2003): 191-210.

Waldron, Jeremy. "Security as a Basic Right (after 9/11)." In *Global Basic Rights*, edited by Charles R. Beitz and Robert E. Goodin, 207-26. Oxford: Oxford University Press, 2009.

Walker, RBJ. "The Subject of Security." *Critical Security Studies: Concepts and Cases* (1997): 61–82.

Walt, Stephen M. "The Renaissance of Security Studies." *International Studies Quarterly* 35, no. 2 (1991): 211-39.

White, Roger. "Evidential Symmetry and Mushy Credence." *Oxford Studies in Epistemology* 3 (2010): 161-86.

Williams, Bernard. *Moral Luck.* Oxford: Oxford University Press, 1981.

Williamson, Timothy. *Knowledge and Its Limits.* Oxford: Oxford University Press, 2000.

Windschitl, Paul D., and Gary L. Wells. "Measuring Psychological Uncertainty: Verbal Versus Numeric Methods." *Journal of Experimental Psychology: Applied* 2, no. 4 (1996): 343-64.

Wolfers, Arnold. "National Security as an Ambiguous Symbol." In *Discord and Collaboration: Essays on International Politics*, 147-65. Baltimore: Johns Hopkins Press, 1962.

Wolff, Robert Paul. *In Defense of Anarchism.* Berkeley: University of California Press, 1998.

Wyn Jones, R. "Security, Strategy, and Critical Theory." *London: Lynne* (1999).

Yudkowsky, Eliezer. "Cognitive Biases Potentially Affecting Judgement of Global Risks." In *Global Catastrophic Risks*, edited by Nick Bostrom and Milan M. Cirkovic. Oxford: Oxford University Press, 2008.

Zack, Naomi. *Ethics for Disaster.* Lanham: Rowman & Littlefield Pub Inc, 2009.

Zimmerman, Michael J. *The Concept of Moral Obligation:* Cambridge University Press, 1996.

Zimmerman, Michael J. "The Good and the Right." *Utilitas* 19, no. 3 (2007): 326-53.

Zimmerman, Michael J. *Living with Uncertainty: The Moral Significance of Ignorance.* Cambridge: Cambridge University Press, 2008.