

The Dynamic Foundations of Epistemic Rationality

Barry Lam

A dissertation presented to the faculty of Princeton University
in candidacy for the degree of Doctorate of Philosophy.

Recommended for acceptance by the Department of Philosophy

January, 2007

©Copyright by Barry Lam, 2007. All rights reserved.

Abstract

Classical theories of epistemic rationality take an agent's individual beliefs to be the only things that are rational or irrational. For them, rationality is wholly *static*. Recent work in epistemology take sets of individual beliefs and also *changes of belief over time* to be rational or irrational. For these theories, rationality is both static and *dynamic*. However, for both groups, static rationality is fundamental. In my dissertation, I argue to the contrary that, in fact, all rationality is dynamic rationality. Epistemic reasons, rationality, and justification as applying *only* to *changes* of belief. This wholly dynamic view of rationality, which I call "Dynamicism" has wide-ranging epistemological consequences. A small set of simple, elegant, and independently motivated principles of dynamic rationality can illuminate and solve otherwise interminable epistemological disputes.

Chapter One refutes the view that dynamic rationality can be reduced to static rationality. Chapters Two and Three extend the arguments from Chapter One to argue for skepticism about the concept of static rationality altogether. Chapter Four adopts Dynamicism as a working hypothesis and presents a dynamic proto-theory of rationality for all-or-nothing beliefs. Chapter Five turns to the philosophical upshots of Dynamicism for some problems of epistemology. Of central importance is the way in which Dynamicism reformulates a variety of issues centered around the problem of skepticism, including the dialectic between skepticism and dogmatism, closure principles for justification, and Conservatism.

Contents

Abstract	iii
Preface	vi
Introduction	ix
1 Against Static Rationalism	1
1.1 Static Rationalism Explored	5
1.1.1 Full Reduction	5
1.1.2 The Preface Paradox [Makinson, 1965]	10
1.1.3 Static Rationalism and The Basing-Relation	12
1.1.4 The Irrelevance of Quality of Basing	17
1.1.5 The Basing-Dilemma	20
1.2 Summary	26
2 Against Static Rationality	
Part I	28
2.1 Minimal Static Constraints	29
2.2 Rational Degrees of Belief	31
2.2.1 A Simple Strengthened Preface Case	34
2.2.2 A Complicated Strengthened Preface	39
2.3 Objections	44
2.3.1 Replies	48
2.3.2 Lessons from the Preface Cases	56
2.4 Ideal Rationality to the Rescue?	57
2.4.1 Formality, Resemblance, and Epistemic Imperfection	59
2.4.2 Alternatives to Resemblance and Formality?	65
2.5 Summary	68

3	Against Static Rationality	
	Part II	69
3.1	Insufficient Evidence and Forgetting	71
3.2	Basing on Defeated Evidence	76
3.3	Basing and Changes of Grounds	77
	3.3.1 Dynamicism and Rational Changes of Grounds	78
	3.3.2 Objective versus Subjective Grounds	80
3.4	Summary: The Dynamicist Argument Schema	84
4	Dynamic Rationality	87
4.1	Replacing Static with Dynamic Rationality	87
	4.1.1 The Dynamic Foundations of Epistemic Norms	88
	4.1.2 The Principles of Dynamic Rationality	91
4.2	The Rationality of Changes of Grounds	101
4.3	Rational Changes of Degrees of Belief	102
4.4	Summary	103
5	Dynamicism Vindicated	104
5.1	Application 1: Consistency and Rationality	104
5.2	Application 2: Skepticism, Dogmatism, and Closure	109
	5.2.1 Closure	111
	5.2.2 Dogmatism Revisited	115
5.3	Application 3: Conservatism in Epistemology	117
5.4	Further Remarks on the Relationship between Logic and Ra- tionality	120
5.5	Side benefits of a Dynamic Theory of Rationality	122
5.6	Conclusion	124
	Afterword: Hard Problems	125

Preface

I came to graduate school as a young philosophical naturalist, suspicious of the contribution that a priori philosophical thinking could make to the “one true source” of human knowledge, the empirical sciences. Much of that naive stance is now gone, eroded by the five-year practice of philosophizing in the Princetonian way. Still, the old naive stance seems to linger in my taste in philosophical problems. I remain preoccupied more with issues concerning normativity, methodology, and inquiry rather than the nature, definition, and identities of things, perhaps betraying some lingering skepticism about matters metaphysical. Nonetheless, upon reflection on the substance of my thesis now that it is complete, it seems to me that I have done as much metaphysics as epistemology.

Graduate school taught me many valuable lessons. As I discovered in presenting the material of this thesis to a variety of audiences, different people place very different values on the problems, arguments, and examples I have constructed. What is compelling to one is unconvincing to another; what seems essential to one is orthogonal to another. The central claim of this dissertation is considered by some to be banal, and others to be barely defensible. In writing a doctoral dissertation, I have learned that philosophy is difficult, not only because it is an unfettered market for innovation where

little is to be had, but also because it is an endless attempt to convince other philosophers that your ideas might be worth their fleeting attention. Contrary to my initial expectations, I discovered that articulating a thesis for an audience of professional philosophers involves as much skill in rhetoric as in rigor. I also discovered that an extended, well-argued, and successful refutation of your work is a high form of compliment and flattery in philosophy, for it implies that some keen mind considered your work valuable enough to be worth their time and attention. More common, however, is unreflective dismissal. Finally, I have learned that time and hard work are never enough. Only when such things are punctuated by episodes of luck and spontaneous insight can good philosophical ideas ever develop.

The questions I address in this thesis are as follows. What are we fundamentally evaluating when we judge people and their beliefs reasonable or unreasonable? What is the relationship between the rationality of belief and the rationality of changes of belief? What kind of rationality should epistemologists focus on as the fundamental notion in epistemological theory, and what form should such theories take? I believe my answers to these questions are true, but I also believe that my arguments are not fully conclusive. My aim is to make my case persuasively without any pretense to its decisiveness. Perhaps the best lesson I have taken away from Princeton is the ability to find merit in philosophy that fails to rest on irrefutable arguments.

I am indebted to audiences at Princeton University, the University of Kansas, Swarthmore College, and Vassar College for discussions of earlier drafts of the Introduction and Chapter One. I am indebted to many discussions of ideas and drafts with Gideon Rosen, Tom Kelly, Brett Sherman,

Mark Schroeder, and Sarah-Jane Leslie. Much of my work betrays my intellectual heritage at the Department of Logic and Philosophy of Science at UC Irvine, where Jeffrey Barrett, Kyle Stanford, and Brian Skyrms were my first epistemology teachers. Most of all I am deeply indebted to my advisor, James Pryor, for attention to the details of my views at every stage of development. I also want to thank David Christensen for publishing a book in early 2005 that helped me develop my ideas when they were in their infancy, and Gil Harman for having a counterexample for every generalization (including this one). Finally, I am indebted to writing-specialist extraordinaire, Shanna Andrawis, and our two furry children, Cody and Ally, for many long, thought-provoking afternoon walks on which I developed many of the views in this thesis.

Everything in this Preface is true. But undoubtedly, one of the claims I make in the upcoming pages will be false. Worse yet, it is bound to have at least one grammatical or spelling error. But of course, none of these facts alone render my views less reasonable, or my thesis less passable.

Introduction

All rationality is dynamic rationality. The next one hundred and fifty pages will be an explication and defense of this thesis, a thesis I call *Dynamicism*. The kind of rationality of concern to me is epistemic rationality (sometimes called “theoretical rationality”). Epistemic rationality is the kind of rationality associated with people as agents who believe, disbelieve, suspend judgment, and have varying degrees of confidence in, the truth of propositions. In contrast, *practical* rationality is the kind of rationality associated with people as agents of purposive actions. Some take believing to be a kind of purposive activity, and thus, would reduce epistemic rationality to practical rationality. A philosophical tradition traceable from William James to Bas van Fraassen understands belief and epistemic rationality in this way. I however do not. I do not think of belief as purposive, even less as an action or activity. But even those who reduce belief to a kind of purposive activity must admit a sense of rationality of belief that is irreducible to the rationality of purposive action. Suppose you offer me a million dollars on the condition that I believe something I currently do not believe, and have much evidence against, for instance, that I can defeat Roger Federer in a tennis match. Suppose further that in addition to the immediate monetary benefit I reap from adopting

such a belief, I can foresee that such a belief will give me great emotional pleasure. Relative to my personal goals for wealth and well-being, I would be very practically reasonable in taking on the belief. Yet there is still a sense in which I would be unreasonable if I believed that I could beat Roger Federer. The sense in which I am unreasonable is epistemic in a way that cannot be reduced to practical rationality. (Compare the case to an offer of a million dollars to eat a delicious slice of pizza, something I clearly *do* rather than believe.) As we can see, epistemic rationality is not tied to my personal goals and aims in any straightforward way. Thus, even if there were a sense of “practically reasonable” that applied to belief, it is still a wholly different sense of “reasonable” from epistemically reasonable belief (see [Kelly, 2003] for a defense of this view along these exact lines). According to Dynamicism, *epistemic* rationality is the kind of rationality that is wholly dynamic.

What does it mean to characterize all epistemic rationality as *dynamic*? Epistemic rationality is wholly dynamic just in case the fundamental object of epistemic evaluation is a *change* of belief (or degrees of confidence) over time. If all epistemic rationality is dynamic, then reasons, permissions, requirements, rationality, and irrationality are normative properties and relations that apply to changes of (degrees of) belief, not to the (degrees of) beliefs themselves. In other words, what we evaluate when we evaluate people epistemically, what we track when we make epistemic evaluations, are changes of mind, from one belief to another, from one set of beliefs to another, from one doxastic state to another. And most importantly, the Dynamicist thesis claims that *we do not track anything else* but changes of belief-states when we evaluate people epistemically. To say that epistemic rationality is

wholly dynamic is to say that changes of belief-states are the fundamental and sole objects of epistemic evaluation. Any other apparent objects of epistemic evaluation can be fully explained in terms of the rationality of changes of belief.

Consider a view analogous to Dynamicism in political philosophy. In order to determine whether a society is just, we might need to take a snapshot of it at a certain time, and look at the distribution of wealth, power, and other internal relationships within that society at that time. A wholly static conception of justice claims that the nature of those relationships fully determine whether a society is just or unjust. On the other hand, how the society develops, changes, or transforms might fully determine the its justness, regardless of the internal relationships that result from such developments or transformations. For instance, a completely uneven distribution of wealth and power might have been the outcome of a free democratic process, an outcome determined completely by the will of a free-thinking and freely-associating majority. To be a Dynamicist about *justice* is to claim that justice lies not in the outcomes, the snapshot of a society, but in the manner in which such societies come about. It is the processes of change and transformation that are just or unjust, not the results of such change or transformation. Whatever we think about the case of justice, we understand what it means to be a Dynamicist about justice, we understand how to argue about its competing views, and we understand its importance for political philosophy. Dynamicism about rationality is precisely the corresponding view in epistemology. Rationality applies fundamentally to the changes and transformations of belief, and not to beliefs or sets of beliefs that are the

results of such changes or transformation.

The thesis I defend, that all rationality is dynamic rationality, is relatively straightforward. But one might be skeptical of the philosophical significance of a project that seeks to identify and understand the fundamental object of epistemic evaluation. Why does it matter whether that object is beliefs, or changes of beliefs, both, or neither? In Chapter 4, I hope to show that Dynamicism is significant in its ability to resolve and dissolve certain epistemological problems. However, some general philosophical motivations for the project are worth mentioning.

In ordinary life, we evaluate many different things as rational or irrational. We evaluate people, their beliefs, and the way in which they reason and change their mind. Similarly, we evaluate many different things for goodness. We evaluate people, their actions, and their reasons for action. Yet a general philosophical account of goodness might reduce or explain the goodness of persons in terms of the goodness of their actions, and the goodness of their actions in terms of the goodness of their reasons for action. That is, the best philosophical theory might tell us that fundamentally, there is one thing that is good, someone's reason for action, that in turn explains the goodness of actions and the goodness of people. The same can be said for the epistemological case. We want to seek the fundamental object of epistemic evaluation simply to establish a metaphysical fact of the matter concerning the order of explanation in the normative realm generally. In the moral realm, a version of this problem goes back to Plato's *Euthyphro*, when Socrates asks Euthyphro whether an act was pious because the Gods commanded it, or whether the Gods commanded it because it was pious. For

Plato, an answer to the problem of order of explanation in the moral realm implied a metaphysical view about the relationship between mind and world, between empirical fact and normative evaluation. The epistemological problems of order of explanation should be of equal interest for those who seek the true structure of the normative realm, and for those who seek the source of normative judgments and normative facts.

But more importantly, when we attend to the fundamental objects of evaluation, the very things that are bearers of goodness, rightness, or rationality, we are seeking a form of methodological guidance. That is, we not only want to know what to look for in the world when we seek goodness, or rationality, but also how to proceed by way of theorizing about such things. We want to know how our theories are going to carve normative reality; we are looking for the fundamental objects of evaluation so that we proceed to carve normative reality at the joints, and not the bone. In other words, how should we formulate the elusive norms that govern our judgments and evaluations of ourselves and others? Are they norms that govern people, actions, or reasons? Answering this question has serious methodological repercussions. It tells us how to formulate thought experiments, what norms to argue about, and how to argue about them. So when I seek the fundamental object of epistemic evaluation, I am looking not only to find out some metaphysical fact of the matter, I am seeking guidance on how to do normative epistemology.

The dissertation to follow is both an attempt to formulate and make vivid the problem of the fundamental object of epistemic evaluation, as well as characterize and argue for a certain solution to that problem, namely Dynamicism. It follows from Dynamicism that epistemology should solely

seek norms for belief-acquisition and revision. My aim in the fourth chapter is to provide the outlines of just such a theory.

Chapter One defines and refutes the view that all rationality is static. Static Rationalism states that rationality and reasons apply only to the beliefs or sets of beliefs at a time. In Chapter One, I argue against a view that follows from Static Rationalism, namely, that the rationality of belief-changes can be explained in terms of the rationality of belief. I present a series of counterexamples that show the incoherence and self-defeating character of reducing dynamic rationality to static rationality. The success of Chapter One shows that dynamic rationality is at least as basic as static rationality.

Chapters Two and Three extend the arguments from Chapter One to argue for skepticism about the concept of static rationality altogether. Any seemingly independent conditions on static rationality are either false or reducible to conditions on dynamic rationality. In Chapter Two, I argue that the Preface paradox is merely an instance of a general form of argument that shows that there cannot be irrationality in beliefs without irrationality in changes of belief. I argue that such paradoxes show that there are no formal constraints on beliefs or degrees of belief at a time because agents can lack any rational fault while having beliefs or degrees of belief that are inconsistent. I then argue that appealing to epistemically perfect, ideally rational agents to respond to such objections cannot work. In Chapter Three, I argue that non-formal constraints for static rationality like appropriate epistemic “basing” either fail or are otherwise reducible to constraints on reasonable changes of belief. I conclude from Chapters 2 and 3 that a strong case can be made that there cannot be any kind of epistemic rationality over and above dynamic

rationality.

Chapter Four adopts Dynamicism as a working hypothesis and presents a dynamic proto-theory of rationality for all-or-nothing beliefs. I define the idea of reasoning from one point to another, where each point is an individual or set of beliefs or disbeliefs. I then argue from simple cases for a series of norms for reasoning, and a relation of *priority* that arranges the norms in an order of applicability.

Chapter Five turns to the philosophical upshots of Dynamicism for some problems of epistemology. Of central importance is the way in which Dynamicism reformulates a variety of issues centered around the problem of skepticism, including the dialectic between skepticism and dogmatism, closure principles for justification, and Conservatism.

Finally, in the Afterword, I list a set of difficult problems that a Dynamicist must solve before the view is fully palatable.

Chapter 1

Against Static Rationalism

Two people can believe the same thing, but differ in the rationality of their belief. Jack and Jill start out believing in God. Jack is a well-studied theologian, while Jill simply believes whatever her friends and family do. Jack's belief in God is reasonable, whereas Jill's is not. Two people can also change their beliefs, but differ in the rationality of their belief-changes. Suppose Jack asked God for a red fire truck for Christmas and did not get it, and decided that God did not exist. Jill, on the other hand, discovered the Problem of Evil in a philosophy class, and realized she believed the premises of the argument much more confidently than she believed in God. So she concluded that God did not exist. Jill changed her mind in a reasonable way, whereas Jack did not.

The case of Jack and Jill suggests that in ordinary, everyday epistemic evaluation, there are at least two different things that we are evaluating; beliefs, a psychological state people have at a certain time, and reasoning, or how people change their mind from one belief to another. Thus, the case of

Jack and Jill suggests that in normative epistemology, we should be asking two different questions.

Question One: *When is a belief of yours rational or justified?*

What makes it the case that Jack's belief in God was reasonable whereas Jill's was not. In philosophical jargon, what are the conditions of *static* rationality?

Question Two: *When is a change of belief rational or justified.*

What makes it the case that Jill's change of belief was reasonable, whereas Jack's was not? In jargon, what are the conditions of *dynamic* rationality?

This chapter will consider the relationship between Question One and Question Two. In other words, what is the relationship between rational belief and rational changes of belief? What is the relationship between static and dynamic rationality?¹The distinction between static and dynamic rationality is not the distinction between the properties of static and dynamic coherence found in probability theory. Static and dynamic coherence are *logical* concepts, stating the conditions under which a set of probabilities, or a pair consisting of such a set and a function from one set to another set, are inconsistent (see [Skyrms, 1993] for discussion of the logical concepts of static and dynamic coherence). Static and dynamic rationality are *normative* concepts that state the conditions under which a certain kind of object is rational or irrational. This distinction between a logical property

¹During the development of this chapter, I discovered that Anthony Gillies has defended a version of the thesis I defend in this chapter in an unpublished paper entitled "Two More Dogmas of Belief-Revision." Our arguments differ, but I owe him much in the manner in which I frame some of the issues I discuss here.

or concept and normative property or concept will be crucial for the rest of the discussion in this chapter.

The appearance that there are two different sorts of things in ordinary life that we evaluate epistemically should not lead us to conclude that at the fundamental level, there are two completely different sorts of things that are rational or irrational, and that the conditions that determine rational belief and rational changes of belief are independent. Indeed it appears that a guiding assumption in current epistemological theory is that static rationality is primary. Static rationality has been the primary object of study in normative epistemology since Descartes inaugurated the subject in contemporary form. In cataloging and rejecting every standing belief for which he had some reason to doubt, the “I” of the first Meditation seemed to be advocating a principle of static rationality: your current beliefs are not fully rational unless you have absolutely no reason to doubt them. Today, the many “isms” that have descended from Cartesian epistemology, such as Foundationalism, Coherentism, Evidentialism, Conservatism, and Infinitism, all share the presumption that rationality primarily applies to an agent’s currently-held beliefs. And while there is much disagreement over substantive epistemological principles, all “isms” agree that such principles are principles of static rationality.

The bias toward thinking about rationality in wholly static terms extends even to contemporary formal epistemology, which purports to be about rational belief-change (see [Hansson, 1999] for an introduction to the subject). Many such theories take static rationality to be conceptually primary. Almost all theories of rational belief-change begin with some version of static Coherentism ([Harman, 1986], [Gardensfors, 1988], many versions of Bayesianism)

or Foundationalism ([Pollock, 1986], [Pollock, 1995]) as the basis of their dynamic theories. According to these theories, rational belief-change is a matter of acquiring, preserving, or restoring static rationality, whatever the latter turns out to be. To change your mind reasonably is to change your mind in such a way as to end up with rational beliefs (or beliefs that are progressively rational). Let us call this bias or presumption toward static rationality “Static Rationalism.” Static Rationalism is both a presumption about the primary objects of epistemological evaluation, and a methodological commitment that arises out of such a presumption. A Static Rationalist assumes that beliefs and sets of beliefs are the primary things that are rational or irrational, and that the task of the epistemologist is to characterize the conditions under which beliefs or sets of beliefs are rational or irrational. Any remaining notion of dynamic rationality, or rational changes of belief, can be understood wholly in terms of static rationality.

Logically speaking, there are two equally coherent competitors to Static Rationalism. Let *Binary Rationalism* be the view that both static and dynamic rationality are conceptually basic and that neither are reducible to the other. Let *Dynamicism* be the view that dynamic rationality is fundamental and exhaustive of epistemic rationality.

My aim in this chapter is to argue that Static Rationalism fails to give the right account of the relationship between rational belief and rational changes of belief. It follows from this failure that dynamic rationality is at least as basic as, and cannot be defined in terms of, static rationality. Static Rationalism is therefore false. The next two chapters will argue against the viability of static rationality altogether, from which it follows that Binary

Rationalism must be false.

1.1 Static Rationalism Explored

Static Rationalism implies that dynamic rationality can be reduced or explained in terms of static rationality. Is this implication true? I will proceed to construct a series of theses reducing rational changes of belief to rational belief on behalf of Static Rationalists. For each thesis, I will present a series of problems that will culminate in a strong case for the irreducibility of dynamic rationality.

1.1.1 Full Reduction

The first thesis I will look at is fully reductionist. It says that the rationality of a change of belief is completely reducible to the rationality of belief.

SR Thesis 1: A change of belief is rational just in case it is a change from a rational (set of) belief(s) to a rational (set of) belief(s).

Thesis 1 is admittedly a straw man, but examining it will be nonetheless instructive. The problem with Thesis 1 is pretty straightforward. It both undergenerates and overgenerates. That is, it both rules out as irrational changes of beliefs that are in fact rational, and rules in as rational changes of beliefs that are in fact irrational. Consider the possibility of Janitorial Jay.

Janitorial Jay: Jay's belief-system is always a complete mess at the start of a research project. Nonetheless Jay manages to

straighten out his thoughts quite well. He decides to research the Southern California freeway system. At t_1 he believes:

1. Interstate 405 runs North-South.
2. The 55 freeway runs North-South.
3. The 405 and 55 intersect in Costa Mesa, CA.

Jay proceeds to reason to these beliefs at t_2 :

4. No two things can run both North-South and also intersect.
5. Either the 405 and 55 do not intersect, or at least one of the two does not run exactly North-South.

Jay seems to have reasoned impeccably, reasonably, from his irrational beliefs to rational beliefs. But of course this is not captured in the fully reductionist Thesis 1.

There is any easy fix to the problem of undergeneration. We simply require that your beliefs have to end up rational, they don't need to start out rational, in order for a change of belief to be rational. Thus Janitorial Jay changes his beliefs rationally because he ended up rational, even though he did not start out rational.

There is a more serious problem with the reductionist thesis. It is an accepted fact that people with varying numbers of beliefs can be statically rational. That is, you can be rational, someone who believes everything you believe, but believes just a little bit more can be rational, and someone who believes just a little bit more than that person can be rational. Similarly, someone who believes less than you can be rational. As a result of these

accepted facts, Thesis 1 runs into a problem of overgeneration. Consider Eager Elle.

Eager Elle loves US geography so much that she'll jump to any conclusions she can to have more beliefs about US cities. Eager Elle believes at t_1 that

1. Chicago is west of Cleveland,
2. Cleveland is west of New York,
3. New York is North of Miami, and
4. Miami is south of Chicago.

From these beliefs, Elle reasons to the following beliefs at t_2 solely from the previously mentioned beliefs:

5. New York is east of Chicago

But also concludes

6. New York is north of Chicago
7. Cleveland is north of Miami.

and she adds any beliefs she needs to add to make 6-7 rational, such as:

8. Cleveland is south of New York and north of Chicago.

Elle very reasonably concludes 5. But Elle is so eager, that from 1-4, she also reasons to 6-7. And furthermore, she reasons to 8, which is a belief that supports 6-7.

Elle ends up with a set of rational beliefs. They are the same beliefs that a very reasonable geographer would have. But Elle is not rational in the way

she has reasoned or managed her beliefs. She has jumped to conclusions, a paradigm case of an unreasonable change of belief! Eager Elle is a case of overstocking yourself with beliefs in an unreasonable way. The mirror case would be one in which you purge your beliefs in an unreasonable way. Consider Eager Elle's brother.

Wishful Will: Will starts with the same beliefs as his sister Elle, and reasons to 5 just like Elle, but not to 6-8. However, Wishful Will also has a nagging belief that his wife is cheating on him. So in reasoning from 1-4 to 5, Will also gives up the belief that his wife is cheating on him, and all the beliefs that support that belief.

Will ends up with statically rational beliefs. They are the beliefs he would have had had he never had any evidence or knowledge of his wife's cheating. But intuitively he has changed his beliefs in an unreasonable way. It is typically unreasonable to go from your beliefs about geographical facts to the revision of beliefs about marital fidelity.

Eager Elle and Wishful Will start with statically rational beliefs, and end with statically rational beliefs. In Elle's case, she has overstocked herself. In Will's case, he has overpurged. But being overstocked or overpurged in these cases does not make the set of beliefs they end up with irrational. Nonetheless, Elle and Will have changed their mind in an unreasonable way. So our full reductionist thesis cannot be true.

The easy fix to the problem of undergeneration will not suffice for the problem of overgeneration. To deal with the problem of overgeneration, we

are going to need to constrain the sets of rational beliefs that Elle and Will can end up with given their starting points. One constraint we find in the literature is this:

Minimality: Your change of mind from A to B is rational only if it is a minimal change. A change of mind is minimal if and only if you do not add too much and you do not give up too much.

(For versions of Minimality, see [Harman, 1986], [Gardensfors, 1988], and Bayesianism ala [Horwich, 1982], [Howson and Urbach, 1993].)

Setting aside the issue of what “too much” amounts to, it seems that Minimality gets the job done in these cases. According to Minimality, Eager Elle is irrational in her reasoning because she added too many beliefs. Wishful Will is irrational because he gave up too much. With the addition of Minimality, we have given up the fully reductionist thesis and arrived at the following principle reducing dynamic to static rationality.

Weaker Static Rationalism: A change of belief from set A to set B is rational if and only if B is statically rational, and the change from A to B is minimal.

Weaker Static Rationalism is actually the thesis that forms the foundation of almost every contemporary theory of rational change of belief today [Hansson, 1999], [Gardensfors, 1988]. Notice that this weaker version of static rationalism will not be fully reductionist. The condition of a minimal change cannot be a notion reducible to static rationality. A set of conditions that tells us what properties a set of beliefs must have in order to be rational will need to characterize many sets of beliefs corresponding to believers

who have more or less information about the world. You cannot derive the idea of one set being a minimal change of another from a set of conditions telling you what the good and bad sets are. However, even though Weaker Static Rationalism is not fully reductionist, the concept of a rational belief still plays an essential role in determining which changes of belief are rational and irrational, and that is consistent with Static Rationalism.

Minimality has been challenged on the grounds that it is either too vague a notion for doing the work required of it in formal theories of rational belief-change, or otherwise descriptively inadequate as a constraint on rational changes of belief (see [Rott, 2000] for arguments to this effect). The details concerning the Minimality Condition will not concern me here. Instead, it is sufficient to focus on the other, often unchallenged condition, namely, that beliefs need to end up rational in order for a change of belief to be rational. There is a well-known paradox that seems to suggest otherwise.

1.1.2 The Preface Paradox [Makinson, 1965]

You are a researcher of Hudson Valley History. There are many things you do not know, but you diligently go through the archives, interview citizens and scientists, and compile your evidence into a book representing everything you now believe about the Hudson Valley. Now you look at the history of history books, and your own past history of error, and conclude that there must be at least one false claim in your book, since it is highly unlikely that you were perfect in your research. You write in the Preface “Because of the inability of anyone to be absolutely perfect in their research no matter how careful, one of the claims in the body my book is bound to be false.” Now

the book represents the set of your beliefs about Hudson Valley history, but no matter what way the world is, the claims in the book cannot all be true. So they are inconsistent.

On the assumption that consistency is a constraint on rational beliefs, you have ended up with irrational beliefs. But the question is, does this show that your reasoning was irrational? Did you irrationally reason from the evidence you had to your new beliefs about Hudson Valley history? Did you irrationally reason from your beliefs about past error to the belief that one of the claims in the body of your book is false? And given that you now have inconsistent beliefs, are you unreasonable in failing to change your mind? It certainly doesn't seem so. It seems like we have set up a case where your changes of mind are flawless. Someone who didn't reason to the preface conclusion is much worse than you with respect to rationality. That person would have had evidence that they were imperfect, but would have nonetheless believed that everything in their book was true. You are much more rational than someone who would have ended up with consistent beliefs. Yet you've ended up with irrational beliefs, though your reasoning was fully justified! ²

There are many different lessons we can draw from the Preface Paradox. Maybe consistency is not a constraint on rational belief after all. This is indeed the view that I like, and the view that I think puts serious pressure on the very theoretical fruitfulness of independent conditions for rational belief.

²Notice that we cannot weaken the condition to the condition that beliefs must end up *more* rational than the starting point. So long as we are assuming that consistency is a constraint on rationality, it isn't clear how to make sense of the idea that starting out consistent and ending up with inconsistent beliefs is ending up with beliefs that are more rational than before.

However, I will return to it later. On behalf of the Static Rationalist, let us assume that logical consistency is a constraint on rational beliefs at a time. In that case, ending up with rational beliefs is not a necessary condition on a rational change of mind. So our weaker thesis of Static Rationalism will not be the right characterization of the way in which rational beliefs determine rational changes of mind. Back to the drawing board.

1.1.3 Static Rationalism and The Basing-Relation

I do not think that Preface cases leave Static Rationalism dead in the water. What is essential to static rationalism is that the rationality of belief is always needed as an explanation for rational changes of mind. When we ask why a change of mind is rational, a static rationalist claims that you always need to cite something about rational beliefs to answer the question. Even if logical consistency turns out to be a feature that matters to rational beliefs but not to rational changes of mind, other features of rational belief might always make a difference to rational changes of mind. So long as there is always some feature of rational belief that makes a difference to rational changes of mind, Static Rationalism can still be true. So we have arrived at the following very weak form of Static Rationalism.

Weak Weak Static Rationalism: Some of the conditions that determine when a belief is rational are always part of the explanation as to why a change of belief is rational (even if the resulting beliefs from a change of mind are not themselves rational.)

This is the version of Static Rationalism that seems to me most promising, and it is the thesis I am going to dedicate the remainder of this section to constructing and refuting. In order to come up with the best possible instance of this thesis, we need to find a feature that plays a necessary, constitutive role in determining the rationality of belief and always makes a difference to rational changes of mind. What might such a feature be? Let me look at a prime candidate, namely, what epistemologists have called “the basing relation.”

You might believe that a Bush is President of the U.S. on the basis of your belief as to who won the last election. Or, you might have been in a coma for the last fifteen years, wake up, watch the news, and believe that a Bush is President of the U.S. on the basis of your belief that we are at war in Iraq. This relationship of believing one thing on the basis of another is what epistemologists call the “basing relation.” The basing-relation plays an essential, even constitutive role in the theory of rational belief, for it is a widely accepted truism that what it is to be a rational belief is to be a well-based belief, a belief based on a good reason. (See [Pollock, 1986], [Korcz, 1997], [Kelly, 2002] for discussion and citations concerning the significance of the basing-relation.)

To illustrate the role that the basing-relation plays in the theory of rational belief, consider the following two philosophy professors.

Professor A. Goodcatch. Professor A. Goodcatch believes that 1. I am a good philosopher. 2. I have published many articles in the best philosophy journals. 3. I am a good philosophy teacher. 4. My students write in their evaluations that they love

me and think I am a genius. Professor A Goodcatch believes 1 on the basis of 2, and 3 on the basis of 4.

Professor A. Littlehoff. Professor A. Littlehoff believes all of the same things as his colleague A. Goodcatch. But he believes 1 on the basis of 4, and 3 on the basis of 2.

Professor A. Goodcatch and A. Littlehoff have the same beliefs, but different basing-relations between their beliefs. And because of this, there is a difference between them with respect to rationality. While we would not go so far as to say that A. Littlehoff is irrational, it is clear that Professor A. Goodcatch's beliefs are more rational than A. Littlehoff's. Professor A. Goodcatch's beliefs are well-based, he believes them on good bases, while A. Littlehoff believes them on bad bases, or at least bases that are not as good. The goodness and badness of your bases for your beliefs determine their rationality. In fact, the normative quality of your bases for your beliefs are constitutive of their rationality. To have a rational belief is to have a well-based belief, to have an irrational one is to have a poorly-based belief, and similarly for all gradations of rationality. Believing that you are a good philosopher based on what your undergraduates say about you isn't as reasonable as believing that you are a good philosopher on the basis of your publication record. Believing that you are a good philosophy teacher on the basis of your publication record isn't as reasonable as believing the same on the basis of what your students say about you.

Importantly, the basing-relation also plays a crucial role in determining rational belief-changes. Recall that we are looking for a feature of rational beliefs that always makes a difference to rational changes of mind.

The following cases are variations from [Tennant, 1994] and [Hansson, 1999] (See [Gillies, 2003] for discussion of these points). Suppose you believe that Cody is at the movies or the music shop (A or B). You also believe that Cody is at the music shop (B). Finally, you believe that there is no music shop in Santa Monica. You then learn that Cody is in Santa Monica. How should you change your mind?

The rational changes of belief depend not only on what you believe, but on the basing-relations that hold between your beliefs. If you believe that Cody is at the movies or the music shop based on your beliefs about Cody's general habits, and you believe that Cody is at the music shop based on what he said about needing a string for his guitar, then the reasonable way to change your beliefs is to give up the belief that Cody is at the music shop, and come to believe that he is at the movies.

However, if you believe that Cody is at the movies or the music shop on the basis of believing that he is at the music shop (A or B based on B), then you should give up both your belief that he is at the music shop, and your belief that he is either at the movies or the music shop.

So both your beliefs and the basing relations that hold between them play a role in determining the rational way for you to change your mind.

The role of basing-relations to rational belief-change is not just idiosyncratic to beliefs in disjunctions. Suppose you believe that that thing (pointing at something someone is holding) is a match. You also believe that that thing will light when struck on the side of a matchbox. Now you acquire some evidence that that thing is a toothpick disguised to look like a match. How should you change your mind?

Again, it depends not only on what you believe, but the basing-relations that hold between your beliefs. If you believe that that thing will light when struck on the basis of your believing that it is a match, then you should give up both the belief that it is a match and the belief that it will light when struck. But suppose you are at a magic show, and the magician just lit ten consecutive things on fire by striking them on the side of a matchbox (including a pencil, a chopstick, and so forth). If you believe that the thing will light when struck on the basis of such observations, then the reasonable way for you to change your mind would be to give up the belief that that thing is a match, but not give up the belief that the thing will light when struck.

We have seen that the basing-relation is normatively significant both to the rationality of belief and to the rationality of changes of belief. A well-based belief makes for a rational belief, and having a bad basis for a belief is sufficient to render that belief irrational. This explains the difference in rationality between Professor A. Goodcatch and Professor A. Littlehoff. The basing-relation also tells you which beliefs you should give up and which ones you should not. If belief A is based on belief B, then having justification to reject B amounts to having justification to reject A, as in the Cody Case and the Match Case.

Let us return to Weak Weak Static Rationalism. We have identified a feature that plays a normative role both for rational belief and rational changes of belief. We are looking for a way to explain the normative role that the basing-relation plays with respect to rational changes of belief in terms of its normative role with respect to rational beliefs. Weak Weak

Static Rationalism will succeed if such an explanation is available. Let us investigate.

1.1.4 The Irrelevance of Quality of Basing

One line of Static Rationalist thought seems promising. A Static Rationalist can state that you are always required to revise irrational beliefs, and irrational beliefs are beliefs that do not have good bases. Therefore, whenever you have given up your good bases for your belief that P, the belief is no longer rational, so you are required to revise it. In both the Cody and the Match Case, you start out with well-based, rational beliefs. You then acquire evidence against the good bases of your beliefs thereby undermining the rationality of those beliefs. In general, a Static Rationalist will require you to revise beliefs that are no longer rational. Thus, you are required to revise your now unjustified beliefs. This is an eminently good line of reasoning for the Static Rationalist. It is precisely the kind of explanation they need for Static Rationalism to succeed.

While it turns out that this explanation gets the Cody and the Match cases right, it turns out that it gets them right only accidentally. The goodness and badness of your bases for belief in fact make no difference to the rationality of belief-change in those cases. The way in which you should change your mind turns out to be insensitive to the goodness or badness of your bases for belief. This point becomes clear when we attend to a pair of cases, one in which you had a good basis for belief, the other a bad basis, where the two cases differ only with respect to those bases. If the rational revisions are identical, then the rationality of belief-change is insensitive to

the quality of your bases.

Consider Judgmental Patty versus Neurotic Patrick, who are on a first date. Judgmental Patty is looking for evidence that Patrick is a weirdo. Patrick has low self-esteem and really wants Patty to like him. After Patty takes a tour of Patrick's house, they have a short conversation. During this conversation they have the following beliefs.

Judgmental Patty believes that (1) Patrick is a weird dog-lover on the basis of (2) her belief that Patrick sleeps with a stuffed-three legged dog (stuffed in the sense of taxidermed). In the course of the conversation, Patty then realizes that the animal is actually a stuffed teddy-bear missing a leg.

The reasonable way for her to manage her opinions is to give up her belief that Patrick is a weird dog-lover.

Neurotic Patrick believes that (1) Patty thinks he is ugly on the basis of (2) his belief that Patty just said "Your hair needs a trim." But then he learns over the course of the conversation that Patty actually said "Your bear needs a limb."

The reasonable way for Patrick to change his mind is to give up his belief that Patty thinks he is ugly.

In Judgmental Patty's case, 2 is a good basis for 1. But in Neurotic Patrick's case, 2 is a bad basis for 1 (Patrick is unreasonably neurotic!) Nonetheless, the reasonable way for Patrick to change his mind is the same as the reasonable way for Patty to change her mind. They should both reject 1 upon being justified in rejecting 2. Patrick is no different from Patty with respect to how he should change his mind upon new information, even though Patrick has a bad basis for his belief, and Patty a good basis. So

the goodness and badness of bases plays no role in determining the rational changes of mind here. The fact that 1 is based on 2 alone determines the reasonable way for Patrick and Patty to change their minds. But basing alone is not supposed to make a difference to the rationality of a belief. *Goodness and badness* of basing makes that difference. So the notion of a rational belief does not explain the rational changes of belief in these cases.

Perhaps the relationship between good and bad basing is much more straightforward than we thought. Perhaps, in Patrick's case, it is because 2 is a bad basis for 1 that he would be unjustified in reasoning from 2 to 1. Similarly in Patty's case, it is because 2 is a good basis for 1 that she would be justified in reasoning from 2 to 1. In other words, maybe the explanation for why you are justified in reasoning from belief A to belief B is because A would be a good basis for B. Similarly for bad bases and unjustified reasoning.

For the Static Rationalist, this would be a very straightforward way to explain rational changes of belief in terms of rational beliefs. To summarize, the Static Rationalist would say:

The reduction of rational change to good basing: It is because a belief that A is a good basis for a belief that B that makes one justified to reason from A to B. (So a constituent of static rationality explains dynamic rationality).

Similarly the Static Rationalist could reduce irrational changes to bad basing.

The reduction of irrational change to bad basing: It is because a belief that A is a bad basis for a belief that B that

makes one unjustified to reason from A to B. (It is because bad basing is a sufficient condition for irrational belief that explains why reasoning from one belief to another is unreasonable).

According Weak Weak Static Rationalism, something's being a good or bad basis determines what is justified and unjustified reasoning. Thus, a central feature of rational and irrational belief determines justified and unjustified reasoning.

I am going to argue that the Static Rationalist has reversed the order of explanation. It is much more natural to explain what basing is in terms of reasoning from one belief to another, and good and bad basing in terms of rational and irrational reasoning. But my thesis is even stronger than that. I think that the Static Rationalist explanation of rational belief change in terms of good and bad basing degenerates into incoherence and self-defeat.

1.1.5 The Basing-Dilemma

Let us look past beliefs for a moment and consider doxastic states in general. In addition to believing and disbelieving things, we also sometimes suspend judgment as to whether some proposition is true or false. And sometimes these suspensions of judgments can be rational or irrational. It is what the Static Rationalist needs to say about these cases that I think undermines their position. Consider the following kinds of cases.

We have the following two ethnic bigots, Deutsch O. Phobe Sr. and Deutsch O. Phobe Jr.

Senior: Deutsch O. Phobe Sr. is a thorough-going bigot. Once

he learns that someone is German, he immediately infers that they're cruel. So Deutsch O. Phobe starts with the belief that Schroeder is not cruel. And for good measure, believes that Schroeder is French, maybe because he saw Schroeder eating a croissant for breakfast. Now Deutsch O. Phobe learns that Schroeder is in fact German, and changes his mind to believe that Schroeder is cruel.

Junior: Deutsch O. Phobe Jr. has inherited only part of his father's bigotry. He starts out with exactly the same beliefs as his father. But once he has learned that Schroeder is German, he simply suspends judgment as to whether Schroeder is cruel.

Now I submit that both Deutsch O. Phobe and his son have engaged in an illegitimate, unreasonable piece of belief-change. There's a complicated explanation for its irrationality. The explanation is something to the effect that both Deutsch O. Phobe and his son have changed their mind in a way that manifest a doxastic disposition that is responsive to ethnic biases rather than evidence of cruelty. Whatever is wrong with Deutsch O. Phobe's reasoning should be the explanation for what is wrong with Jr's reasoning, and something like this explanation must be correct.

What is the Static Rationalist's explanation of the case? For the Static Rationalist, Deutsch O. Phobe's reasoning is unjustified because the belief that Schroeder is German is a bad basis for the belief that Schroeder is cruel. Only by way of this Static Rationalist explanatory bridge do we understand why the reasoning is irrational. But recall that whatever is wrong with Deutsch O. Phobe's reasoning should be the explanation for what is wrong

with Jr's reasoning too. So the Static Rationalist is committed to saying that Deutsch O. Phobe Jr's reasoning is bad because the belief that Schroeder is German is a bad basis for suspending judgment about Schroeder's cruelty.

This explanation, however, renders the Static Rationalist position unstable. Consider the case of Mr. I.M. Impartial.

Impartial: Mr. Impartial hasn't considered whether Schroeder is cruel or not. But he does believe that Schroeder is French. Then he learns that Schroeder is German, and changes his mind, concluding that Schroeder is German and not French. Still he hasn't considered whether Schroeder is cruel or not. Mr. I.M. Impartial is currently watching a World War II era film with the stereotypical cruel German villain, and it occurs to Mr. Impartial to think about the Germans that he knows. He remembers that he believes that Schroeder is German, but concludes that he doesn't know whether Schroeder is cruel or not. That is, he suspends judgment as to Schroeder's cruelty.

Notice that Mr. Impartial ends up with the same doxastic state of mind as Deutsch O. Phobe Jr. Both believe that Schroeder is German, both suspend judgment as to whether Schroeder is cruel. But Mr. Impartial, unlike Deutsch O Phobe Jr. has reasoned impeccably, rationally, justifiably. There is nothing wrong with the way in which Mr. Impartial has managed his opinions about Schroeder's nationality or his cruelty.

The case of Mr. Impartial raises a question. Does Mr. Impartial base his suspension of judgment that Schroeder is cruel on his belief that Schroeder

is German? This question poses a serious dilemma for the Static Rationalist. Either way the static rationalist answers this question will undermine his explanation of rational belief-change in terms of rational belief. If the Static Rationalist says yes, then since Mr. Impartial has reasoned justifiably, it would follow that the belief that Schroeder is German is a *good* basis for suspending judgment as to Schroeder's cruelty. But the Static Rationalist is already committed to it being a bad basis! Why? Because we've already established that Deutsch O. Phobe Jr. is unreasonable in the way he reasons from the information that Schroeder is German, to his suspension of judgment as to Schroeder's cruelty.

On the other hand, the Static Rationalist might claim that Mr. Impartial did not base his suspension of judgment as to Schroeder cruelty on his belief that Schroeder is German. In that case, the Static Rationalist has some more explaining to do. What's the difference between Deutsch O. Phobe Jr. and Mr. Impartial such that Jr. *did* base his suspension of judgment on his belief but Mr. Impartial *did not*? It seems that the only answer is that the line of reasoning that led to their beliefs was different. Jr. changed his mind *from believing that Schroeder was not cruel*. He changed his mind as a result of changing his mind from believing that Schroeder is French to believing that Schroeder is German. *Mr. Impartial did not change his mind in such a way*. He never reasoned from the belief the Schroeder is not cruel to his suspension of judgment. So the difference between Deutsch O. Phobe Jr. and Mr. Impartial is simply a difference in inferential history, a difference in the way in which they reasoned to their current beliefs.

But notice that this sort of explanation of why Jr. based his suspension of

judgment on his belief, whereas Mr. Impartial did not, undermines the Static Rationalist position. The way you change your mind, or the way that you changed you mind in the past, determines what is based on what. What it is to base one doxastic state on another is determined by how one reasoned from one doxastic state to another. So reasoning determines basing. But if reasoning determines basing, then we should most naturally understand good and bad reasoning to determine good and bad basing. The natural understanding of the normative quality of basing derives from the normative quality of the reasoning that determine it. Thus, there would be a Dynamicist explanation of basing and good and bad basing. To resist this line of thought, the Static Rationalist will be forced into the intractable position of requiring a dynamic explanation of what basing is, but refuse a Dynamicist explanation of the goodness and badness of basing. This Basing-Dilemma pushes the Static Rationalist to a Dynamicist view, or no view at all.

Objections, Replies, and Clarifications

One might object to the set-up of the Basing-Dilemma by claiming that in the case I have described, both Sr. and Jr. have tacit beliefs to the effect that Germans are cruel or likely to be cruel. Thus, we only think that their reasoning is unjustified because they have such unjustified, unreasonable beliefs, i.e., their belief that Germans are cruel. Thus, the irrationality of their reasoning is explained in terms of the poor bases they have for their racist beliefs about Germans. We trace their irrational piece of reasoning back to an irrational belief, contrary to Dynamicism.

The objection here focuses on an inessential detail of the particular case,

namely, that both Jr. and Sr. have general dispositions to infer from someone's being German to their cruelty (or in Jr's case, to reason from Germanness to withholding judgment about cruelty). I agree that it is immensely plausible that people who have such dispositions thereby reveal to us that they have beliefs, perhaps unjustified ones, that Germans are cruel. One will find this view plausible if one held a dispositional account of belief; if a belief just is a disposition to behave or infer in a certain way, then of course reasoning or inferring in such a way amounts to a belief. I want to reject both the dispositional account of belief, and the strategy of postulating implicit beliefs in general to rationalize or make sense of inferential patterns. But for the purposes here, it is not necessary to engage in any dispute over beliefs and dispositions. The general disposition to reason from the belief that a person is German to the belief that they are cruel is an eliminable feature of the case. My point would hold even in cases of people who, lacking such general dispositions, just so happen to engage in a one-off piece of reasoning like that of Sr. and Jr. Even a person who just so happens to reason like Sr. and Jr. in a single case-but has never and will never again engage in such reasoning-has changed their mind in an unreasonable way. People can still engage in unjustified reasoning occasionally even if they lack the disposition to do so generally. In cases of one-off pieces of reasoning, there is no temptation to ascribe to agents anything like an implicit belief. Thus, the dilemma posed by the basing-problem for Static Rationalism does not depend on the fact that Jr. and Sr. have general dispositions to infer in one way or another. The Basing-Dilemma still holds when formulated as a case of one-off pieces of reasoning.

One point of clarification is in order regarding the basing-relation. Some of the time, our beliefs have bases. Some of the time, they might not. The same holds of our suspensions of judgments, disbeliefs, or degrees of belief. The Basing-Dilemma and the other considerations against Static Rationalism do not presuppose that we always have a basis for our doxastic states of mind. The problems hold for anyone who considers that in some cases, the bases of our beliefs have some role in making them rational. While there are some who deny this (for instance, Minimal Static Rationalists as I discuss below, and perhaps a radical Dynamicist who denies that beliefs are rational at all), many epistemologists agree that having evidence or a coherent body of beliefs is insufficient to render a belief reasonable. One must also base that belief on such evidence or body of beliefs [Pollock, 1986], [Korcz, 1997], [Kelly, 2002]. If we accept this basing-constraint and agree with the intuitions in the above counterexamples, Weak Weak Static Rationalism is in trouble. And by implication, so is Static Rationalism.

1.2 Summary

Let us take stock of where we have ended up so far. With our discussion of the Preface Paradox, we have ruled out that changes of belief need to end up statically rational in order to be rational. With our discussion of the basing-relation, we have seen that a crucial feature of static rationality cannot be used to explain the conditions of dynamic rationality. In fact, if anything, this crucial feature must be explained dynamically. Thus, we have reason to conclude that the rationality of changes of belief is at least as fundamental

as the rationality of belief. The Static Rationalist attempt to reduce dynamic rationality to static rationality fails. In fact, we have seen evidence that the rationality of belief is explanatorily posterior to the rationality of changes of belief. Thus Static Rationalism, the idea that static rationality is primary and fundamental, is false, and dynamic rationality is at least as basic. The argument cannot end there, however. Static Rationalism does not fail simply because we have a series of Static Rationalist proposals that run afoul of clever counterexamples. In the next two chapters, I will show that the counterexamples have a unified structure. We find the counterexamples compelling precisely because our judgments of epistemic rationality fundamentally track changes and only changes of belief. My argument will depend on a Dynamicist analysis of both the Basing-Relation and the Preface Paradox.

Chapter 2

Against Static Rationality

Part I

In Chapter One, we found a counterexample to every plausible proposal reducing dynamic rationality to static rationality. The counterexamples warrant us to conclude that dynamic rationality is at least as basic as static rationality. In the next two chapters, I will argue that some of the counterexamples in Chapter One point to an even stronger conclusion; our judgments of epistemic rationality fundamentally track only changes of mind and not states of mind at all. There is one unifying explanation of our intuitions behind the counterexamples to Static Rationalism. That explanation is Dynamicism.

My argument in this chapter will depend on an argument that the Preface Paradox can in fact be generalized to show that there cannot be any rational failures without a failure of dynamic rationality. My argument in the next chapter is that the normative significance of the basing-relation has a Dy-

dynamicist explanation. (1) and (2) provide evidence for the Dynamicist thesis that only failures of dynamic rationality are failures of epistemic rationality, and also instantiate a general schema for establishing that thesis.

2.1 Minimal Static Constraints

We have established that the rationality of belief-change is basic and cannot be analyzed in terms of the rationality of belief. However, we have not ruled out that there are independent conditions for the rationality of belief that are irreducibly static. In our current dialectical situation, we are left with Binary Rationalism, or the view that both static and dynamic rationality are basic. Binary Rationalism differs from Dynamicism only with respect to the claim that there must be wholly static conditions on rational belief that are independent of the conditions for rational changes of belief. My aim, then, is to argue that, on the contrary, there are no such conditions.

Epistemology is replete with proposals for conditions on static rationality. Some classic proposals include the condition that all current beliefs must be *coherent* in some technical sense [Bonjour, 1985], that they must be reliably formed [Goldman, 1979], that they are default rational absent defeaters [Harman, 1986], and so forth. One almost universal starting point for theories of static rationality is *logical* coherence of some sort or other. Logical coherence has been understood to mean bare consistency, consistency and closure, or probabilistic consistency. Let us call Minimal Binary Rationalism any view of the rationality of belief that takes logical consistency to be a constraint on static rationality. (Stronger constraints, like logical closure,

will be ignored for my purposes).

As we have seen, the Preface Paradox already poses a problem for those who aim to reduce rational belief-change to rational belief. Beliefs can end up inconsistent while still being the product of rational changes. But even more, we do not think people with Preface-type beliefs are guilty of any rational failure whatsoever, even static ones. David Christensen, in his book *Putting Logic in its Place* [Christensen, 2005], contains an extended defense of the Preface Paradox [Makinson, 1965] as a decisive counterexample altogether for logical coherence as a constraint on rational belief. Not only do we fail to find Preface-type agents guilty of irrationality, they seem rationally superior to evidentially-identical counterparts who either fail to adopt the Preface-type beliefs, or instead give up their beliefs in order to restore consistency. It is better to be inconsistent than consistent when one is in the Preface situation. No theory that requires logical consistency of rational beliefs can explain this. Thus, logical consistency cannot be a necessary condition of statically rational all-or-nothing beliefs.

Gilbert Harman has been most insistent in arguing against straightforward logical constraints on all-or-nothing belief. Harman notes that many inconsistencies are difficult, perhaps even humanly impossible, to detect. For example, if arithmetic were inconsistent, as it might be, very few mathematicians who accept the axioms of arithmetic can be charged with any kind of *rational* failure. Furthermore, some inconsistencies, even upon detection, can be rationally held. Harman illustrates this point with examples of philosophical paradoxes. Paradoxes like the Liar and Zeno's paradox have eminently believable premises and unbelievable, even inconsistent conclusions. Yet,

Harman argues, the evidence we have for the premises, and our inability to independently reason ourselves out of any particular premise, might justify the retention of all them, even when we simultaneously reject their conclusions. In my opinion, Harman overstates his case. It might be that in light of paradoxical arguments, where we believe the premises and deny the conclusions, yet acknowledge that the arguments are valid, we are required to suspend judgments on all of the premises until we find special reason to favor some premises over others. However, it is sufficient for Harman's point that those who maintain all of their inconsistent beliefs in light of paradoxical arguments, while lacking any evidence that favors some premise over another, are *more* rational than those who choose to reject some particular premise over another. This latter kind of agent is consistent, but clearly less rational than the former. And from this it follows that consistency is not a necessary condition for the rationality of beliefs at a time.

If logical consistency is not a constraint on statically rational all-or-nothing belief, then does it follow that there are no constraints on static rationality at all? We have seen that logical consistency seems to be a minimal condition; absent that, what other conditions could there be?

2.2 Rational Degrees of Belief

One standard response to the Preface Paradox and various objections to propositional consistency constraints on all-or-nothing belief is to “go probabilistic.” These Probabilistic Static Rationalists maintain that there are formal constraints on static doxastic states, but only on degrees of belief. The

other response to the Preface is to abandon formal constraints completely for other static constraints. This section will argue against the first strategy, the next section will argue against the second.

Probabilistic Static Rationalists (or Probabilists), who accept that logical consistency does not rationally constrain all-or-nothing belief generally conclude that the concept of all-or-nothing belief must be flawed, or that there is no direct connection between logic and rational all-or-nothing belief [Kaplan, 1996]. Instead, they maintain that there are only formal constraints on your degrees of belief. According to Christensen's view, for instance, my degrees of belief need to be statically coherent in order to be epistemically rational.¹ Static Coherence is a matter of one's degrees of belief satisfying the Kolmogorov axioms of probability. Informally, your degree of belief on any proposition must be greater than or equal to zero. Your degree of belief in a tautology is one. Your degree of belief in a mutually exclusive disjunction is the sum of your degrees of belief in each disjunct. Anyone whose degrees of belief fail to abide by any principle that follows logically from the axioms will be statically incoherent.

Static Coherence of one's degrees of belief requires that one must be fully opinionated, fully believe every tautology, fully deny every contradiction, and be mathematically precise in all of one's degrees of confidence. Legitimate complaints have been filed against all three of these requirements ([Levi, 1987], [Maher, 1992], [vanFraassen, 1995]). I will address one of these complaints in the next chapter. However, familiar complaints aside,

¹ [Horwich, 1982], [Howson and Urbach, 1993], [Kaplan, 1996]. Compare with the work of Skyrms, ex. [Skyrms, 1987] and Jeffrey, ex. [Jeffrey, 1988, Jeffrey, 1965] who characterize probabilistic consistency as a constraint on practical or pragmatic rationality.

Probabilistic Static Rationalists take it to be a benefit of their view that they can escape the original Preface paradox. Probabilists explain the original Preface case as a case in which agents simply have high degrees of belief in each proposition in the body of their book while having a low degree of belief in their conjunction. Not only is this state of mind perfectly coherent, it is indeed mandated by the axioms of probability theory (See [Glymour, 1980], who uses this point as a complaint against probabilism). According to the Probabilist, the reason we do not judge Preface-type beliefs as unreasonable is because they are not; they are perfectly probabilistically coherent, and only probabilistic coherence is required for static rationality. Probabilism not only fails to generate the original Preface-type paradox, it explains our judgments with regards to the paradox.

Is Static Coherence subject to the same kinds of objections that render logical coherence problematic as a constraint on all-or-nothing belief? Is there a generalization of the original Preface Paradox that poses the same problem for static coherence as logical consistency? I believe that, unfortunately, the ability of Probabilistic Static Rationalism to give a solution to one Preface-type situation fails to immunize it from Preface-type considerations generally. Preface-type problems are simply instances of a more general kind of phenomena in which we can have strong evidence for each of many propositions, while also having evidence of our own fallibility, from which it follows that we should believe we are wrong about one of our beliefs. Such problems apply to degrees of belief as much as to all-or-nothing belief. The next two sections will illustrate how evidence of our own fallibility can lead to reasonable, but statically incoherent, degrees of belief.

2.2.1 A Simple Strengthened Preface Case

There are many different domains in which I consider myself fallible. Any domain in which I have a history of error should count. Of course, this leaves open the possibility that I should consider myself fallible in a great many more domains, including those in which I have been lucky enough never to err. Unlucky enough for me, I have a history of error with respect to some of my logical beliefs. It's not that I have been wrong about whether Tracy is a dog or Tracy is not a dog, or that I have been shown that some contradiction is true. Rather, I wrongly believed in the past that a certain proposition was a tautology, when in fact it was not, or that a proposition followed from another, when in fact it did not. These are still logical beliefs, albeit they are beliefs about the logical forms and properties of propositions. Given this history of error, I should currently recognize the likelihood of error with respect to some of my current logical beliefs.

For instance, what degree of belief should I have in the proposition that its not the case that (if Bill or Joe went to the store, then Jack or Phil didn't unless Mary went, and Bill or Joe went to the store and Jack and Phil did too but Mary didn't)? Let us call this proposition P. It seems to me now that P has the logical form of a tautology. Thus, according to static coherence, I should be absolutely certain of it. However, I have been mistaken about the logical form of propositions in the past. Perhaps I was too confident in my assertion that P is a tautology. Many people just as smart if not smarter than I have been mistaken about such matters in the past also. I recognize that I am not sure as to the precise logical form of some of the propositions I believe. Could any of them be contradictory? Might some of them be

tautologous though I do not recognize them as such? Even if we grant that logical knowledge is more certain than empirical knowledge, that logical facts are all knowable a priori, and that logical errors can always be corrected a priori, it remains true that our history of logical error can and should affect the attitudes we have toward a putative logical fact.

Here is one line of reasoning that appears to me perfectly reasonable. My history of error with respect to my beliefs about the logical form of propositions pressures me to be less than perfectly certain of the tautological form of proposition P. Given that I am less than perfectly certain that P has a tautological form, I should revise down my degree of belief in P itself from one to something slightly less than one. In other words, the following seems to be a rationally permissible principle to follow: My degree of belief on P, given that I am less than perfectly certain that P is a tautology, is less than 1. I update my opinions following such a principle in the usual Bayesian way. Now, as a matter of fact, P is a tautology. I have ended up statically incoherent.

The case of tautologies is in no way special in illustrating our fallibility with respect to logical facts. We have all erred, and can be systematically prone to err, about whether a proposition is contradictory, whether it follows from one or many other propositions, whether it is consistent with another, and so forth. We believe and have degrees of belief in a great deal of propositions, many of which are logically related to each other. But I can also have beliefs and degrees of belief about how they are logically related, and here such degrees of belief must respect my history of error. Such a history of error constitutes evidence I should respect even if, as a matter of fact, *I am*

not currently making any logical errors whatsoever. I might have all true and no false beliefs about the logical facts concerning the propositions I believe. My history of logical error itself gives me sufficient evidence that I should not be perfectly certain regarding the logical facts of the propositions I believe. This imperfect certainty then affects the rationality of my first-order degrees of belief in those very propositions.

If the previous lines of reasoning are correct, then imperfect certainty about the logical facts about any proposition P can constrain my degrees of belief that P. Imperfect certainty about the tautological form of P constrains my degree of belief in P. Imperfect certainty about the contradictory form of some other proposition, Q, constrains my degree of belief in Q. Imperfect certainty that P follows from Q constrains my degrees of belief in the conjunction of P and Q. How? Some of the relevant principles relating our higher-order degrees of belief about logic to our first-order degrees of belief might be as follows:

Tautology: $pr(P|pr(P \text{ is a tautology} < 1)) < 1$

Contradiction: $pr(P|pr(P \text{ is a contradiction} < 1)) > 0$

Consequence: For $pr(Q) > 0$, $pr(P \& Q|pr(P \text{ follows from } Q) < 1) < pr(Q)$

In ordinary English, Tautology states that if I am less than perfectly certain that P is a tautology, I should be less than perfectly certain that P. Contradiction states that if I am less than perfectly certain that P is a contradiction, I should have some degree of belief in P. Consequence states

that if I am less than perfectly certain that P follows from Q, then my degree of belief in P & Q should be less than my degree of belief in Q. Such principles, when combined with the thesis that one must be less than perfectly certain of a logical fact in order to respect evidence of one's logical fallibility, imply that agents can be reasonable but statically incoherent. In cases in which a certain proposition P in fact has a logical property that I am less than certain it has, adhering to Tautology, Contradiction, and Consequence can lead me to be less than perfectly certain of a tautology, have some degree of confidence in a contradiction, or lead me to believe a logical consequent of Q with less confidence than I believe Q.

We do not even need to interpret the principles as being rational requirements to yield the required result. Even if I am simply rationally permitted to conform to such principles (if we switch the "should"s to "can"s), I can rationally have statically incoherent degrees of belief. Static coherence requires that our degrees of belief in a proposition P respect what are in fact the logical properties of P. In effect, static coherence requires my unconditional degrees of belief to be *as though* I were absolutely certain and truly believed all of the logical facts about the propositions I believe. But in any instance in which I am less than perfectly certain about such logical facts, it can be rationally permissible that I fail to have perfectly coherent degrees of belief.

Two points of clarification are in order here. First, I am not presupposing or arguing that static coherence requires agents to have absolutely certain degrees of belief in all of the true logical facts concerning the propositions they believe. Agents can very well have no beliefs whatsoever about such logical

facts. Rather, I am arguing that, *if* an agent has such degrees of belief, static coherence requires that they fail to respect a certain kind of evidence they can acquire, namely, evidence of their own logical fallibility. Respecting such evidence can be incompatible with maintaining static coherence. Second, the objection is simply not the familiar objection that humans in their limited capacity cannot be logically perfect in real life, and can therefore be rationally incoherent. An agent can fully respect logic in the way she reasons, can fully have all of the true beliefs (and no false ones) about the logical forms of the propositions she believes, and still, due to her history of logical error, be reasonably less than perfectly certain of her degrees of belief in the logical facts. Even an agent who *can* be logically perfect can simulatenously have evidence that they are not. Such evidence is sufficient to make reasonable certain statically incoherent degrees of belief.

Notice also that in these simple cases, the relevant intuitions about rationality parallel the original Preface case. An agent who has evidence of past error with respect to the logical facts, but nonetheless remained perfectly certain that a proposition P had logical property ϕ , would clearly be *less reasonable* than an agent who became less than perfectly certain that P has ϕ . Similarly, an agent who remained statically coherent while violating Tautology, Contradiction, and Consequence would be less reasonable than one who respected such evidence and revised her opinion according to those principles. These judgments suggest that the principles are genuine requirements on rationality. They also suggest that rationality and static coherence simply come apart. We can still accept that the Kolmogorov axioms give us the correct axiomatization of the concept of probability and give us

coherence-conditions for degrees of belief. An agent is certainly in an *illogical* state of mind when obeying Tautology, Contradiction, and Consequence at the expense of static coherence, for they are vulnerable to a static Dutch book. They are, however not in an *irrational* state of mind. If anything, they are rationally impeccable.

2.2.2 A Complicated Strengthened Preface

I believe that there can even be a Strengthened Preface Case that does not rely solely on our history of logical error, one that simply exploits our fallibility with respect to getting right a large number of propositions at one time. This type of Strengthened Preface would perfectly parallel the original Preface paradox, and would therefore be on sounder footing. I will attempt to construct this more complicated version of a Strengthened Preface in this section. Readers have previously noted that this section is controversial; I have tried to respond to the relevant objections below.

I believe it is possible that, with respect to our conditional degrees of belief, we can very reasonably take any two propositions in a series of propositions to be pairwise independent, but consider a proposition to be *dependent* on the truth of the conjunction of every other proposition in the series. Such a situation arises because we can have evidence of our own fallibility that leads us to a high degree of belief in a higher-order proposition, namely, that one proposition in a set of propositions is false. Such a situation would generate a statically incoherent set of degrees of belief. I believe that respecting evidence of one's fallibility can lead to such a situation, and that agents in such situations, like the original Preface Paradox, can be fully reasonable.

Suppose I am weighing 1000 cinder blocks individually on a scale. Trial runs on the scale reveal that it does not get any more or less accurate with repeated weighings, so the accuracy with respect to weighing cinder block n does not render the scale more or less accurate with respect to weighing cinder block $n+1$. Thus, from the trial runs, I have good reason to think each weighing is independent. The cinder blocks themselves are also completely random. I am required to number all of the cinder blocks and place them into a ten-pound pile and a non-ten pound pile. After finishing my job, I am required to list in my report the cinder blocks that weigh ten pounds, in case the piles get mixed up. My report is thus a list of the form.

Cinder block $\#i$ weighs ten pounds.

Supposing there are $\#n$ cinder blocks that I weigh to be ten pounds, let P_1 through P_n be the names of the propositions expressed by each sentence in the report.

For simplicity sake, let us assume that my unconditional degrees of belief in each of these propositions are uniform, say, .9. Therefore by their independence, my conditional degrees of belief for each pair of propositions is equal to my unconditional probabilities for those propositions. That is, for all P, Q , $pr(P|Q) = pr(P) = pr(Q) = .9$. The assumptions about uniformity and the precise numbers are inessential. The assumptions about independence is crucial.

It is consistent with evidence of the independence of each pair of proposition that I am very likely wrong about one of the many propositions in the report. So suppose, as in the original preface case, past error and the

errors of others, along with my background beliefs that I am in no better epistemic position than my past self or others, I come to believe that I am wrong about at least one of the many propositions in the report. Given this belief, how do I judge the conditional probability of some proposition, say, P_{17} , in my report given the conjunction of all of the other propositions? I reason that, since I am probably wrong about one of the propositions, if the conjunction of every proposition but P_{17} were true, then P_{17} is false. Thus, my $pr(P_{17}|P_1, \dots, P_{16}, P_{18}, \dots, P_n)$ is very low.

Such a probability will generate an inconsistency with any plausible assignment of unconditional probabilities to the propositions in my report. By uniformity and independence, for all n, m , $pr(P_m|P_n) = P_m = P_n$. By the Kolmogorov axioms, this fact implies that the probability of P_m given the conjunction of every other proposition in the report is equal to P_m . (To see this, it is sufficient to see that if P is independent of two conjuncts that are independent of each other, then P is independent of their conjunction.) Thus, my conditional degrees of belief $pr(P_{17}|P_1, \dots, P_{16}, P_{18}, \dots, P_n)$ is inconsistent with my unconditional degrees of belief for P_{17} and every conditional degree of belief paired with P_{17} , $P_{17}|P_n$ for all n .

The argument generalizes. For all of the propositions P in my report, my degrees of belief on P conditional on the conjunction of all the other propositions should be low, since I have the preface belief. By the Kolmogorov axioms, this degree of belief should also equal my unconditional degree of belief on P because of my pairwise conditional degrees of belief are equal to the unconditional degrees of belief. Any degree of belief lower than $pr(P_n)$ will generate the inconsistency.

I believe there are two plausible and compatible explanations of this Strengthened Preface Case. The first is that it is simply a straightforward probabilistic application of the same kind of reasoning we find in the original Preface case. My evidence sufficiently justifies me in taking the propositions in my report to be pairwise independent, therefore requiring that $pr(P|Q) = pr(P)$ for all P,Q. That cinder block #1 weighs ten pounds given that block #2 weighs ten pounds should be as likely as that 1 weighs ten pounds. But other evidence, the preface-belief, requires my degree of belief in a proposition to be dependent on the truth of the conjunction of all the others, therefore requiring my conditional degree of belief that P given the conjunction of all the other propositions to be very low. This latter bit of evidence does not defeat my original belief in pairwise independence. Evidence of pairwise independence does not defeat the evidence of my epistemic fallibility. With respect to individual, pairwise, or small numbers of beliefs, my higher-order preface belief does not require me to give up my commitment to these beliefs. But with respect to a large set of beliefs (where it is highly vague how large is large enough), my preface belief does change how I must commit to my entire set of beliefs.

There is one other suggestive explanation of our intuitions in this Strengthened Preface. This Strengthened Preface might be exploiting a distinction between two different kinds of statistical dependence and independence. Independence requires that your conditional probabilities for each pair of propositions is equal to your unconditional probabilities for those propositions. If this is not the case, your degrees of belief are manifesting your perception that two propositions are dependent. But there are at least two different

kinds of dependence between P and Q that your degrees of belief might be manifesting. The first is a kind of metaphysical or ontological dependence. It might be that you believe that whenever Q, P, (or not P) or generally whenever Q, P, (or not P). Such dependencies might be explained by causal relations, common causes, nomic laws, or other kinds of real relations between events or things you take to exist. But I might also believe that I tend to be wrong, or right, about P whenever Q. I might believe this even if I know there is no ontological dependence between Q and P. My conditional and unconditional degrees of belief on P and Q do not themselves determine what kind of dependence or independence it is that holds between P and Q. Such dependencies can come apart. I can consistently believe an ontological dependence to obtain where an epistemological one does not, and vice versa. Thus, my conditional and unconditional degrees of belief can be torn both ways from such consistent background beliefs. When ontological dependencies and epistemological dependencies come apart in a reasonable way, I can end up with first-order degrees of belief that are statically incoherent, and such degrees of belief do not manifest any kind of irrationality. In the Strengthened Preface Case, my conditional degrees of belief in each pair of propositions manifest their ontological independence, but my degree of belief in each proposition conditional on the conjunction of all the others manifest an epistemological dependence that stems from my higher-order preface belief.

2.3 Objections

As many readers have been insistent on pointing out, any attempt to show that one can be rationally inconsistent will be liable to a charge of fallacious reasoning. Somewhere along the line, anyone who reasons and ends up probabilistically inconsistent can be charged with committing a fallacy, at exactly the point in which their opinions become inconsistent. The open question is whether, at that particular point in their piece of probabilistic reasoning, the logical error is also a normative error. Many readers of this section have insisted that the argument fails precisely because there is in fact a normative error. In effect, they deny the intuition I attempt to elicit in the case I describe.

To one kind of reader of this section, nothing I say will be satisfactory. This is the reader who insists that the normative error arises because and only because of the logical error. Anyone who insists, come what may, that there are static probabilistic constraints on degrees of belief at a time can easily deny any attempt to show otherwise, simply by invoking static probabilistic constraints. I do not have much in the way of a reply except that I hope such a reader will find the first, simpler Strengthened Preface paradox more convincing.

More interesting is someone open to the possibility that there are intuitive counterexamples to static probabilistic rationalism, but insists that at the basic, intuitive, pre-theoretical level, this more complicated Strengthened Preface case I have described is not such a counterexample. Thus, the objection goes, the agent I have described really is guilty of a basic, intuitive failure of rationality. Because Static Probabilistic Rationalism explains such

a failure, I have failed to give any counterexample to such a view.

On behalf of my opponent, let me construct two arguments showing how I might have failed in this Strengthened Preface Case to prove the possibility of rational first-order inconsistent degrees of belief. (Variations on such cases were given to me by Gideon Rosen, Tom Kelly, and Jim Pryor).

Argument #1: Imagine someone who engages in the following line of reasoning; he starts off believing that he is likely wrong about one of the propositions in a set, he places a high probability in each proposition in the whole set, and he takes all of the propositions to be independent of each other. He begins to acquire absolutely conclusive evidence that the first proposition is true, then the second, then the third, then the fourth, and so on up until the last one. At this later point in time, how should he manage his degrees of belief? Clearly, this person's unconditional degree of belief in the last proposition should be exactly as it was in the beginning. However, such a person should revise his original high degree of belief that he is likely wrong about one of the propositions in the set. As this person acquires conclusive evidence that each proposition is true, he acquires evidence against his belief that he is likely wrong about one of the propositions in the set. The fallibilist line he takes toward his beliefs before he engages in the series of changes of belief gets defeated as the truth comes in.

Moreover, such a person can easily foresee everything I described in the previous paragraph as one possible future state of his degrees of belief. Before this hypothetical person acquires any further evidence that conclusively tells him, one at a time, that the propositions in the set are true, he can easily foresee that, *were* he to acquire the evidence of the relevant type, the

reasonable way for him to update his degrees of belief is exactly as I described above. Since this person can foresee all of this, his initial conditional degree of belief that, say, proposition 17 is true, given that all of the other ones are true, should equal precisely his unconditional degree of belief that proposition 17 is true. This should be his conditional degree of belief precisely because he can foresee that, in all future reasoning in which he can foresee acquiring decisive evidence that every other proposition is true, he ought to revise his preface-type, fallibilist degree of belief that one of the propositions is false, and therefore by Independence, retain his original degree of belief in P_{17} . If he foresees all of this as the right way to reason given the acquisition of decisive evidence for each proposition, then his conditional degrees of belief ought to reflect his foresight, for what is a conditional degree of belief but foresight into how one ought to reason upon learning a proposition?

Argument #1, if true, contradicts the complicated Strengthened Preface Case, and points out exactly where the fallacy occurs. The fallacy occurs because the Preface-type belief cannot be rationally held to a high degree of belief simultaneously with Pairwise Independence and the low conditional degree of belief in P_{17} given the conjunction of all the other propositions. Your conditional degree of belief on P_{17} , given the large conjunction, ought to reflect how you foresee you will change your degree of belief on P_{17} when you acquire decisive evidence of the large conjunction. Given Pairwise Independence, you should foresee that your future degree of belief on $pr(P_{17})$ in this situation will be your current degree of belief on $pr(P_{17})$.

Argument #1 creates a special difficulty for me because I embrace the conclusion that you should revise your fallibilist or preface-type degree of

belief upon acquiring cumulative conclusive evidence for all but one of the propositions you believe in a set of propositions. In effect, I agree with the claims that Argument #1 make about the appropriate ways for us to reason. Before I respond to Argument #1, let me construct another case that shows how my complicated Strengthened Preface fails.

Argument #2: Imagine that yesterday there were a hundred flips of a two-sided coin. I am told that the coin is perfectly symmetrical with respect to weight, that each flip was done by a different person chosen from a room of a thousand people at random, and that the results of the flips were recorded by an impartial observer. With all of this information, I justifiably have a high degree of belief that the coin is fair, that at least one of the flips landed tails (Preface), that the probability that the coin landed heads on flip n given that it landed heads on flip m is equal to the probability that it landed heads on flip n (Pairwise Independence).

In this case, my initial degrees of belief, conditional and unconditional, about the probability of heads for each flip should clearly be probabilistically coherent. Nothing weird is going on in this case, just merely a textbook example of probability theory meant to illustrate how the Kolmogorov axioms capture the consistency constraints on degrees of belief. But what is the difference between this case and the Strengthened Preface case? Make the number of flips arbitrarily large. It will make no normative difference to the fact that this is a textbook example of how your degrees of belief in heads, and all your other degrees of belief, ought to be Kolmogorov-consistent. The challenge, then, is to show that the complicated Strengthened Preface is more than a mere re-description of a textbook argument *for* probabilistic

consistency as a constraint on rationality. If I cannot meet this challenge, the complicated Strengthened Preface clearly fails.

2.3.1 Replies

I have stated above the extent to which I agree with advocates of Argument #1. I believe they are stating the correct results regarding how we ought to reason in a hypothetical situation. I also believe that we can reasonably foresee such situations, and that such foresight might be reflected in my beliefs and degrees of belief now. However, proponents of Argument #1 and I seriously disagree about *the nature of conditional probabilities*. For me, a conditional judgment is basic, a kind of opinion one has that is irreducible to other opinions you have.²

My conditional probability on P given Q is not reducible to my current opinions about how I should update a certain unconditional degree of belief P given the acquisition of decisive evidence of Q. My current conditional judgment that P is the case, given that Q is the case, is one kind of judgment. It is the likeliness I give to P assuming that Q is in fact the case. It is *not* the judgment about what my future unconditional probabilities on P will be or should be on the assumption that I will acquire decisive evidence for Q, or learn that Q. This latter kind of judgment is also a conditional judgment, but a very different kind of conditional judgment. It is a conditional judgment on

²We know that certain reductions of conditional probabilities do not work. For instance, from the results of David Lewis [Lewis, 1986], we know that a conditional probability is not an unconditional probability on any truth-conditional indicative conditional. We know from Alan Hajek [Hajek, 2003] that a conditional degree of belief is not analyzable as a ratio of two of your unconditional degrees of belief (What exactly is a ratio of two of your mental states anyway?).

a proposition very different from P, given a proposition very different from Q. When I make the first kind of judgment, I assume that Q is the case, and make a judgment about how likely something else is the case. When I make the second kind of judgment, I get to assume that I will get decisive evidence that Q is the case and learn that Q, and I ask how I would or should reason about P given that assumption. The second kind of judgment is much richer than the first in a variety of ways. In simply assuming Q to be the case, I am not ipso facto entitled to assume anything about myself, that I would also learn or get decisive evidence that Q is the case. I am not ipso facto entitled to assume that I will in any way even learn anything about whether Q. On the other hand, if I am asked to assume I will learn or get decisive evidence that Q is the case, I am making a much richer assumption, an assumption that entitles me to make use of my beliefs, if I have any, about the nature of such evidence, or the way in which I would acquire it, and so forth. Simply assuming that Q is the case does not involve, and ought not involve, such a richer assumption.

Consider a case, independent of any issue concerning the Strengthened Preface, in which my conditional probabilities intuitively come apart from the judgment about how I ought to reason given decisive evidence. Annette and Betty have a mutual friend, Claire who hates telling Annette anything personal, and Annette knows this. So Annette gets all of her Claire-gossip from Betty. But Betty only tells Annette something about Claire when the gossip is particularly juicy. Annette currently has a degree of belief that Claire is dating Fred, a degree of belief that Claire is pregnant, and a conditional degree of belief that Claire is pregnant, on the assumption that

she is dating Fred. Suppose further that the fact that Claire is dating Fred isn't intrinsically a particularly juicy piece of gossip. From Annette's current point of view, how should she update her degree of belief that Claire is pregnant upon learning that Claire is dating Fred? Annette should, I think, justifiably conclude that she probably got the information from Betty, in which case something has made the fact that Claire is dating Fred particularly juicy. If that is true, Annette can quite reasonably raise her degree of belief that Claire is pregnant above her previous conditional degree of belief that Claire is pregnant, given that she is dating Fred. In other words, for Annette, her current $pr(P|Q)$ can come apart from her judgment of her appropriate future degree of belief P given that she learns that Q .

At this point, a Bayesian will stop me and claim that none of this can be right because, if it were right, I have just given a counterexample to conditionalization, and have therefore proven too much! The Naive Bayesian will simply say, very plausibly, that Annette has conditionalized on a richer sample space, on the conjunction that Claire is dating Fred and that Betty tells her that Claire is dating Fred. But my point here is simply to illustrate how a two-place conditional probability differs from a hypothetical judgment of how one ought to revise one's opinions upon learning a piece of information. Annette can very plausibly now give a conditional judgment about how likely it is to her that Claire is pregnant, given that she is dating Fred, that comes apart with how likely she ought to think that Claire is pregnant, given that she has learned that Claire is dating Fred. When Annette is asked to make the simpler conditional judgment, she adds the hypothesis to her current stock of beliefs to get a judgment about what is the case. When Annette

supposes that Claire is dating Fred, she supposes that it is the case without supposing that she would know or find out about it if it were the case. When Annette makes the more complex judgment about how she would change her mind given her acquisition of a certain piece of evidence, she is allowed to make more hypothetical revisions to her stock of beliefs than in the simpler case in order to generate a conditional judgment. When assuming that she has learned about Claire's dating Fred, Annette is allowed to appeal to her beliefs about the ways in which she would learn such things, and such beliefs, together with her original assumption, can make a difference to the conclusion she might draw concerning how she ought to change her mind. Simple two-place conditional probabilities come apart from judgments concerning hypothetical situations of belief-change given the acquisition of evidence.

In respond to Argument #1, I contend that the complicated Strengthened Preface paradox is a case in which our conditional judgments come apart from our judgments of how we ought to revise our opinions upon a certain kind of learning. When we make a judgment conditional on Q , we take only for granted its now being the case that Q , adding that assumption, but nothing stronger or richer, as a hypothesis to our stock of opinions to judge what else is likely or unlikely. When I make a conditional probabilistic judgment on P_{17} given the conjunction of all the other propositions, I currently believe very confidently that I am wrong about one of them. My degree of belief in this proposition plays a role in how I make this conditional judgment on P_{17} . When I add the conjunction of every proposition except P_{17} as an assumption to my stock of beliefs to judge the likelihood of P_{17} , one of the

beliefs in my stock is my high degree of confidence that I am wrong about one of them. That is why I have such a low degree of confidence in this conditional degree of belief. However, when I am asked about how I ought to change my opinions given a learning experience of Q , I am required to be a little more sophisticated in how I generate my answer. I may see that the ways in which I learn that Q make a difference to how likely I would then take other things to be the case.³ When I am asked how I would hypothetically change my mind if I acquired information that the whole conjunction is true, I can judge that if I learned the conjunction of all of the other propositions one by one, I would and should give up confidently believing that I am wrong about one of them, and conclude that, lucky for me, this is an occasion in which I am right about every proposition in a set. In such a case, such a hypothetical judgment requires me to first *hypothetically* revise my high degree of belief that I am wrong about one of the propositions, *and then* conclude hypothetically that my future degree of belief in P_{17} would be my current degree of belief in P_{17} . This is the way in which I can reasonably have a current conditional degree of belief that comes apart from my conditional judgments about how I would appropriately reason when I learn something.

Such a response, I take it, will not be fully satisfactory to someone starting from a certain kind of Bayesian position. It may be claimed that distinguishing between a basic conditional probability and a (true, justified) judgment about how one ought to update given one's learning experiences already up-roots much current thinking about rational updating of degrees of belief. If

³We may now have information that our future selves will lose information, or will be subject to odd learning situations. Such information about our future selves can lead us to expect ourselves to revise in such a way that comes apart from our current conditional judgments. See [Arntzenius, 2003].

we take conditionalization to be the only rational way of updating, then one's conditional probabilities better line up with the right ways to update one's unconditional probabilities, given a learning experience. Otherwise, the right way to update will require a new unconditional probability different from an old conditional probability. Hence, from a Bayesian perspective, there can be no distinction between conditional probabilities and one current judgments about how one would and ought to update given a learning experience. I do not have any stronger a response to Argument #1 that can satisfy such a Bayesian. I recognize that we may simply have a clash of intuitions in this case.

Argument #1 might be called a Dynamic Bayesian objection in that it exploits Bayesianism's identification of intuitions about rational updating upon learning with static judgments of conditional probabilities. I have defused the objection by arguing against such an identification. Argument #2 relies on nothing of the sort. It simply challenges me to point out a significant enough disanalogy between the complicated Strengthened Preface case and paradigm examples of learning situations in which our degrees of belief one a set of propositions ought to be Kolmogorov-consistent.

I want to reiterate that the dispute between myself and my opponent is not whether an agent in a Strengthened Preface case is probabilistically inconsistent. He clearly is. He has reasoned to an inconsistent set of degrees of belief, Pairwise Independence, The Preface Belief, and the conditional judgment that I take to be motivated from the Preface Belief (that the degree to which I believe that P_{17} given that all but P_{17} are true, is low.). Rather, my opponent claims that, since it is clearly unreasonable to hold onto the

analogous inconsistent triad with respect to the coin tosses in Argument #2, it is also unreasonable to hold onto the inconsistent triad in the complicated Strengthened Preface.

I contend that the relevant disanalogy between the Strengthened Preface case and the coin-toss case is simply the presence of a higher-order degree of belief about one's lower-order degrees of belief being wrong in the Strengthened Preface case. Reasonable, evidence-respecting higher-order beliefs about being likely wrong with respect to one's lower-order beliefs are sometimes sufficient to generate reasonable inconsistencies in one's lower-order beliefs. There is no such higher-order degree of belief in the coin-toss case, and thus, no way for any higher-order degree of belief to generate inconsistent lower-order degrees of belief. In order to respect the evidence of my own fallibility with respect to my epistemic situation, I must have a high degree of belief in a proposition about myself, namely, that I am wrong about one of the propositions in my report. Merely being less than fully confident in each of the propositions in the report themselves is insufficient. Instead, I think the less than one degree of belief in each proposition is required by rationality from my high degree of belief in my own fallibility (which is not to say that I can only justifiably infer the former from the latter). In contrast, Argument #2 does not require that I have any degree of belief about myself and my fallibility. I simply have to have degrees of belief about the outcome of individual coin tosses, and about the outcomes of a series of coin tosses. The evidence in Argument #2 alone does not require the adoption of any degree of belief about yourself or your epistemic situation with respect to the coin tosses. Respecting the evidence in Argument #2 does not require any

higher-order probabilities, the complicated Strengthened Preface does.

The next question one might raise is whether having a higher-order probability is sufficient to rationally justify both my conditional degree of belief on the large conjunction and Pairwise Independence? Isn't it possible to be consistent while still having such a higher-order probability? To this kind of question, I simply ask the objector to attend to the details of the case. Suppose I have a high degree of belief that I am wrong about one of the propositions in the book. On the assumption that I am wrong about one of the propositions in the report, it seems clear that it will be P_{17} if all the rest are true. So I should have at least as high a probability that P_{17} is false given that the rest are true, as I do that I am wrong about one of the propositions in the report. Now I look at the sentence in my report that expresses P_{17} , and it says that cinder-block #20 weighs 10 pounds. It seems to me highly likely that cinder-block #20 weighing ten pounds in no way affects how much any of the others weigh. The same holds for all the other cinder-blocks. The content of these beliefs together imply that the weights of the cinder-blocks are pairwise independent, and this should be reflected in my degrees of belief. But of course, I don't think the propositions are *group-wise* independent precisely because of the higher-order preface-belief.

If you followed the previous line of reasoning and found it justifiable, you accept that someone can end up probabilistically inconsistent though reasonable. It doesn't follow from my being wrong about one of the propositions that the weight of one affects the weight of the other, i.e., my degrees of belief should reflect any pairwise independence of the propositions even given my higher-order degree of belief. I believe that my being wrong about one propo-

sition does imply that P_{17} is wrong given the truth of all the others. This latter line of reasoning I defend above in my discussion of conditional probability. I do not think that the line of reasoning is obviously unreasonable, even though inconsistent.

2.3.2 Lessons from the Preface Cases

One usual way of treating the original Preface case probabilistically is to show that one's Preface belief can be modelled as a low degree of belief in the conjunction of all of the propositions believed. This has led many people to think that "going probabilistic" is enough to solve the Preface Paradox. But it is not. Respecting evidence of my fallibility even probabilistically can lead me to static incoherence. The Preface Paradox is not a paradox to be solved. It is simply a counterexample to a certain thesis, the thesis that there are formal constraints on sets of rational beliefs or degrees of belief at a time. Indeed, I think the lessons from the Preface and Strengthened Preface are even more general, so as to extend not just to any formal constraints, but any rational constraint on sets of doxastic states of mind at a time.

Preface-type considerations show that even the most minimal versions of Static Rationalism are prone to counterexample. These considerations give us prima-facie reason to think that Static Rationalism, at least about all-or-nothing belief, is a degenerating research program. But such counterexamples point to a deeper problem with Static Rationalism. The Preface-Paradoxes share a striking similarity. Such examples begin with the hypothesis that we are justified in reasoning to some currently held beliefs from the evidence that we acquire. Our considered judgments are that people who are justified

in reasoning to their conclusions are rational. If rationality applied first and foremost to our currently held sets of beliefs, our considered judgments would co-vary with the states of mind that result from justified reasoning, not with the reasoning itself. But in fact, our judgments about rationality track the reasoning and not the resulting state of mind. We feel no pressure to charge some inconsistent agents with irrationality when their inconsistencies are upshots of fully justified reasoning. It is precisely the good reasoning that we evaluate for rationality, not the states of mind that result. No irrationality in the dynamics of belief-change, and there is no irrationality in the statics of belief states.

Thus, the form of Preface-type counterexamples indicate a general problem with arbitrary versions of Static Rationalism about all-or-nothing beliefs. Those problems carry-over to the probabilistic case. From the fact that we are not prepared to criticize dynamically rational agents who are statically “irrational,” we have powerful reason to conclude that our judgments of epistemic evaluation are only tracking the ways that agents reason and change their mind, not their beginning or end states-of-mind. If there are any conditions on static rationality to be had, and we already have reason to be skeptical, they must in the end be dynamically explicable.

2.4 Ideal Rationality to the Rescue?

In his classic defense of skepticism, Peter Unger [Unger, 1975] argued that nothing we ordinarily take to be flat is in fact perfectly flat. Nothing we ordinarily take to be dry is perfectly dry. Similarly, we are not perfectly

certain about the things we believe. For Unger, knowledge requires perfect certainty, from which it follows that we do not know anything we believe. Today, epistemologists tend to think of rationality the way Unger thought about flatness, dryness, and certainty. None of us are perfectly rational. This claim would be a harmless platitude if not for the accompanying methodological presumption that some epistemologists carry with it. Epistemic norms, so the presumption goes, are characterizations of perfect rationality, and ordinary imperfect people are rational to the degree to which they approximate or resemble perfectly rational agents. This conception of epistemic rationality, which I will call the Resemblance View, has been used to (1) motivate logical and probabilistic constraints on static rationality, and (2) resist intuitive counterexamples to such constraints. My aim in this section is to challenge the Resemblance View and its uses.

I will argue that, contrary to popular perception, the Resemblance view and the view that there are formal constraints on static rationality, a view I will call “Formality”, are in fact incompatible. My argument will center around the claim, established in the previous section, that it is sometimes more reasonable for real human agents to revise their beliefs in a way that renders them less like a logically perfect agent. Such situations are typically cases in which you must reason under circumstances of known epistemic imperfection. Such situations are pervasive, and norms for such reasoning make up a large part of rationality. These normative facts will undermine the conjunction of the Resemblance View and Formality. The Resemblance View itself is simply incapable of explaining such norms, and when conjoined with Formality, generates the wrong results.

2.4.1 Formality, Resemblance, and Epistemic Imperfection

How might appealing to ideally rational, epistemically perfect agents help someone who advocates Formality? Here is one way. Suppose that you only have two beliefs, P, and the belief that you have been wrong in the past. Does this second belief require you to believe that P is false? If you wrote P in a book, and you had to write a preface for the book, would it be reasonable to write that, although you investigated all of the claim in your book to the best of your ability, there is definitely an error in it? It seems like such a prefatory note would be quite unreasonable. Why? Because the belief that P is right there before you, you know all the evidence for it, and you fully believe it. To say that that you believe it and to say that it is false would be completely unreasonable. Similarly for two beliefs P, Q, or three, P, Q, and R. You believe that they are all true, hence you believe their conjunction. You do not believe that one of them must be false, even with your belief that you have erred in the past. If there were five hundred such propositions, then such a preface would be justified. But with respect to just one, two, or three beliefs, having a belief that you have been wrong in the past does not render you justified in believing that one of your beliefs is false.

The number of propositions you believe seems to matter to whether you can be justified, on the basis of past error, in concluding that one of them must be false. This is because we can have one, two, or three propositions “in full view,” so to speak, but not, say, five hundred. But epistemically *perfect* agents can have all five hundred in full view in the same way that we can have one in full view. So just as it would be crazy for me to believe one

proposition P , and use my history of error to conclude that P is also false, it would be crazy for an ideally rational agent to conclude that one of the five hundred propositions she believes is false. Thus, the Preface paradox is no counterexample to the view that ideally rational agents must be logically consistent.

The same argument holds for probabilistic consistency. The complicated Strengthened Preface can fail for an epistemically perfect agent for the same reasons. In addition, an epistemically perfect agent has an infinite capacity for information storage and computation. Hence, the inability of a human being to fully believe every logical truth is no counterexample to the view that rationality requires probabilistic coherence. Moreover, ideally rational agents, being logically perfect, do not have any history of logical error, so they do not have any evidence that can lead them to be less than perfectly certain in any logical fact. Thus, they are never in any position to be statically incoherent. As long as formal constraints are constraints on ideally rational agents, the considerations in the previous section do not apply.

But what is the relationship between such logically perfect agents and human beings? What do they tell us about the conditions under which *we* are reasonable. The answer is the Resemblance View. We are rational to the extent that we resemble such agents. Taking Formality as a starting point, we can use the Resemblance View to generate rational norms for human beings.

But can the Resemblance View adequately give us the norms for human beings if Formality is true? Anyone who takes our normative judgments in the Preface and Strengthened Preface at all seriously must answer in the

negative. Clearly we are not always consistent, not always probabilistically coherent, and not always reasonable for that matter. Given such imperfect epistemic situations, there must be better or worse ways for us to change our minds, both in light of new information, and in light of discovering our own internal failures. These norms that dictate the better or worse ways to revise seem to be incompatible with the Resemblance View if Formality is correct. The conjunction of the Resemblance View and Formality implies that rationality requires revisions that better approximate compliance with Formality. In other words, a belief-change that renders me more like the formally coherent ideal agent will always be more rational than one that renders me less like the ideal agent. This latter consequence of Resemblance+Formality is clearly false.

Consider the simplest of cases of rational reasoning upon discoveries of our own inconsistencies. Resemblance + Formality implies that, for any pair of agents A and B where A maintains her inconsistent beliefs, and B revises them in any way as to be consistent, B is more rational than A. Let A and B be two people with identical beliefs and evidence who encounter your favorite philosophical paradox. A and B both accept all of the premises, deny the conclusion, and recognize that the argument is valid. Neither A nor B have any further evidence about which of the premises is false. Similarly, we can consider the Preface Case, where we end up inconsistent, but have good evidence for each individual belief and the Preface belief. We can grant for the sake of argument that agents are statically irrational in such cases. We can even grant for the sake of argument that in all such cases, the *best way* to revise your beliefs is one in which agents end up consistent, that is, they

perfectly resemble the ideal agent. Even under these dubious assumptions, it is clearly false that anyone who ends up consistent is more rational than everyone who remains inconsistent.

Arbitrary Al may decide that in all cases in which he encounters valid arguments with premises he believes and conclusions he denies, he will give up believing Premise 1. Arbitrary Al always ends up consistent, better resembling the formally consistent ideal agent. Yet, he is much less reasonable in the way he manages his beliefs than Cautious Cal, who maintains his inconsistent beliefs for lack of evidence against any one of them. This is most clearly true in the Preface Case, where giving up the Preface belief is worse than believing it even though it renders you inconsistent. In general, in all ordinary cases in which one ends up with inconsistent beliefs by respecting conflicting evidence, arbitrarily giving up one to become consistent is a worse way to change your mind than maintaining your beliefs.

My opponents now can object to this argument by claiming that Formality is an incomplete characterization of ideally rational agents. My opponents can claim that there are other features that ideally rational agents might have, and that we must also resemble those features in order to be rational. These other features might account for our intuitions in cases of reasoning from imperfect situations. For instance, perhaps ideally rational agents cannot be arbitrary in the way they reason. Arbitrary Al is worse than Cal with respect to resembling some feature of ideally rational agents (call it Non-Arbitrariness), whereas Cautious Cal is worse than Al in respecting Formality. Ideally rational agents have many features, and we must resemble all of these features in order to be rational. Let us call this the

Many-Features View of ideally rational agents. Sometimes, given conflicts, we ought to privilege some features over others.

However, the Many-Features View cannot be independently motivated. Arbitrary Al might aim to reject the first premise of a paradoxical argument because he wants to restore consistency to his beliefs. Revising one's beliefs with the aim of restoring consistency cannot be considered arbitrary by my opponent's lights, for someone with such an aim is simply attempting to maximize resemblance to a perfect agent. For what feature of an ideally rational agent can Arbitrary Al be failing to resemble here? Yet, even with such an aim, Arbitrary Al is worse off by way of rationality than Cautious Cal. Moreover, the normative facts of the case do not motivate a Many-Features view. We do not think that Arbitrary Al is better off rationally than Cal in one respect, but Cal better than Al in another respect. Al is consistent and Cal is not, but Al manages his belief less reasonably than Cal. These are the facts that a theory must account for. There are no more normative judgments here for the Many-Features View to explain other than the fact that Cal is more reasonable than Al.

The most decisive objection to the conjunction of Formality and Resemblance, whether it is a Many-Features view or not, is the existence of evidence of our own fallibility. We do in fact have evidence of our own fallibility, especially our logical fallibility. Our past failures give us some reason to believe that we are currently failing in some respect, whether it is in some respect failing to believe something true, or failing to be perfectly rational. We do in fact need to respect such evidence in our beliefs and degrees of belief in order to be reasonable. We can respect such evidence by believing that we

are not perfectly rational now, or not perfectly coherent now, or currently in possession of false beliefs. These evidence-respecting beliefs cannot be said to be unreasonable. How can any Resemblance view account for this normative fact? Any agent who has such evidence, and respects it in the way they manage their beliefs, is not ideally rational. So what feature of an ideally rational agent can I possibly be resembling when I respect such evidence? Possessing such evidence renders me unlike an ideally rational agent, revising my opinions in light of such evidence can render me even less like an ideally rational agent. As we have seen in the simple Strengthened Preface, I can reason from a state of mind that better resembles an ideal agent with respect to Formality to one that less resembles an ideal agents with respect to Formality. The need to respect evidence of my own fallibility undermines any Resemblance view, and it undermines any Resemblance View conjoined with Formality.

In the end, Resemblance and Formality generates the wrong results in cases of reasoning under circumstances of epistemic imperfection, and in situation in which we must respect evidence of such imperfection. Resemblance and Formality imply that certain ways of reasoning under such circumstances are better than others, when they are clearly not. Resemblance and Formality also fail to generate norms for reasoning upon evidence of fallibility. These failures to capture pervasive norms for rationality suggest that Resemblance and Formality should be abandoned as a way of explaining any norms for rationality. There is nothing special about norms for reasoning under circumstances of epistemic imperfection. If we are always imperfect, we are always reasoning in such circumstances. Respecting evidence of our own fallibility

is just a special case of revising one's opinion to respect evidence generally. Why should general norms for rationality admit of an explanation according to the Resemblance View, but special cases of it not?

2.4.2 Alternatives to Resemblance and Formality?

Perhaps there are nearby alternatives to the Resemblance view that saves Formality in light of these considerations. In the realm of practical rationality, Michael Smith advocates *The Advisors View* [Smith, 1995] according to which you have reason to do what an idealized version of yourself would advise you to do. With respect to epistemic rationality, the analogous view would be that the norms or constraints on rationality for real agents are determined by the hypothetical advice of an ideally rational transformation of yourself. According to the Advisor View, normative epistemology is doing two things, describing an ideally rational agent, and describing the advice such an agent would offer their real, non-ideal counterparts.

The Advisor's View seems to be true of a certain class of reasons for action. Is it true in the case of epistemic rationality? Does the Advisor's view conjoined with Formality do any better at dealing with the problems of reasoning under epistemic imperfection? It is unclear to me. Even supposing that we have an independent grasp of the descriptive characteristics of ideally rational agents—they are computationally unconstrained, perfectly consistent, and so forth—we do not seem to have any independent grasp as to how such agents would advise us to manage our opinions in epistemically problematic circumstances. We have even less of a grasp as to how we should make sense of the better or worse ways to reason oneself out of inconsistencies

on the Advisor's view. For instance, suppose that as a matter of fact, in a Preface-type case, all of my beliefs except my Preface belief are true. Since I have the Preface-belief, I am inconsistent. What would an ideally rational transformation of me advise me to do? By the Formality hypothesis, the ideally rational version of me is consistent, so he does not believe exactly what I believe. Would he advise me to give up my Preface belief? If he would, then it follows that I would be more reasonable in giving up my Preface-belief than retaining it. This is clearly the wrong result. Does he advise me to stay as I am? It is not clear, especially because it is not clear what such an ideal agent believes. He can be consistent by not having the Preface-belief, but he can be consistent a variety of other ways also. Simply lacking the Preface-belief in no way suggests that the ideal agent will advise me to maintain all of my beliefs as they are. It also in no way tells me that it is better to maintain all of my beliefs as they are than to give up my Preface belief.

The murkiness of how an ideal agent would advise us carries over to non-preface-type cases of reasoning under inconsistency. If I believe inconsistently due to conflicting evidence, it is sometimes better to keep believing what I believe rather than giving up one particular belief. If the ideally rational agent is consistent, he will be consistent albeit with some false beliefs, or consistent with all true and no false beliefs. In either case, it is unclear how such an agent would advise me to change my mind. If he has all the true beliefs and no false beliefs, does he advise me to give up the false belief, and believe as he does? That would, again, be the wrong result. If I believe an inconsistent triad P, Q, R, based on conflicting evidence, and I discover this,

giving up R without special reason is worse than maintaining belief in all three, even if R turns out to be the false one. Whatever the advice would be, the Advisory View still cannot make sense of why some ways of reasoning are more reasonable than others. Do we also need a notion of an ideally rational agent giving both better and worse advice? Our theoretical postulations here seem to have gone off the rails.

To do any adequate work, the Advisory View must be supplemented with a satisfactory account of the advice we would get from an ideally rational agent. If such an account is available, the Advisory view could be a promising way to save the project of appealing to ideally rational agents to explain epistemic norms, and could also save the view that ideally rational agents adhere to Formality. But to adopt such a view is already to acknowledge that human beings do not need to be formally coherent to be rational, and that they need not move closer and closer to formal coherence in order to revise their beliefs rationally. Claiming that perfectly rational agents need to be formally coherent would amount merely to claiming that formal coherence is a required feature of agents who give all and only the right epistemic advice. And while such a view may end up useful in epistemological theory, I see no reason to accept it as anything that saves the spirit of Static or Binary Rationalism. We were looking for independent constraints on *our* sets of beliefs or degrees of belief at a time. Instead, we are told that the agents who can tell us what such norms are must be formally coherent. What were considered formal constraints on our static rationality are still no constraints on rationality at all. We arrived at such a result attending to intuitive cases that elicited our ordinary normative judgments. Appealing to ideally rational

agents simply seems superfluous.

2.5 Summary

Binary Rationalism requires static and dynamic rationality to be equally basic, neither reducible to the other. In this chapter, we began our search for independent conditions on static rationality, and concluded that the minimal formal constraints like logical or probabilistic coherence are inadequate. The problem with such formal conditions is that dynamically rational agents can end up formally inconsistent, but intuitively exhibit no rational failure. This constitutes good evidence that there cannot be any irrationality without dynamic irrationality. If there are no formal constraints on static rationality, are there any independent constraints on static rationality? The next chapter will argue that even a minimal non-formal constraint fails, and fails for the same reasons as formal constraints.

Chapter 3

Against Static Rationality

Part II

In Chapter One, we noted that Static Rationalists consider “good basing” to be a constitutive feature of rational belief. A belief must be based on good enough reasons to be rational. We have also seen that we cannot give the right analysis of rational change of belief in terms of good basing. In the previous section, we concluded that formal coherence, broadly construed, cannot be a condition on static rationality. If good basing requires formal coherence of some form, it follows that even good-basing is no condition on static rationality. Such a conclusion seems to have some intuitive support. If I base my belief that P on my belief that Q but simultaneously believe a set of things inconsistent with P, or inconsistent with Q, it seems that my belief that P cannot be well-based on Q. However, such abstract considerations are inconclusive. Static Rationalists might argue that believing that P on the basis of my belief that Q can be appropriate even when I believe something

inconsistent with Q. According to a Static Rationalist, ruling out formal coherence constraints does not thereby exclude non-formal constraints like good-basing. And if good basing turns out to be irreducibly static, static rationality would have an independent, non-formal constraint, contrary to Dynamicism.

The case of *good-basing* will provide an interesting test case for the Dynamicist hypothesis about the Preface paradox, namely, that it is an instance of a more general phenomenon in which there cannot be irrationality in belief-states without irrationality in belief-changes. Dynamicism implies that you cannot have a poorly-based belief and also be (1) fully dynamically rational and (2) epistemically unreasonable. Binary Rationalism requires the possibility of a case in which a person reasons rationally and justifiably to beliefs that are poorly-based, and that agents in all such cases are unreasonable in some way. Test cases that tell in favor of one view or the other can therefore be easily constructed.

Let's consider some paradigm cases of poorly-based beliefs; (1) a belief based on insufficient evidence, (2) a belief based on defeated evidence, and (3) a belief based on beliefs that do not constitute evidence at all. Good test cases should involve agents who end up with beliefs as in (1)-(3) but also exhibit perfect dynamic rationality. If agents in such cases exhibit some sort of rational failure, and such rational failure does not arise from a failure in dynamic rationality, then the Static Rationalist wins. If there are no such cases, then Dynamicism wins. I will argue through the test cases below that dynamically rational agents can end up as in (1), but are guilty of no rational failure at all. In addition, dynamically rational agents simply cannot end up

as in (2) and (3). Any rational failures exhibited in cases that purport to be like (2) and (3) must be instances of *dynamic* irrationality. In the end, all of the paradigm cases of poorly-based beliefs will be cases that favor Dynamicism.

3.1 Insufficient Evidence and Forgetting

Let us consider the possibility of Case #1. Is it possible for me to be dynamically rational and end up believing that P, but on the basis of insufficient evidence? It is clear that I cannot do this while simultaneously keeping track of my reasons from which I concluded that P. If I have sufficient evidence or justification to conclude that P is the case, and I reason to P from such evidence, then my bases for my belief that P consists in the evidence that I reasoned from. If, for some reason, I currently base my belief that P on *insufficient evidence*, I must now lack some evidence I once had. If basing is wholly static, a relation that holds between states of belief not analyzable in terms of belief-change, then forgetting a piece of my justification for concluding that P seems to be the only way I could reason justifiably to the belief that P but now lack some required evidence for P. Otherwise, I will be fully justified in reasoning to my belief that P, and remember so reasoning, and possess all of the information from which I reasoned to P, but also be basing my belief on insufficient evidence. This latter kind of situation seems impossible.

We must now ask, if a person uses a piece of evidence that sufficiently justifies him to conclude that P, and subsequently forgets that piece of ev-

idence while maintaining the belief that P, is he thereby guilty of any kind of rational failure? If the answer is no, Case #1 is no counterexample to Dynamicism. If the answer is yes, there seems to be an irreducibly static constraint on rationality that is incompatible with Dynamicism.

Let us distinguish Case #1 from a case in which I recognize that I lack a crucial piece of evidence for P, and now believe that P, but nonetheless refuse to revise my belief or degrees of belief in P. This is a case of unreasonableness on my part, but such unreasonableness can be explained in terms of dynamic irrationality. In this latter kind of case, a certain piece of evidence, a recognition of lacking evidence, justifies a certain revision of belief, or degree of belief. A failure to abide by such a norm of belief-revision results in a rational failure, a failure to respect reasons for belief-revision. But in the original case, Case #1, there is no such recognition. All we have is a case in which an agent justifiably reasons to a certain conclusion, P, and happens to forget some piece of evidence that justified that reasoning. In that case, is there any kind of rational failure?

Let us also distinguish Case #1 from two other nearby kinds of cases. The first is a case in which I cannot cite a piece of evidence that I used to reason to my belief that P when I am challenged to do so. The second is one in which I have evidence that I have lost or forgotten a piece of evidence. I think in both of these cases, I count as having evidence that I now lack evidence. In these cases, one can make a strong case that I ought to revise my belief or degree of belief in P, and that failure to do so would indicate unreasonableness on my part. I personally do not want to make such a case, though certain epistemologists do. But regardless, we can all agree that *if*

agents are guilty of a rational failure in such situations, then they are guilty of a dynamic failure, that is, a failure to revise in light of reasons to revise. We can also agree that this kind of case is different from the test case at hand, a case in which I simply forget some piece of evidence without having evidence that I have forgotten.

Now, in our original test case, Case #1, am I guilty of a rational failure or not? Once we distinguish the test case from nearby cases of dynamic irrationality, I believe that the intuitive answer to this question is “no.” However, I cannot depend solely on intuitive answers here because some epistemologists purport to have differing opinions, for instance, Evidentialists like Feldman and Conee [Feldman and Conee, 1985]. However, I think a good argument can be made that we in fact do not convict people of rational imperfection in cases of mere evidence-forgetting. For if we did find some kind of rational failure in such cases, there should be some pressure to believe that agents in such situations have reason to revise their “poorly-based” belief. A rational failure should be a reason for belief-change. Now in the original case, an agent who justifiably reasons to a conclusion that P, but now forgets a piece of evidence constituting that justification, does not ipso facto have any reason to change her mind. For that agent, there is no evidence that she lacks a crucial piece of evidence, no evidence of forgetting, no evidence against P, and her reasoning to P was completely justified. Imagine an agent who, in just such a situation, did in fact revise her belief (or degree of belief) that P (or some other relevant belief). Such an agent would have changed her mind in an unjustifiable way, with no relevant evidence and background beliefs that justify that change. Therefore, the original agent in a case of

evidence-forgetting does not have any justification to change her mind, and is therefore guilty of no rational failure, not even a static one.

To give a more concrete example, imagine two jurors in identical evidential situations prior to concluding that a defendant is guilty. For simplicity, imagine that there are exactly 17 pieces of evidence, and both jurors begin with the same relevant background beliefs about the case. As a matter of normative fact, the evidence together with such background beliefs justifies the conclusion that the defendant is guilty. Both jurors justifiably conclude that the defendant is guilty, and for good measure, have the same degree of confidence in their conclusion. Subsequent to their conclusions, juror one forgets one piece of evidence; juror two does not; and everything else is equal. Since everything else is equal, juror one does not have any more evidence of her own forgetting the evidence than does juror two. Neither has any evidence that he lacks an ability to cite all 17 pieces of evidence he originally had. Neither has any new evidence of the defendant's innocence, nor evidence that their original evidence was misleading. Both reached the conclusion for the same reasons, both have the same doxastic attitudes toward the same propositions. It seems clear that juror one does not have any more reason to change his mind than does juror two. For juror one, since there is no reason for belief-change, there is not rational failure to begin with.

Those who feel pressured to judge that juror one does have some reason to lower her confidence simply in virtue of forgetting evidence should consider an alternative explanation of such a judgment. One may simply be conflating the Two Juror case with one of the nearby but distinct cases I mentioned above. It is true that when asked to cite reasons for his belief, juror one is not

in a position to cite the same reasons as juror two. Similarly, if asked to re-deliberate from memory alone as to whether the defendant is guilty, thereby starting from a point in which one suspends judgment as to guilt, juror one and two will be in different positions with respect to their evidence. From memory alone, redeliberation might require that juror one and two differ in their conclusions. But such situations do, and should, *change* juror one's reasons for belief-change. For any specific piece of evidence, if juror one is unable to cite such evidence due to memory loss, that inability is evidence for him that a piece of evidence is absent or has been forgotten. Similarly, when re-deliberating from memory, the inability to conjure up a piece of evidence renders that piece of evidence unable to give juror one reason to conclude the defendant's guilt with the same original degree of confidence. In both such cases, an agent *acquires* a reason to revise her doxastic attitudes in the process of reexamining her evidence. We are correct to think that juror one has some reason to revise some of his beliefs or degrees of belief once he is prompted to defend them, or to re-deliberate from memory alone. But that reason is not there solely in virtue of the fact that juror one forgets a piece of evidence. Without a failure of dynamic rationality, mere forgetting of evidence is insufficient for any form of static rational failure, even if the memory loss leads to a "poorly-based belief".

Thus, Case #1 is no counterexample to Dynamicism. Agents with beliefs based on insufficient reasons who are dynamically rational, as in cases of forgotten evidence, are perfectly reasonable. Indeed Case #1 lends support to Dynamicism. Any nearby cases of genuine irrationality happen to be explicable in terms of dynamic irrationality. When I have beliefs based on

insufficient reasons and am unreasonable, I am always unreasonable due to some failure of rational belief-change.

3.2 Basing on Defeated Evidence

Case #2: The other typical case of a dynamically rational but “poorly-based” belief is a belief based on defeated evidence. If I conclude, justifiably, that Prince of Darkness will win the race from my belief that Prince of Darkness is the fastest horse in the group, and acquire evidence that Prince of Darkness is slower than Queen of the Night, my belief that Prince of Darkness will win is now “poorly-based” in the sense that its basis is now defeated. Am I guilty of a rational failure if I still believe that Prince of Darkness will win the race, even upon acquisition of such evidence? Yes I am. However, any rational failure in such a case seems straightforwardly dynamic. A defeated piece of evidence is a piece of evidence you now have reason to reject. Basing a belief on a defeated piece of evidence is simply a case of having justification to reject something you believe by having justification to reject a belief on which it is based. Failing to revise the beliefs for which you have justification to revise is a straightforward case of a failure of dynamic rationality, and therefore no counterexample to Dynamicism. Case #2, then, fails to support Static Rationalism.

3.3 Basing and Changes of Grounds

Case #3: Finally, a belief can seemingly have been arrived at rationally but which is now based on bad evidence, not because one has reason to reject that evidence, but because the bases of my beliefs are no evidence for them at all. Is this possible? Consider the case where I justifiably conclude that a defendant is guilty from the forensic evidence, exhibiting dynamic rationality *par excellence*. While maintaining my belief to the same precise degree of confidence, I now base my belief on a prejudicial attitude against people with long hair and tattoos, forgetting or dismissing my original reasons. My belief in the defendant's guilt was arrived at rationally (dynamically rational) but is now based on bad evidence or reasons. Am I guilty of some rational failure here? The answer of course is yes. There appears to be a clear case of a failure of rationality without a failure in the rationality of how one originally reasoned to a belief. Dynamicism needs to account for this kind of case.

Again, I think the rational failure in this case is a failure of dynamic rationality. The dynamic irrationality involved here concerns a dynamically irrational change of grounds or bases for my beliefs. When I change my grounds from my beliefs about the forensic evidence to my beliefs about the tendencies of people with long hair and tattoos, I am guilty of a kind of unreasonable change, not with respect to my beliefs, but with respect to my bases or grounds for beliefs. It is the irrational change of grounds that explains why I am guilty of a rational failure, not an irreducibly static condition on static rationality.

The idea that we can reasonably or unreasonably change our grounds for our beliefs is not a merely ad hoc postulation meant to save Dynamicism.

Evaluating changes of grounds or bases is a general and familiar feature of normative evaluation, moral, practical, and epistemic. Bush Administration officials are famously rigid with respect to their proclamations of the moral imperative to invade Iraq. They are equally famous for radical shifts in what they present as their grounds for the invasion. One can very reasonably support the invasion while criticizing the Administration for unreasonable shifts in their grounds for the invasion (assuming that one takes the Administration at their word at any given moment regarding their grounds for the invasion). People may have originally believed much about the health benefits of green tea largely on faith and unreliable testimony. These people could very reasonably have updated their grounds and base their beliefs on subsequent scientific research that have verified the health benefits. Post-hoc searching for evidence for one's deeply held beliefs is very familiar. Some argue that much of scientific beliefs work this way. Anselmian methodology in theology seems to be explicit in endorsing the search and revision of grounds for belief. The phenomenon of changing grounds or bases for beliefs is familiar to ordinary life, not a mere philosophical construction meant to save a theory.

The challenge for the Dynamicist is to ensure that the rationality of a change of grounds is not reducible in any sense to static rationality. An account of what makes for a rational change of grounds must be irreducibly dynamic. Such an account is clearly available.

3.3.1 Dynamicism and Rational Changes of Grounds

Two approaches are available to explain rational changes of grounds in terms of dynamic rationality: we can take changes of grounds as basic and explain

the goodness of grounds in terms of the rationality of its changes, or take the goodness of grounds as basic and explain the rationality of changes of grounds in terms of their goodness. I prefer the latter strategy, for it allows us to reduce the rationality of changes of grounds to the rationality of changes of belief. As we saw in Chapter One, we can explain the goodness of grounds in terms of rational changes of belief. For an agent A, a ground G for a belief B is better than another G' for B when it is more reasonable for A to reason from G to B than G' to B on the assumption that A has not made up his mind about B. A change of ground is rational if it is a change to better grounds, where the goodness of a ground is determined by the rationality of belief-change. Once we reduce the rationality of changes of grounds to the rationality of changes of belief, the principles of dynamic rationality carry-over to the rationality of changes of grounds. (I give such principles in the next Chapter.)

It follows on my analysis that one cannot change grounds reasonably and also end up with a belief based on bad grounds or reasons. This consequence raises the following objection: Suppose I believe that Asian-Americans are good at math, and meet Fred, who looks to me to be Asian-American. I justifiably conclude that Fred is good at math on the basis of my belief that he is Asian-American. Given my state of belief antecedent to my change of belief, I justifiably ground my belief that Fred is good at math on my belief that he is Asian. But haven't I ended up with a belief based on a bad reason? Beliefs about Fred's ethnicity are a bad basis, for me and for anyone, for beliefs about their mathematical ability. This notion of a bad basis, so the objection goes, is not captured by justified and unjustified belief-change

and change of grounds. There is a potential problem for Dynamicism here that needs to be discharged before we proceed.

3.3.2 Objective versus Subjective Grounds

One seeming weakness of reducing the rationality of changes of grounds to the rationality of changes of belief is that the goodness of a ground for belief would then depend on an agent in the same way that the rationality of changes of belief depends on the agent. For some, the goodness or badness of grounds or bases for belief ought to be an objective, agent-neutral matter. Facts about ethnicity are *objectively* bad grounds for beliefs about mathematical ability. The rationality of changes of belief are subject-relative. Thus, the rationality of changes of belief seems to be unable to explain the goodness of grounds.

To illustrate the issue at hand, consider the case of the celebrated 16th-century astronomer Tycho Brahe. According to textbook descriptions of the history of science, Tycho's recorded observations of celestial bodies justified the adoption of Kepler's laws for planetary motion. We might say that Tycho's records were good objective grounds for believing that Kepler's laws hold. Of course, Kepler devised his laws, from Tycho's records, *after* Tycho's death in 1601. Tycho himself knew nothing of Kepler's laws (he wasn't even a Copernican!), nor of any other theory that explained the motion of the heavens better than his own version of geocentrism. In one sense of grounds, or reasons for belief, Tycho had in his hands objectively good reasons for adopting Kepler's laws. He failed to do so for two reasons; first, he could not discover the phenomena of stellar parallax, a consequence of heliocen-

trism; and second, he had no access to Kepler's laws that better explained his records. Setting the issue of stellar parallax aside, Tycho had good objective justification for believing that Kepler's laws held of the planetary heavens, despite the fact that he had no access to the propositions expressed by Kepler's laws. But relative to Tycho's beliefs and the alternative hypotheses he entertained at the time, he had no subjective grounds to come to believe Kepler's laws.

Tycho Brahe is a person who seems to fail to change his beliefs when confronted with objectively good grounds. The original "Fred-is-good-at-math Case" is a case in which I change my beliefs on objectively bad grounds. It is clear that the notion of an objectively good and bad ground for belief-change cannot be explained in terms of whether it would be rational for a certain kind of agent to reason from the ground to the belief. Such rationality depends on what else the agent might believe and the propositions he has access to at the time of deliberation.

The inability of Dynamicism to explain the objectivity of grounds for belief is no mark against it. I contend that Dynamicism cannot explain objective grounds for belief-change because people cannot in general be faulted with a failure of *rationality* for failing to respect such grounds. Tycho Brahe cannot be faulted for failing to come to believe Kepler's laws; he did not think of them, he did not know of them, and he had many true beliefs about the failures of 16th-century heliocentrism to explain planetary motion. Tycho's case indicates that failures of respecting objective grounds or reasons are due to familiar sorts of factual or normative ignorance. When you do not have a belief in a relevant empirical fact -that stellar parallax is indeed ob-

servable with powerful telescopes-, or access to some alternative explanation -that planetary orbits are ellipses and not circles-, or fail to know that a certain kind of data supports a proposition -that recorded planetary movements overwhelmingly support a heliocentric model with elliptical orbits-, you will fail to respect an objective reason for belief. However, simply lacking a belief in some empirical or normative fact, or lacking access to some alternative explanation, does not in itself render you unreasonable. Similarly, an agent cannot be faulted for irrationality simply in reasoning from an objectively bad reason that follows from all of the other things he believes. Whatever normative failure is involved in failing to respect objective grounds or reasons, it is not a failure of rationality. Dynamicism is a thesis about epistemic rationality. Objective grounds or reasons for belief are outside of its domain.

On the other hand, in the epistemic realm at least, one's ignorance or failure to believe can be an upshot of an antecedent failure of dynamic rationality. I might fail to have a relevant belief I have sufficient subjective justification to have, or fail to revise a belief I have sufficient subjective justification to revise. Such a failure of dynamic rationality can lead to the kind of ignorance that results in a failure to respect an objective reason. Tycho Brahe might have had data for stellar parallax that he refused to see as relevant to the heliocentrism-geocentrism debate. He could have formulated the beginnings of a Keplerian hypothesis but stopped out of fear of the Inquisition. This hypothetical Tycho could have had the same bits of ignorance and beliefs as the historical Tycho, but his failures to respect objective reasons would have been unreasonable. The hypothetical Tycho's failure to respect objective reasons would have arisen from an antecedent failure to respect

subjective reasons. So whereas we are not generally unreasonable in failing to respect objective reasons, we can be epistemically unreasonable when such failures are due to antecedent failures in dynamic rationality. In such cases, Dynamicism can explain why a failure to respect objective grounds seem to be unreasonable; such failures are upshots of antecedent failures of dynamic rationality.

So far we have distinguished objective from subjective grounds for belief-change, identified subjective grounds as the difference-maker to rationality, and offered Dynamicist reductions of the goodness of subjective grounds for belief-change, and the rationality of changes of grounds. Let us return to the original objection that prompted this discussion. The question was whether I can have a belief based on a bad reason without any sort of failure of dynamic rationality. The answer is yes and no. When I believe that Fred is good at math on the basis of my belief that he is Asian-American, I do base my belief on a bad objective reason (assuming of course that being of Asian descent does not matter to mathematical ability). But such a bad basis for belief does not in itself make a difference to rationality unless accompanied by a failure of dynamic rationality. I either fail to respect subjective reasons I in fact have, ones that defeat my line of reasoning, or I am ignorant of some empirical or normative fact that defeats this line of reasoning because of some other failure to respect my subjective reasons. In conclusion, while it is possible to believe on the basis of bad objective reasons, only when such beliefs are upshots of antecedent failures to respect subjective reasons is there any failure of epistemic rationality.

Thus, Case #3 is insufficient to support Static Rationalism and under-

mine Dynamicism. We can explain the irrationality of agents in situations like Case #3 as failures of those agents to reasonably change their grounds for belief. The rationality of changes of grounds can be fully explicable in terms of dynamic rationality, and can be subject to the same principles as reasonable changes of beliefs.

3.4 Summary: The Dynamicist Argument Schema

For a Static Rationalist, good-basing is considered to be a necessary condition on static rationality. Hence, poorly-based beliefs ought to be sufficient for a failure of epistemic rationality. However, we have seen that the paradigm cases of poorly-based belief do not undermine Dynamicism and support Static Rationalism. Cases #1-#3 fail to refute the thesis that there can be no failure of rationality without a failure of dynamic rationality. In fact, Cases #1-3 support Dynamicism by pointing out that any case in which we find poorly-based beliefs, they are either cases in which there is no rational failure, or a case in which such rational failures are dynamically explicable. A non-formal constraint on static rationality like “good basing” fails to save the idea that there are independent conditions on static rationality. All counterexamples point to Dynamicism.

Chapters Two and Three have instantiated a general strategy, an argument schema, for resisting the idea that there are independent conditions of static rationality. Someone offers a proposal for a proposed condition on rational belief. We have already established in Chapter One the autonomy and independence of dynamic rationality. So we first ask whether this proposed

condition is reducible or explainable in terms of the conditions for rational change of belief. If it is, Dynamicism wins. If the proposed condition isn't explicable in terms of dynamic rationality, then there will exist a possible scenario in which an agent reasons rationally to a set of beliefs that fails to satisfy the proposed condition. We should ask, then, whether or not agents in such scenarios are guilty of any rational failure whatsoever. Dynamicism states that no conditions will generate an answer of "yes" to the preceding question.

We have seen much evidence for Dynamicism in the previous two chapters. The Preface Paradox is an instance of such a possible scenario, where logical consistency is the proposed condition. The Strengthened Preface paradoxes illustrate that even probabilistic coherence fails to be such a condition. Since logical and probabilistic coherence are the minimal formal constraints on static rationality, it is likely that there are no such constraints. Finally, even minimal non-formal constraints like "good basing" fail to survive the Dynamicist Argument Schema. The Dynamicist Argument Schema is not simply a general methodology for refuting alternative views. It is an argument whose strength flows from the fact that our judgments of rationality solely track changes of mind. Thus, any successful application of the schema to a proposed condition C will support Dynamicism twofold: it would refute an alternative to Dynamicism, and it would show that any apparent feature of epistemic normativity is either real and dynamic, or only apparent because it is not dynamic.

While the Dynamicist thesis has not been conclusively demonstrated, general skepticism concerning the project of seeking independent conditions

for rational belief now seems warranted. Without independent conditions for rational belief, the only conditions on epistemic rationality are conditions on rational changes of belief.

Chapter 4

Dynamic Rationality

4.1 Replacing Static with Dynamic Rationality

The previous three chapters established a prima facie case against the theoretical fruitfulness of seeking conditions on static rationality. There are no minimal formal constraints on rational sets of doxastic states. There does not seem to be any irrationality in beliefs without irrationality in belief-changes. If there is such a thing as being statically rational or irrational at all, the conditions seem definable or reducible to rational and irrational changes of belief. Appealing to idealized epistemic agents does not help. Nonetheless, the arguments for Dynamicism in the previous chapters were not absolutely conclusive. In the next two chapters, I will bolster the case for Dynamicism by replacing static concepts and theories with dynamic ones, and show how this can solve and dissolve outstanding problems in epistemology. In

effect, I will offer Dynamicism as a better theory and better explanation of epistemological judgment and evaluation.

A normative property or relation is dynamic when it applies to changes of doxastic states, not the doxastic states themselves. Changes of mind might be changes in individual beliefs or degrees of beliefs, small sets of (degrees of) beliefs, or shifts of one's total beliefs. Changes of doxastic states include the addition of new doxastic states, the elimination of old doxastic states, and the revision of existing doxastic states. Existing Bayesian and non-Bayesian theories of belief-revision recognize the normative importance of reasoning and changes of mind. The spirit of such theories, as theories of belief-revision, is admirable; but they universally presume that the principles of dynamic rationality are principles of the conditions under which static rationality is preserved or restored over time. Previous chapters suggest that we must treat dynamic rationality as basic.

4.1.1 The Dynamic Foundations of Epistemic Norms

What are the fundamental notions in a purely dynamic theory of epistemic rationality? Let us begin with the notion of *being justified to reason from point A to point B*.¹ A point *from which* you reason is an individual or set of contentful doxastic states you currently possess. A point *to which* you are justified to reason is an individual or set of contentful doxastic states that you are in a position to adopt. In my terminology, you can be justified to reason

¹According to many readers, this locution “being justified to” is controversial both in its connotations and in its grammaticality. I adopt the term on the advice of Jim Pryor, and keep it throughout for consistency, but I acknowledge its oddness. Those who find it linguistically unpalatable can substitute it with “having sufficient justification to” or “being licensed to.”

from point A to point B without actually so reasoning. As a child, when my mom arrived home ten consecutive Fridays with Cantonese roast duck for dinner, I was justified to conclude that I will have roast duck for dinner on Fridays. I was so justified even when I failed to draw such a conclusion, perhaps believing in vain that we would have fried chicken instead. I can be justified to come to believe something without coming to believe it. Similarly, I can be justified to revise my beliefs without actually revising. I currently believe that my dog is loyal. I continue to believe this even though I also believe he is willing to walk away with any stranger with a biscuit. This latter belief justifies me to revise my belief in my dog's loyalty, though in fact I do not revise my belief. Being justified to reason from point A to point B implies that you inhabit point A, for they are beliefs you currently have. However, it does not imply that you in fact reason from point A to point B, nor does it imply that you inhabit point B.

If you do reason from point A to point B justifiably, I will say that you are justified *in* reasoning from point A to point B. If I reason from point A to point B, then I am justified *in* so reasoning just in case I am justified *to* reason in that way. Thus, being justified in reasoning from point A to point B is simply a matter of being justified to reason in that way, and actually reasoning in that way.

On a wholly dynamic theory, principles of rationality will be principles of justification to reason from one point to another. I will use the locutions "point A justifies you to reason from point B to point C," interchangeably with "you are justified to reason from point A and point B to point C." In doing so, I am assuming for now that beliefs are what justify you in reasoning

from one point to another. However, a dynamic theory of rationality is not limited to reasoning from beliefs to beliefs. Though my focus will be on belief-changes that are based on existing and newly acquired beliefs, my account below in a way assumes that starting points from which you reason can include experiential states of mind like perception (a la Pollock's *percepts* [Pollock, 1995]), memory, and reflection. A total theory of rational belief-change will need take into account the rational ways to respond to your experiential states of mind.²

I will write $b(P)$ as the belief that P , $db(P)$ as the disbelief that P , and $a(P)$ agnosticism about P . The notation $\{b(P), b(Q)\}$ will represent the set of beliefs containing the belief that P and the belief that Q . Hence, you can reason from $b(P)$, or you can reason from $\{b(P), b(Q)\}$. You can similarly reason from $db(P)$ or $a(P)$. When I say that you are justified *to reject* a point, I intend to be neutral with respect to whether you are justified to disbelieve something you already believe, or whether you are justified to become agnostic about something you already believe. For example, I can be justified to reject the belief that P by being justified to disbelieve that P , or by being justified to become agnostic about whether P . In addition to being justified to reject an individual belief, I can also be justified to reject $\{b(P), b(Q)\}$, meaning that I am justified to disbelieve or be agnostic that both P and Q are true. In general, "rejection of $\{b(P_1) \dots b(P_n)\}$ " will be agnosticism or disbelief that all of the propositions are true.

I am treating disbelief as a distinct *sui generis* state from belief and agnosticism. I think that one can take a distinct attitude toward a proposition

²A total theory might even need to take into considerations other states of mind, like interests and goals. Thanks to Gil Harman for this point.

in rejecting its truth without thereby believing its negation, or thereby being merely agnostic about its truth or falsity. It will be difficult to make a noncontroversial case for such an attitude here, but I think we can disbelieve things like “This sentence is false” without believing it, believing its negation, or being completely non-committal about it(like we are the exact population of Kazakhstan). Once we allow disbelief into our ontology of mental states, there is no principled reason to preclude it from being a prevalent attitude in our doxastic lives as the contrary to belief. Those who accept degrees of belief will find none of this controversial, as having a low degree of belief in a proposition in no way (metaphysically or psychologically) necessitates a high degree of belief in the negation of that proposition.

4.1.2 The Principles of Dynamic Rationality

Epistemic rationality at least in part depends on evidence. On standard views, current states of evidence determine current states of rationality (Feldman and Conee [Feldman and Conee, 1985]). On the Dynamic view, evidence constitutes reasons for belief-change. Let me begin with a series of cases that motivate some principles governing the rationality of belief-changes in response to normal cases of visual and doxastic evidence.

Match 1: I take a match out of the box and examine it, thereby coming to believe that it is a normal match. I also believe that generally, normal matches light when struck, though abnormal matches do not. I am justified to conclude from such beliefs that the match will light when struck. However, I then read from the box that these are trick matches made to look like normal

matches, so I believe the match is a trick match. I am now justified to give up my belief that the match is normal. Therefore, I am justified to disbelieve that the match will light when struck. (That is, unless somehow I also have evidence that there is a special trick-match lighting strip that will light trick-matches.)

This familiar kind of case suggests the following series of dynamic principles:

Closure: You are justified to reason from the belief that P and the belief that generally Q whenever P, to a belief that Q.

Defeater: Believing P and some Q that is in general incompatible with P justifies you to reject P.

I use the term “generally” in such principles with its ordinary generic meaning. “Generally whenever P, Q” and “Generally, P and Q are incompatible” can be true even if there are exceptions. Also, “generally” implies “sometimes” and is entailed by our ordinary uses of “always” and “necessarily.” For all of the principles that I formulate in terms of “generally,” I take the principle that results from replacing “generally” with a stronger term like “always” and “necessarily” to also yield a true principle (for more on generics and generalizations with exceptions, see Leslie’s dissertation [Leslie, 2006]).

Match 1 suggests that Closure and Defeater hold only under certain conditions. Having a defeater for my starting points defeats my justification to reason from my starting points to any end points. The following principle captures this fact:

No Defeaters: For all points A, if there is a point C such that you are justified to reason from point C to a rejection of point A, then for no point B are you justified to reason from point A to B.

No Defeater cites a specific condition under which Closure and Defeater no longer hold, even though they generally do. When one principle P states a condition under which another principle P' no longer holds, even though P' generally holds otherwise, I will say that the former has priority over the latter. As a metaphor, consider principles to be related like a hierarchy of courts, where one principle having priority over another is akin to one court having authority over another. Higher courts take the rulings of lower courts and either leave them untouched, use them to make judgments in other cases, or overrule them in light of their own reasoning. We want to combine the three principles, Closure, Defeater, and No Defeaters, in such a way as to explain why you are *not* ultimately justified to reason from your belief that the match is normal to the belief that the match will light when struck. To predict the right result in the Match 1 case, No Defeaters must have priority over both Closure and Defeater, and Defeater must have priority over Closure:

No Defeaters > Defeater > Closure

Given these priority relations, we can derive the desired result in the Match 1 case as follows:

1. Starting Point={b(Match m is normal), b(Normal matches generally light when struck)}.

2. You are justified to reason from 1 to $b(\text{Match } m \text{ will light when struck})$ by Closure.
3. From experience, you add to starting point $\{b(\text{Match } m \text{ is a trick match}), b(\text{trick matches do not light})\}$.
4. You are justified to reason from $b(m \text{ is a trick match}), b(\text{trick matches do not light})$, to reject $b(\text{Match } m \text{ is normal})$ by Defeater.
5. You are no longer justified to reason from $b(\text{Match } m \text{ is normal})$ to $b(\text{Match } m \text{ will light when struck})$ by No Defeater from 4 and 1.
6. By No Defeaters $>$ Defeater $>$ Closure, 4 and 5 stand and 2 is defeated.

Notice that different priority relations among the principles will predict different results. Closure $>$ No Defeater $>$ Defeater expresses the norm that Defeater generally holds, but fails to when the condition in No Defeater is satisfied, yet Closure always holds, even when the condition in No Defeater is satisfied. Closure $>$ No Defeater $>$ Defeater implies that you would be ultimately justified to conclude that the match will light when struck. Closure $>$ Defeater $>$ No Defeater is a norm that implies that you would be ultimately simultaneously justified to reject your belief that the match is normal, *and* to reason from your belief that the match is normal to the belief that the match will light when struck. Since both of these results are incorrect, No Defeaters $>$ Defeater $>$ Closure is likely the correct priority ordering.

Defeater as I have stated it is not wholly adequate and therefore requires slight revision. Suppose I come to believe something inconsistent with something I already believe. From the existing principles, I will be both justified

to disbelieve that the match is normal, and justified to disbelieve that the match is a trick match. By No Defeaters, which has priority over Defeater, I will not be justified to disbelieve either. If I have inconsistent beliefs, by the three principles, my beliefs do not justify me to believe anything else, and my beliefs do not justify me to change my mind about anything I already believe. The three principles imply that the acquisition of defeaters for my beliefs leaves me in a normative tie, where I cannot justifiably reason from my beliefs nor reject any of my beliefs. This result arises because incompatibility is a symmetric relation, so that if P is incompatible with Q, then Q is incompatible with P. On the one hand, this is clearly a good result. Merely believing two things P and Q that are incompatible is not enough to tell us which belief ought to be rejected. It simply tells us to reject that the belief that P and the belief that Q are both true. Incompatible beliefs sometimes create a *normative tie*. On the other hand, it is clear that in the Match 1 case, your experiences of reading the match box justify you to reject your belief that the match is normal, whereas your experiences that justify you to believe that the match is normal do not justify you to reject your belief that the match is a trick match. We have a firm judgment that reading about the match from the match box is better evidence than your experience in individuating matches from trick matches, thereby breaking any kind of normative tie. We might want to separate Defeater into the following two principles:

Inconsistency: You are justified to reason from your belief that P, your belief that Q, and your belief that P and Q cannot both be true, to a rejection of $\{b(P), b(Q), b(P \text{ and } Q \text{ cannot both be true})\}$

true)}.

Defeater*: You are not justified to reason from the rejection of $\{b(P), b(Q)\}$ to a rejection of $b(P)$ unless you inhabit a point C (or possess the states of minds in C) such that you are justified to reason from point C and the rejection of $\{b(P), b(Q)\}$ to a rejection of $b(P)$.

Recall that “rejection of $\{b(P_1)...b(P_n)\}$ ” is the disbelief or agnosticism that all of the propositions in the set are true. In the Match 1 case, reading that the box is a box of trick matches is sufficient to justify you to change your belief that the match is normal. On the other hand, the belief that the match is normal is insufficient to justify you to reject the belief that the match is a trick match because there is no other belief that you have that is sufficient to justify you to reject that the match is a trick match. Similarly, you are not justified by your evidence to reject your belief that, generally, a match cannot be both normal and a trick match. Defeater* must have priority over Inconsistency to yield the desired result. We thus have the following series of principles:

No Defeaters > Defeater* > Inconsistency > Closure

Other Useful Principles

The No-Defeaters principle captures cases in which you come to possess a defeater for something you already believe. In virtue of possessing that defeater, you are no longer justified to reason from that belief to other points. No-Defeaters is a dynamic version of a familiar No-Defeater principles given

by Static Rationalists (for example, [Bergmann, 2005]) to the effect that unjustified beliefs cannot justify beliefs. The dynamic version states that you cannot justifiably reason from beliefs that you have sufficient justification to reject, where such justification arises from other things you believe or disbelieve. However, there are other kinds of defeaters of justification, defeaters that do not necessarily justify you to reject something you already believe. These kinds of defeaters are familiar in cases of “non-montonic” reasoning.

Match 2: I take a match out of the box and examine it, thereby coming to believe that it is a normal match. I am justified to conclude from such beliefs and perceptions that the match will light when struck. However, my friend tells me he had just dipped the match into a glass of water. Knowing how matches work, I am now no longer justified to conclude that the match will light when struck. (That is, unless somehow I have evidence that this is a special water-resistant match.)

I need not, in Match 2, actually believe that the match will light to be justified to believe it, and to subsequently lose that justification upon acquisition of new evidence. Instead, my learning can be cumulative, while my justifications for belief-changes shift. I come to believe that it looks like a normal match, that it is a normal match, and then, that the match has been dipped into water. Such cumulative learning need not involve acquiring justification to reject something I already believe. Instead, one piece of learning justifies me to conclude that the match will light, while a subsequent piece of learning preserves my original belief, while rendering me unjustified to

conclude that the match will light. Match 2 suggests the following general principle about “non-monotonic” cases:

No Defeaters II: You are justified to reason from point A to point B only if there is no C such that you are justified to reason from point C to a rejection of point B.

This No-Defeater principle states a familiar condition of defeasible justification. Believing that P can generally justify you to conclude that a proposition Q is true, but in some cases, believing that P and another proposition P' that is compatible with P is sufficient for you to lose justification to conclude that Q is true.

Can the two No-Defeater principles be unified into a single, simpler principle? In both cases, one piece of learning justifies one to reason to some point B. Another piece of learning renders one unjustified to reason to the same point. The similarities in both cases suggest that we can unify the No Defeater principle as follows:

No Defeaters*: For all points A and B, you are justified to reason from point A to point B only if you are not required to reason from any point to the rejection of point A or point B.

I introduce the notion of a rational requirement here to capture the sense in which the justification we possess can sometimes compel us, on pain of irrationality, to revise our beliefs. It is possible that justification works this way generally, but I leave open the possibility that we can possess justification that permits but does not require a belief-change.

With No Defeaters* in place, we have the following proto-theory of how evidence determines the rationality of belief-change:

No Defeaters* > Defeater* > Inconsistency > Closure

This theory states that, with respect to an arbitrary change of belief from A to B, the change is reasonable if these principles and priority relations do not rule it out as unjustified.

Another range of cases suggest a further kind of principle for dynamic rationality that appears useful. In general, we cannot rationally change our minds willy nilly absent any evidence to do so.

Willy Nilly 1: I catalog all of my beliefs and decide to reject every third one I can remember.

Willy Nilly 2: I catalog all of my beliefs and decide to reject my beliefs about tapestries.

Willy Nilly 3: I catalog all of my beliefs and decide to reject the beliefs about the external world.

Willy Nilly 4-6: I consider all of the propositions I disbelieve and am agnostic about. I decide to believe every third one I can think of, every one about tapestries, and every one about the external world.

The absurdity of the reasoning in these Willy Nilly cases suggests that I cannot change my mind, either in rejecting or adding beliefs, without justification to do so. In none of these Willy Nilly cases am I justified to change

my mind in the ways described. All six Willy Nilly cases have in common the appearance that I am simply changing my mind as I wish, without any reason or evidence sufficient to justify the way in which I manage my beliefs. The Willy Nilly cases suggest the following principle:

Demanding: All changes of mind require justification.

Demanding states that any change of belief that occurs without justification is unreasonable. Is this principle in some sense trivial? It is not, for we can imagine at least one coherent alternative to Demanding. For instance, one might take the view that changes of belief about a special domain can be rational without justification. You might think that one needs evidence to change one's beliefs about the natural world rationally, but no such evidence is required to change one's beliefs about the supernatural world rationally. If someone believes that there are ten flowers in the flower bed, and changes his mind without any justification, he is changing his mind unreasonably. However, if that same person believes that there are exactly ten angels in the flower bed, and changes his mind without any evidence, he is not changing his mind unreasonably. Demanding states that *every* change of belief from a state of mind in which one lacks justification for that change is unreasonable.

Tentatively, then, we have arrived at a simple dynamic theory of rational changes of mind:

Demanding > No Defeaters* > Defeater* > Inconsistency > Closure

The list of principles and their priority relation expresses a norm for rational belief-change, a norm that implies, in every situation, the rational changes of belief for an agent in that situation. The principles are lacking

in specificity as to what constitutes evidence for what, and what a piece of evidence justifies and what it does not. However, it states enough of a general form of a theory to yield testable results, as we have seen in the toy cases we used to generate the theory. Applications of this theory for the epistemology of all-or-nothing belief will be the subject of the first half of the next chapter. In the interim, I will discuss how one might extend such a theory of rational changes of grounds, and then to a degree-theoretic framework. My remarks will be the groundwork for future work.

4.2 The Rationality of Changes of Grounds

In Chapter Three, we noticed that we can sometimes maintain the same beliefs or degrees of belief while changing our reasons for such beliefs or degrees of belief. Some of these changes are more reasonable than others. I also offered a reduction of the rationality of changes of grounds to the rationality of changes of beliefs. We are now in a position to see some of the principles governing rational changes of belief. Thus, rational changes of grounds can be considered a special case of rational changes of belief. A ground for a belief is good when one is justified to reason from the ground to the belief, on the assumption that one does not already have the belief. A ground G is better than a ground G' for a belief just in case it would be more reasonable to reason to the belief from G than G' . A change of grounds for a belief can also be rational whenever the beliefs that constitute the grounds are rationally changed. A change of grounds can be rational whenever you have some higher-order belief that some belief you have now constitutes good

grounds, reasons, or evidence, for another. Finally, a change of grounds can also be rational when you have a defeater for such a higher-order belief. New beliefs and evidence can justify a change of grounds in the same way that they can justify a change of belief. The rationality of changes of grounds is therefore unmysterious, and reduces naturally to the rationality of changes of belief.

Let me now proceed, in even sketchier form, to the theory of rational changes of degrees of belief.

4.3 Rational Changes of Degrees of Belief

We know that people can be more or less certain about the propositions they believe and disbelieve. After completely settling that I believe both P and Q, an open and sensible question remains as to whether I am more certain that P than that Q. Some of the time, we can even recover precise numerical values for these levels of confidence. It also seems clear that, just like beliefs, we can reasonably or unreasonably change our levels of confidence in propositions. Seeing that I have been bitten by a tick gives me justification to raise my degree of belief that I have been infected with the bacteria that causes Lyme Disease. It is not sufficient justification, however, for me to believe that I have Lyme Disease. Seeing a ring-shaped rash at the site of the tick-bite, however, is sufficient justification, and thereby warrants me in seeking medical attention. Bits and pieces of evidence seem to give me justification to raise and lower my confidence; some pieces of evidence are not strong enough to justify full belief, while others are. These comparative normative facts can

best be captured in a theory of rational changes of degrees of belief. Such an account will be the subject of much future work, work that is motivated by a basic understanding of dynamic rationality, one that is not defined as the preservation of static probabilistic coherence over time. Such future work will take very seriously the idea that rational-changes of degrees of belief need not be coherence-preserving. Strengthened Preface considerations and principles like Tautology, Contradiction, and Consequence will take center stage on my view.

4.4 Summary

This chapter offered the beginnings of a theory of rational changes for all-or-nothing belief. The theory differs from existing alternatives primarily in taking the lessons from previous chapters seriously: dynamic rationality is taken as basic, and formal constraints are invoked only piecemeal when they appear appropriate. The theory is incomplete, must be extended to the probabilistic cases, and requires more defense. However, if we assume that it is roughly correct, the applications to problems of epistemology are wide-ranging. In the next chapter, I will illustrate some applications to a range of traditional epistemological problems, the problems surrounding the issue of skepticism.

Chapter 5

Dynamicism Vindicated

5.1 Application 1: Consistency and Rationality

Let us start with the problems of characterizing the principled relationship between logical consistency and rationality from a wholly dynamic point of view.

In rejecting Static Rationalism, we are no longer forced to evaluate as irrational someone with logically inconsistent beliefs, simply because they are logically inconsistent. Preface-type cases and cases of hard to detect inconsistencies prove to be counterexamples to the view that rationality requires consistency. However, we must still make rational sense of some changes of mind given inconsistencies, some changes of mind that are irrational given inconsistencies, and some pieces of reasoning that lead to inconsistencies that are nonetheless rational. The dynamic view accomplishes all of this.

On the theory that I have proposed, you are always justified to reject a set of beliefs that is inconsistent. (Recall that rejecting a set consists in disbelieving that all of the beliefs in the set are true). According to the No Defeaters* principle, you are not justified to conclude the logical consequences of your beliefs when you possess justification to reject the logical antecedents. Thus, you are not justified to reason from inconsistent sets of beliefs to anything else. The dynamic theory makes sense of why Closure can be true without *ex falso quodlibet* being a good principle of reasoning. The dynamic theory also makes sense of why you can have reason or be justified to conclude all of the trivial tautologies and disjunctive additions to your beliefs, without being irrational in failing to believe such things. Because all rationality is dynamic rationality, you are justified to conclude such things, but are not irrational for failing to be in a state of mind that includes all such beliefs.

Also according to my dynamic theory, you are not justified to reject a particular belief B in an inconsistent set simply because it is part of an inconsistent set. In order to be so justified, you must acquire some justification to reject B simpliciter, whatever such justification might turn out to be. Defeater* gives one condition under which you have sufficient justification to reject a single belief in an inconsistent set. There could be more. If none of these conditions is satisfied, you are in a state where you have justification to reject an inconsistent set containing B, but no justification to reject B. By Demanding, you need justification to reject B in order to reject it rationally. Thus, you can be in an inconsistent state of mind but not be justified to reject any particular belief you have. My dynamic theory can therefore make

sense of reasoning and rationality in Preface-type cases. By hypothesis, we are justified in Preface-type cases to reason to each individual belief and to the Preface belief. We are not justified to reason from such an inconsistency to the rejection of any particular belief in our inconsistent set. Nonetheless, we are justified to disbelieve that all of our beliefs are true. Furthermore, the restriction on *ex falso quodlibet* prevents us from being justified to conclude anything and everything from our inconsistent beliefs. In Preface-type cases, all of our reasoning is justified, and maintaining our beliefs is justified, because we do not possess reason to reject any particular belief.

Finally, and most controversially, my particular dynamic theory allows you to justifiably reason from beliefs even when they are simultaneously held with other beliefs that are inconsistent with them. For instance, if point A is inconsistent with point B, but you are not ultimately justified to reject point A, and you are not ultimately justified to reject point B, you can still justifiably reason from each point individually, though not both jointly.¹I allow this possibility because it seems legitimate to reason from your beliefs even when you take a Preface-type stance toward them. Allowing for such a possibility also makes sense of certain controversial cases from the history of science in which scientists seemed to believe inconsistent theories, and recognized them as such. Nonetheless such scientists reasoned from their theories to novel predictions and novel hypotheses about how the world

¹Allowing this possibility does not result in *ex falso quodlibet*. To see why, note that in order to justifiably believe an arbitrary proposition P from A and B, one needs to justifiably conclude P from some point C that one is justified to conclude only *from* a set containing A and B. But by No Defeaters*, one is never justified to reason from a set containing A and B. Therefore, one is not justified to conclude any such C. Thus, one is never justified to conclude an arbitrary P.

works [Norton, 1987], [Kitcher, 1993]. Such scientists seemed to maintain their beliefs in their theories, and did not reason from them to the view that the world is every which way and no way at all. As controversial as such cases are, if they are even possible, we need a theory of rationality that makes sense of them.²I take it to be a benefit of my dynamic theory that it can explain and rationalize such phenomena without resorting to adopting heretical alternative logics.

This controversial aspect of my view has certain consequences that some may find objectionable. Take the following cases:³

1. A and B are mutually defeating. C defeats B. A defeats C.

2. A and B are mutually defeating. A defeats D. D defeats C. C defeats B.

3. A and B are mutually defeating. C defeats D. D defeats C. D defeats B.

My view implies that, in Case 1, you are justified to reject C, but are in a normative tie with respect to A and B. In Case 2, you are justified to reject B and D, but not to reject A and C. In Case 3, you are justified to reject B but not to reject any of the other beliefs. I leave the short derivations for

²The issue is complicated in that there are disputes about whether the theories themselves were genuinely inconsistent, and if so, whether scientists *believe* their theories. All that matters for my purposes is that it is possible for some scientists to end up with inconsistent opinions, even recognize them as such, and nonetheless reasoned from them in a rational way *even if they were not rational in arriving at their inconsistent opinions in the first place*.

³Thanks to Jim Pryor for these schematic cases.

the Appendix. These results all arise because my theory permits reasoning from some point A inconsistent with another point B, provided you are not otherwise justified to reject point A. To some, the correct results ought to be a normative tie in all of these cases. I actually find my results defensible, or at the very least, acceptable given the need to explain and rationalize aspects of scientific practice. However, those that find my such results too troubling can easily replace my No Defeaters* with a more restrictive principle that suits their tastes. For instance, the following no-defeater principle precludes you from justifiably reasoning from any point A that is held simultaneously with a point you recognize to be incompatible with A.

Unpermissive No-Defeaters: For all points A and B, you are justified to reason from point A to point B only if you are not justified to reason from any set of beliefs you have to a rejection of any point containing A.

I do not think Unpermissive No-Defeaters turns out to be defensible. (In fact, it is paralyzing. It follows from this principle that any person with inconsistent beliefs cannot reason from any other belief. There also appears to be no ways of making the principle more permissive that isn't just ad hoc. That's why I prefer my No-Defeaters* to anything like Unpermissive No-Defeaters). Unpermissive No-Defeaters nonetheless has the benefit of being a dynamic explanation of the otherwise puzzling phenomena of Preface-style reasoning. Dynamicism seems to give us the right tools for dealing with rational inconsistencies.⁴

⁴see Afterword for further discussion of this section.

5.2 Application 2: Skepticism, Dogmatism, and Closure

As much as we would hate to admit it, skepticism is what drives epistemology. The history of epistemology is a history of responses to a variety of powerful skeptics, skeptical arguments, and skepticisms. What does Dynamicism have to say about the problems of skepticism? In itself, Dynamicism has nothing to say. The view that all rationality is dynamic is consistent both with the view that none of us are ever rational in our belief-changes, and the view that by and large our changes of belief are rational. However, Dynamicism focuses the skeptical question, the skeptical dialectic, and the dogmatic response to skepticism. It in fact makes dogmatist responses to skepticism more philosophically palatable.

To illustrate, suppose you look out of the window and identify a fence, judging that it is made of Western Red Cedar. A skeptic comes along and asks you to consider the possibility that it is made of reclaimed redwood instead, a possibility that you failed to rule in, and therefore failed to rule out. The skeptic then makes a claim that we seem to accept, namely, that your failure to consider the possibility that the fence is made of redwood renders your judgment that it is made of Red Cedar problematic, unreasonable, unjustified. Now the skeptic claims that for every single one of your beliefs, there is a live possibility that you failed to consider, and therefore failed to rule out, rendering that belief problematic, unreasonable, and unjustified. Your belief that you're reading this sentence right now, your belief that your name is "Joe", your belief that today is sunny, are all beliefs you formed

while failing to rule in and rule out some possibility of delusion, hallucination, dreaming, mass conspiracy, and so forth. The skeptic claims that these possibilities, however remote or improbable, are possibilities that nonetheless render your beliefs unreasonable in exactly the same way as the possibility of redwood renders unreasonable your belief that the fence is made of Red Cedar.

When the skeptic challenges our rationality, the Dynamicist understands the skeptic as offering his own principle of rational belief-change. For the skeptic, you need to rule out every logically possible alternative in order to justifiably reason from one belief to another. Thus, you need to rule out every logically possible alternative to the fence's being Red Cedar, in order to be reasonable in concluding that it is made of Red Cedar. According to the skeptic, you need justification to reject some things before you have justification to add some things to your stock of beliefs.

The Dogmatic Dynamicist can offer an alternative to the skeptical principle. For the Dogmatic Dynamicist, you don't need to have justification to first reject alternatives before you have justification to add beliefs. You can have justification to add some beliefs without first needing justification to disbelieve alternatives to such beliefs ⁵. For the Dogmatic Dynamicist, this norm can apply across the board, so there is no charge of epistemological hypocrisy, where some norms apply to some cases and not others. In the fence case, the nonskeptical Dynamicist can claim that we are reasonable in concluding that the fence is Red Cedar even if we have not ruled out the possibility of redwood. Maybe finding out that it could be made of redwood

⁵For a version of such a view see [Pryor, 2000]

would require us to retract that belief. But just because we need to retract it later does not mean that our original piece of reasoning was problematic or unjustified. We find out things all the time that require us to change our mind. Changing our minds on new information doesn't show that we were unreasonable in the way we originally came to our beliefs! So the Dogmatic Dynamicist has one principle that applies to all belief-changes, to the redwood and Red Cedar case, as well as to the case that you believe that you are reading this sentence right now. In effect, the Dogmatic Dynamicist accepts the principle that you can justifiably reason to a conclusion without needing to rule out every logical possible alternative to that conclusion.

Two problems arise for Dogmatic Dynamicism. The first is the problem of Closure Principles. The second is the issue as to how much Dogmatic Dynamicism ultimately helps with the problem of skepticism. Let us take both in turn.

5.2.1 Closure

A closure principle for justification states the relationship between consequence- the logical relationship between propositions- and rational justification- a normative property of beliefs or belief-changes. Standard discussions of closure principles for justification either assume that such principles characterize static rationality, or that such principles are exception-free. The simplest principle states that you are not fully rational unless you currently believe all of the logical consequences of your current beliefs. As we have seen in Section 1, such static closure principles are problematic. The simplest dynamic formulation of a closure principle states that you always have reason

or justification to believe all of the logical consequence of your beliefs. As we have seen, such principles cannot be true either, for *ex falso quodlibet* cannot be a good principle of reasoning. My dynamic Closure Principle is neither static nor exception-free, rules out *ex falso quodlibet*, and rules out justification to believe the logical consequences of your defeated beliefs. Still, there are remaining problems with closure principles, some of which my theory can solve, some of which it cannot. Yet, even those disputes that it cannot solve can be better illuminated.

Consider the following arguments [Dretske, 1970a], [Cohen, 2002].

1. That is a zebra.
 2. If that is a zebra, then that is not a cleverly disguised mule.
 3. That is not a cleverly disguised mule.
-
1. o is red
 2. If o is red, then o is not white and illuminated by red-colored lights.
 3. o is not white and illuminated by red-colored lights.

Deniers of Closure claim that we are not justified in general to believe the conclusions of these arguments even if we are justified to believe the premises. At the very least, such arguments differ in some normative respect with ordinary arguments from the following:

Ordinary Closure

1. The gas gauge in my car is reading empty.

2. If the gas gauge in my car is reading empty, then I need to stop for gas.
3. I need to stop for gas.

Though such cases have been standardly discussed in terms of knowledge rather than justification, some epistemologists nonetheless consider the reasoning that these arguments represent to undermine closure principles for justification generally [Dretske, 1970a, Dretske, 1970b]. According to standard closure principles, if α is a valid argument, then you are justified to believe the premises only if you are justified to believe the conclusion. Cases like the above arguments purport to be counterexamples.

One can rescue standard closure principles by denying that we are in fact justified to believe the premises. Such rescue attempts seem to lead directly to skepticism about justification. One can also rescue closure principles by denying that there is anything wrong with the reasoning that these arguments represent, contrary to intuition [Klein, 1981]. Finally, one can rescue intuition, as well as deny skepticism, by diagnosing such arguments as cases of “transmission failure” [Wright, 1985], [Davies, 2000]. In cases of transmission-failure, closure holds, so you are always justified to believe the conclusion when you are justified to believe the premises. But the arguments above show that you are not justified to believe the conclusion *on the basis* of your justification to believe the premises. So while such arguments are no counterexample to closure, they differ from Ordinary Closure in failing to “transmit” justification from premise to conclusion.

On the dynamic view, the question regarding these cases reduces to the question of whether you can be justified to reason from beliefs in (1) and (2)

to a belief that (3) in the above arguments. On the dynamic view, there is no debate over whether one can simultaneously justifiably believe the premises without being justified to believe the conclusions. Once we reject the notion of static rationality, there is no issue as to what you can rationally and justifiably believe simultaneously; there is only the issue of what pieces of reasoning are rational or justified. Moreover, formulated in wholly dynamic terms, there clearly seems to be something wrong with reasoning from the beliefs that (1) and (2) to the belief that (3). “Transmission failure” is, at the very least, a good label for the phenomena. On the dynamic view, we can also make sense of how competing views need to make sense of the phenomena. Deniers of Closure claim that it is false that, in general, you are justified to reason from your beliefs to their logical consequents. They are therefore required to explain cases of why Ordinary Closure represents a justified way to reason. Defenders of Closure claim that pieces of reasoning in these cases are always justified. They are therefore required to explain away the appearance of dubiousness in cases like the above arguments.

On my dynamic theory, Deniers of Closure can satisfy their explanatory requirements in a way that should satisfy a Defender of Closure. Deniers of Closure need not really *deny* closure in the sense of rejecting a Closure Principle as part of a total dynamic theory. Rather, they can and should endorse some other principle that has priority over Closure. Such a view allows the Denier of Closure to easily explain cases like Ordinary Closure. One such principle that Deniers of Closure can adopt is:

Circularity([Davies, 2000], [Harman and Sherman, 2004]): If you are only justified in reasoning to the belief that P in virtue of

believing or assuming that Q, then you are not justified to reason from the belief that P to the belief that Q.

Accepting that Circularity has priority over Closure does not amount to a denial of Closure. Thus Defenders of Closure can consistently accept this position advocated by a Denier of Closure, and the debate can be closed. Closure holds in the sense that it is a real and useful principle in a full theory of dynamic rationality. But just as it is defeated in cases where logical antecedents are defeated, it is also defeated in cases like the Zebra argument.⁶

The dynamic view of epistemic rationality clarifies the debate over closure principles and focuses the phenomena to be explained. As a result, both Defenders and Deniers of Closure end up with views that the other side can consistently accept. From the dynamic point of view, there need not be a debate over closure principles and transmission-failure as such. We need only debate the acceptability of the further dynamic principles we need to explain the phenomena we find in the Zebra and Easy-Knowledge arguments.

5.2.2 Dogmatism Revisited

Does Dogmatic Dynamicism really rescue us from skepticism? Even if we accept the dogmatic dynamic principle, it remains open as to what counts

⁶A persistent Defender of Closure [Klein, 1981] will not be satisfied with my foregoing attempt to alleviate the dispute about closure principles for justification. Such a Defender of Closure will insist that there is nothing wrong with reasoning from (1) and (2) to (3) in Skepticism and Easy Knowledge cases, despite the appearance of illegitimacy. I think it is imperative for anyone toeing this line to offer some explanation of the apparent normative difference between the reasoning in Skepticism and Easy Knowledge and the reasoning in Ordinary Closure.

as a defeater for our beliefs, even if they are rationally formed. If skeptical arguments, once encountered, constitute sufficient evidence for us to revise our ordinary beliefs, yet we maintain them, we would be guilty of a rational failure. Nothing in the Dogmatic Dynamic Principle says anything about whether skeptical arguments justify a change of belief. Nothing about the Dogmatic Dynamic Principle states that you must, or must not, or are merely permitted to, revise your beliefs in the face of a skeptical request to consider logically possible alternatives to your beliefs. So is there really Dynamicist solution to skepticism?

The aim of this section has been to show how Dynamicism reformulates a certain long-standing philosophical problem, not to solve that problem. We do not solve the problem of skepticism by holding that only changes of belief are rationally evaluable. Rather, we know how to formulate the solutions, we show how proposed solutions become philosophically palatable or unpalatable when formulated dynamically, and we show exactly what the debate ought to be about. In the case of Dogmatism versus Skepticism, the debate is over a Dogmatic versus a Skeptical dynamic principle, and over what kinds of arguments provide sufficient defeaters for beliefs, whether they have been arrived at rationally or not. These are quite substantive debates over the correct principles of dynamic rationality. One can remain neutral about them and still acknowledge that Dynamicism clarifies the debate.

5.3 Application 3: Conservatism in Epistemology

Skepticism always seems to loom whenever a Static Rationalist proposes that your beliefs are unjustified unless proven otherwise. Proving otherwise turns out to be very difficult. Since Quine [Quine, 1951], a less demanding alternative has emerged: your beliefs are fully justified unless proven otherwise. One slogan of such a view is that you are justified in believing what you believe because you believe it. The corresponding dynamic view, also endorsed by some Conservatives, is that you are justified in continuing to believe as much of what you currently believe because you currently believe it. Why? Because you are currently justified in believing it. Conservatism seems to be a good position for the avoidance of skepticism about rational justification. If all of my current beliefs are justified because I believe them, then no skeptic can successfully argue in a non-question-begging way that I should abandon all of my beliefs because they are unjustified. The skeptic is required to show, from my own standards for irrationality, that my beliefs are irrational. But the mere fact that I believe a proposition already gives me justification for believing it. Conservatism proves to be a nice way to turn the table on the skeptic. (Besides Quine, prominent conservatives include [Sklar, 1975], and [Harman, 1986, Harman, 1999]).

Conservatism can be a static doctrine about what you are justified to currently believe and why, or as a hybrid doctrine about how you are justified to change your mind given a justified starting point. But Conservatism so understood is susceptible to serious charges: why should the fact that

I believe something justify me in believing it or continuing to believe it? Why am I not justified in believing everything that George Bush believes instead? (See [Christensen, 1994, Christensen, 2000].) Thus, Conservatism faces charges of bias in favor of oneself. If I favored the opinions of Jones and not Smith simply because they are Jones' opinions, I would be guilty of bias against Smith. Similarly, if I favor my own opinions over those of Jones simply because they are mine, then I am guilty of partiality to my own opinions. And isn't partiality a paradigm case of unreasonableness?

As a hybrid thesis, the standard formulation of Conservatism runs into similar charges of bias. If I changed my mind to maximally preserve Jones' beliefs and not Smith's simply because they are Jones' beliefs, then I am unreasonably biased against Smith. Similarly, if in changing my mind, I seek to maximally preserve my opinions simply because they are mine, then I am unreasonably partial to myself.

If the kind of partiality advocated by Conservatism is permissible, that is, if everyone is perfectly justified in being self-partial, then justification seems too cheap. Too many people with completely conflicting opinions can be equally justified. Thus, Conservatism faces a serious charge of relativism. Conservatism appears to be unable to adequately sort the rational from irrational agents, so the charge goes. These charges of partiality and relativism seem to undermine Conservatism as a viable epistemological position.

A wholly dynamic understanding of rationality and justification has a way of settling such a dispute. There is something right and something wrong about Conservatism. Conservatism is right in that my starting points are in some sense epistemically privileged. Conservatism is wrong qua doctrine of

static rationality, and wrongly formulated as a doctrine of dynamic rationality. Rejecting the notion of static rationality allows us to reap the benefits of Conservatism without facing such apparent problems. On anyone's theory of dynamic rationality, either (i) changes of mind require justification or (ii) preservation of beliefs requires justification. I endorse *Demanding*, which is the view that changes of mind require justification. I will call the alternative view *Crazy*. We have seen that Demanding is independently motivated. I am simply not justified to change my mind willy nilly. It follows that Crazy is not a viable position. The alternative to maintaining your beliefs is to give them up, either by becoming agnostic about or disbelieving the propositions you believe. Crazy amounts to the position that changing your mind willy nilly is permissible absent any justification for not doing so. But it is clear from our judgments on even the simplest of cases that Crazy is false.

A dynamic theory of rationality also precludes a mixed position to the effect that changes of mind *and* maintaining beliefs both require justification. Starting with my beliefs, I either keep them or change them. If I need to be justified in keeping them, then I am by default permitted to abandon them. If I need to be justified in abandoning them, then I am by default permitted to keep them. It is incoherent to think that I need justification to maintain my beliefs and justification to abandon them.

As a result of looking at rationality from a wholly dynamic point of view, we can see that we should not formulate Conservatism as the view that you are justified to believe everything you currently believe *because* you currently believe it. A dynamic theory of rationality thus escapes charges of partiality or bias. Rationality requires everyone to be justified in changing his be-

liefs. Our reasons for belief-change can be as agent-neutral, universal, and objective as we would like. Whatever justifies you in reasoning from point A to point B also justifies me. All of the unjustified ways for me to change my mind are unjustified for you. Thus, there is no partiality or relativism. A purely dynamic theory of epistemic rationality can reap the benefits of Conservatism without the costs.

Left over from the dispute between conservatives and nonconservatives is the issue of what counts as a reason to change your mind. This dispute is a dispute about the nature and extension of reasons for belief-change, an issue that everyone in epistemology needs to settle. A skeptic might claim that any mere possibility, however remote, that not P gives everyone reason to give up the belief that P. Others might claim that any psychologically real sense of doubt that P justifies you to give up the belief that P. However, there is no special dispute about whether all or some or none of your current beliefs are justified in virtue of X, Y, or Z. In particular, there is no special dispute as to whether I am justified in currently believing that P because I believe it. Without a notion of static rationality to characterize, there is no special controversy over Conservatism, just a general issue concerning what constitutes sufficient reason for belief-change.

5.4 Further Remarks on the Relationship between Logic and Rationality

In section 3.3.1, I gave the beginnings of a theory as to the relationship between logic and rationality, a relationship that is difficult to characterize

from a static point of view. Logic is the study of certain modal relationships between propositions, to the effect that some propositions must be true if others are, and some propositions cannot be true if others are. The norms for dynamic rationality tell us that crucial to the conditions under which we are justified to change our mind are our beliefs about when some things are generally true or false given the truth or falsity of other things. Thus, the norms for reasoning when you recognize logical relationships between propositions will be a special case of norms for reasoning when you recognize generic dependencies and incompatibilities between propositions. In this sense, the logical facts play the same role as general facts about the relationship between propositions (or facts, events, states of affairs). When you think that one proposition is necessarily true given another, or necessarily false given another, your belief gives you justification to reason according to the principles of dynamic rationality, just like any other belief about the relationship between propositions (or facts, or events, or states of affairs).

Many people think that logical facts play some special role in rationality that other kinds of facts do not. On the dynamic view, there is one way in which believing a dependency or incompatibility to be logical differs from believing it to be generic. When we recognize a relationship between our beliefs to be logical, we can lose justification for adding beliefs only when we are justified in changing existing beliefs. However, when a dependency or incompatibility relationship is strictly generic, we might lose justification for belief change without being justified to give up anything we already believe. These facts capture the sense in which there is a distinction between what has been called “*montonic*” and “*nonmonotonic*” reasoning. The distinction is not the

distinction between defeasible and indefeasible reasoning. The conclusions of pieces of defeasible reasoning are beliefs that you can be later justified in rejecting. The conclusions of indefeasible reasoning are beliefs that you cannot be later justified in rejecting. I endorse a version of fallibilism according to which all reasoning is defeasible. Anything I currently believe I could under some circumstances be justified to give up. Still there is a distinction that lines up with the logical distinction between monotonic and non-monotonic relations. One kind of defeasible reasoning, call it “shmonotonic”, is the kind of reasoning that can be defeated only by defeaters of things we already believe. Another kind of reasoning, call it “non-shmonotonic” defeasible reasoning, is the kind of reasoning that can be defeated by the acquisition of new information that defeats justification without necessarily defeating anything we already believe.

In one sense, the relationship between knowledge of logical facts and rationality is simply a special case of the relationship between our knowledge of general facts and rationality. In another sense, the relationship between logic and rationality is special.

5.5 Side benefits of a Dynamic Theory of Rationality

There is one final benefit of a purely dynamic theory of rationality. Purely dynamic theories of epistemic rationality are better suited to be unified with theories of rational nondoxastic reasoning, and with theories of rational action. Not all of the forms of reasoning from one point to another involve

some form of doxastic acceptance or rejection. I can reason *from* a proposition taken as hypothesis or a hypothetical. I can reason *to* what might be the case given the hypothesis, without accepting either the hypothesis or the conclusions. Such hypothetical reasoning involves reasoning from propositions on the assumption of their truth or falsity, to what is possible, what is necessary, and what is probable given those propositions. I ultimately think that even these cases are cases of doxastic reasoning, where your conclusions are *conditional* beliefs or probabilities. But not everyone needs to accept this claim. For those who think that there is nondoxastic reasoning, my dynamic theory of doxastic rationality is beneficial in that it can be unified with a dynamic theory of hypothetical rationality using the same family of foundational concepts. The same holds true for justified changes of intentions or plans.

A dynamic view of rationality also lines up better with theories of rational action in taking something essentially active rather than stative to be the object of normative evaluation. Of course, saying that belief-change is active, and not stative, does not in any way imply that it is an *action* subject to voluntary control. I do not pretend that a dynamic view of epistemic rationality is nearly sufficient to unify theories of epistemic and practical rationality. However, the prospects look better when the primary objects of epistemic appraisal no longer differ radically in ontological category from the primary objects of practical appraisal.

5.6 Conclusion

Exactly how revisionary is the wholly dynamic conception of epistemic rationality? In a sense, it is not very revolutionary. If ordinary everyday epistemic evaluation is the evaluation of reasoning and changes of mind, then all along, we have latched onto the right phenomena. It is only our epistemological theories that have carved things up the wrong way. Somehow in the process of philosophical epistemology, we tried to look at reasoning and evaluation of such reasoning as a wholly static process. We asked whether our current beliefs have such and such a property, in virtue of which they are rational. Rationality became tied to the beginning and end products of reasoning, rather than to the process itself. However, as we have learned from the history of science, carving things up the wrong way does not render a theory fruitless. I contend that all of the insights of contemporary epistemology can be rescued from a wholly dynamic point of view. Static Rationalism was a successful theoretical enterprise to the extent that it rightly identified the normative epistemic judgments that constitute data for our normative theories. The task of epistemology from the dynamic point of view will be to unify and explain the phenomena by carving things up correctly. I have illustrated what the epistemological project will look like for all-or-nothing belief changes. My theory of rational changes of degrees of belief will further illustrate the power of understanding rationality from a dynamic point of view.

Afterword: Hard Problems

Keeping Track of Reasoning

The following objection has been posed by both Gil Harman and Gideon Rosen. Suppose I believe P and I believe Q, which is inconsistent with P. I reason from P to P or R. Days or years later, upon forgetting how I reasoned to P or R, I reason from P or R, and my belief that Q, to R. Of course, R is arbitrary. Thus, if this line of reasoning is justified, then reasoning to anything can be justifiable so long as I believe two inconsistent propositions, and my line of reasoning from P and Q to R is separated by enough time that I lose track of how I reasoned to P or R. Of course, this would be an awful result.

Do I have the resources to block the reasoning from P to an arbitrary R? I do. You believe Q but you do not believe that P is false, and you do not disbelieve that P. You cannot reason from Q to a rejection of P. So you cannot rule out P and to reason from P or R to R. All of this is consistent with the theory from Chapter Four.

But the objection can be pressed even further. What if Q just is \tilde{P} . So I believe that P and I believe that \tilde{P} . What principles preclude me from

justifiably reasoning from P to P or R today, and then from \tilde{P} and P or R, to R tomorrow, when I forgot my line of reasoning from P to P or R?

To this objection I propose three answers, all three of which are promising in their own way, but none of which I yet fully endorse.

As Gideon Rosen points out to me, it is easy to resolve this problem if I require agents to keep track of their belief-changes, if not generally, then at least in cases in which they reason from inconsistent opinions. If I require agents to keep track of their belief-changes, then it follows that any agent who reasons in the manner described above violates a principle of dynamic rationality, namely, the principle that you cannot reason from a point you are justified in rejecting. However, the cost of this response is that it adds a psychologically unrealistic constraint on justifiable reasoning. As Harman points out to me, we almost never keep track of our reasoning. Why should that psychological fact preclude us from ever reasoning rationally? Of course, reasoning from opinions that are inconsistent with other opinions is a tricky and delicate matter. It might be completely plausible that rationality requires that we keep track of our ways of reasoning whenever we aim to reason from *inconsistencies* without thereby requiring that we keep track of our reasoning generally. Keeping track of our own reasoning need not be a general requirement for dynamic rationality, it only needs to be a requirement in very delicate situations.

Another response to the objection is to claim that one can still count as having reasoned both from one's belief that P and one's belief that \tilde{P} in one's reasoning to R, even if one does not remember reasoning from P to P or R. The fact that P or R is the result of one's reasoning from P is

sufficient to preclude it from being a candidate belief to join with my belief that \tilde{P} to get to R. I do not have to remember or know that P or R is the result of reasoning from P to count as so reasoning as far as the norms are concerned.⁷ If this response is correct, then I can block the objection and the principles in Chapter Four are still adequate. However, the response completely defeats the spirit of other views I hold, and even defend, in other parts of the dissertation. Nonetheless, it is a promising line of thought for Dynamicists who do not share my idiosyncratic views on other matters.

The final response I can imagine to this objection seems a bit more ad hoc, but nonetheless promising. I can endorse a principle that believing that P rationally precludes you from using the formal contradictory \tilde{P} as an essential step in any line of reasoning. When I reason from \tilde{P} and P or R, one concludes that R only by ruling out that P. It is essential to the line of reasoning that P be rejected. This principle says that whenever one must essentially rule out P in a line of reasoning, one must not also believe that P. This principle prevents one from reasoning from \tilde{P} if you also believe that P, and also from P if you also believe that \tilde{P} . But it still allows one to reason from P and any proposition Q inconsistent with P, so long as Q is not of the form \tilde{P} , and as long as one does not reason from Q to \tilde{P} on the way to some other conclusion. Such a principle restricts the kind of inconsistent opinions from which you are allowed to justifiably reason, without restricting all cases of reasoning from inconsistent opinions. This response will require independent motivation and defense, but it appears promising to me as a

⁷The fact that I arrived at my belief that P or R from my belief that P might be inaccessible to me as I reason, but it nonetheless make a normative difference to what I have justification to conclude.

response to this problem.

Arational Changes of Doxastic Attitudes

For Reliabilism and other externalist theories of justification like the theories of Alvin Goldman and Alvin Plantinga, epistemic rationality is a property of transitional states of mind like belief-generating rules, methods, processes, or procedures. In one sense, then, these theories are fully compatible with Dynamicism; they take an extra step and offer controversial characterizations of the properties that make belief-changes rational, but they are Dynamicist in spirit. But there is a serious difference between my Dynamicism and theories like Reliabilism. For me, not every piece of mental activity that results in an addition, deletion, or revision of a belief is rationally evaluable. Some are *arational*, neither rational nor irrational. But every such piece of mental activity can be said to have or not have some external property like reliability or proper function. For instance, I think that forgetting that P is a mental activity that leads to a change in one's doxastic state. But I do not consider this kind of change of belief either reasonable or unreasonable in the epistemic sense. Similarly, I do not consider cases in which one is force-fed a "belief that P" pill a case of a reasonable or unreasonable change of belief.

However, what should I think about cases in which you voluntarily take the belief-that-P pill, or take a pill that makes you forget that P? If these are rationally evaluable cases of belief-change, (and by this I mean evaluable for *epistemic* rationality,) then how do they differ from being force-fed such a pill? If they are not rationally evaluable, then how do they differ from clear

cases of rationally evaluable changes of belief, like believing that P because you desire that P, or giving up the belief that P because you desire that not P?

I believe there is a distinction between changes of beliefs that are reasonable and unreasonable, and changes of doxastic states that are neither. Yet I cannot begin to characterize it. Doxastic voluntarists have an answer: they can say that only the changes under one's voluntary control count as rationally evaluable changes of belief. But I do not want to be a doxastic voluntarist, so I lack any principled way to distinguish between rationally evaluable and arational changes of belief. I believe such a characterization is needed for Dynamicism to get off the ground.

Rational Transitions

Beliefs are not the only things we reason from. We add and revise beliefs from perceptions, imagination, and even assumptions. What are the other mental states from which we reason, and what are the principles governing the rationality of such reasoning? Probably most of our beliefs are conclusions we draw from what we see, hear, smell, drink, and the like, rather than other things we believe. The theory I propose therefore only covers a small class of cases.

Evidence

What justifies me to conclude that global warming is reversible, and what justifies me to give it up? Given a certain starting point, some learned propositions will justify me to raise my confidence that global warming is reversible, others will justify me to lower it. But which ones? Is it the proposition that carbon emissions can be solidified and stored underground? The principles in Chapter Four are nowhere near specific enough to give answers to these questions. But aren't such questions delegated to scientists and not philosophers? No. A scientist might propose that two facts, Q and R, support P, and two other facts, Q' and R' disconfirm P. But the epistemologist must tell us what it is about Q and R that justifies concluding that P from the scientist's starting point, and what about Q' and R' justifies disbelieving P from the scientist's starting point. What do Q and R have in common, and what do Q' and R' have in common, that makes them the types of justification for belief-change that they are. In effect, what constitutes evidence for and against? This is the most substantive piece of epistemology that must be added to Dynamicism. Nothing I have said in Chapter Four gives a fully satisfactory answer about this matter.

Appendix: Derivation of Cases in 3.3.1.

Case 1: A and B are mutually defeating. C defeats B. A defeats C.

Result: You are justified to reject C. You are not justified to reject A. You are not justified to reject B.

Proof: By Inconsistency, you are justified to reject $\{A,B\}$. Since there is no D such that you are justified to reason from D to reject A, and no D such that you are justified to reason from D to reject B, you are not justified to reject A, and not justified to reject B. By hypothesis, A defeats C. By No-Defeaters*, you are not justified to reason from C to reject B. So you are not justified to reject B.

Case 2: A and B are mutually defeating. A defeats D. D defeats C. C defeats B.

Result: You are justified to reject D. You are justified to reject B.

Proof: Since A and B are mutually defeating, by Inconsistency, you are justified to reject $\{A,B\}$. There is no E such that you are justified to reason from E to reject A. By Defeater*, you are not justified to reject A. By hypothesis, A defeats D, so you are justified to reason from A to reject D. By No Defeaters*, you

are not justified to reason from D to reject C. By hypothesis, C defeats B.

Case 3: A and B are mutually defeating. C and D are mutually defeating. D defeats B.

Result: You are justified to reject B.

Proof: Since A and B are mutually defeating, by Inconsistency, you are justified to reject {A,B}. The same reasoning applies to C and D. By Defeater*, you are not justified to reject A, you are not justified to reject C, and you are not justified to reject D. But by Defeater*, since there is a D such that you are justified to reason from D to reject B, you are justified to reject B.

Bibliography

- [Arntzenius, 2003] Arntzenius, F. (2003). Some problems for conditionalization and reflection. *Journal of Philosophy*, C(7):356–371.
- [Bergmann, 2005] Bergmann, M. (2005). Defeaters and higher-level requirements. *Philosophical Quarterly*, 55:419–436.
- [Bonjour, 1985] Bonjour, L. (1985). *The Structure of Empirical Knowledge*. Harvard UP.
- [Christensen, 1994] Christensen, D. (1994). Conservatism in epistemology. *Nous*, 28:69–89.
- [Christensen, 2000] Christensen, D. (2000). Diachronic coherence versus epistemic impartiality. *Philosophical Review*, 109:349–371.
- [Christensen, 2005] Christensen, D. (2005). *Putting Logic in its Place*. Oxford UP.
- [Cohen, 2002] Cohen, S. (2002). Basic knowledge and the problem of easy knowledge. *Philosophy and Phenomenological Research*, 65:309–329.
- [Davies, 2000] Davies, M. (2000). Externalism and armchair knowledge. In Peacocke, P. B. . C., editor, *New Essays on the A Priori*. Oxford UP.
- [Dretske, 1970a] Dretske, F. (1970a). Epistemic operators. *Journal of Philosophy*, 67:1007–1023.
- [Dretske, 1970b] Dretske, F. (1970b). The pragmatic dimensions of knowledge. *Journal of Philosophy*, 67:1007–1023.
- [Feldman and Conee, 1985] Feldman, R. and Conee, E. (1985). Evidentialism. *Philosophical Studies*, 48:15–34.

- [Gärdenfors, 1988] Gärdenfors, P. (1988). *Knowledge in Flux*. MIT Press.
- [Gillies, 2003] Gillies, A. (2003). Two more dogmas of belief revision: Justification and justified belief change. *Unpublished manuscript*.
- [Glymour, 1980] Glymour, C. (1980). *Theory and Evidence*. Princeton UP.
- [Goldman, 1979] Goldman, A. (1979). What is justified belief? In Pappas, G., editor, *Justification and Knowledge*, pages 1–23. Kluwer Academic Publishers.
- [Hájek, 2003] Hájek, A. (2003). What conditional probability could not be. *Synthese*, 137:273–323.
- [Hansson, 1999] Hansson, S. O. (1999). *A Textbook of Belief Dynamics*. Kluwer Academic Publishers.
- [Harman, 1986] Harman, G. (1986). *Change in View*. MIT Press.
- [Harman, 1999] Harman, G. (1999). Rationality. In *Reasoning, Meaning, and Mind*. Cambridge UP.
- [Harman and Sherman, 2004] Harman, G. and Sherman, B. (2004). Knowledge, assumptions, and lotteries. *Philosophical Issues*, 14:492–450.
- [Horwich, 1982] Horwich, P. (1982). *Probability and Evidence*. Cambridge UP.
- [Howson and Urbach, 1993] Howson, C. and Urbach, P. (1993). *Scientific Reasoning: The Bayesian Approach*. Open Court.
- [Jeffrey, 1988] Jeffrey, R. (1988). Conditioning, kinematics, and exchangeability. In Skyrms, B. and Harper, W., editors, *Causation, Chance, and Credence*. Kluwer Academic.
- [Jeffrey, 1965] Jeffrey, R. C. (1965). *The Logic of Decision*. New York:McGraw-Hill.
- [Kaplan, 1996] Kaplan, M. (1996). *Decision Theory as Philosophy*. Cambridge UP.
- [Kelly, 2002] Kelly, T. (2002). The rationality of belief and some other propositional attitudes. *Philosophical Studies*, 110:163–196.

- [Kelly, 2003] Kelly, T. (2003). Epistemic rationality as instrumental rationality: A critique. *Philosophy and Phenomenological Research*, 66(3).
- [Kitcher, 1993] Kitcher, P. (1993). Review of Levi's *The Fixation of Belief and its Undoing*. *Journal of Philosophy*, 90(8).
- [Klein, 1981] Klein, P. (1981). *Certainty: A Refutation of Skepticism*. University of Minnesota Press.
- [Korcz, 1997] Korcz, K. A. (1997). Recent work on the basing relation. *American Philosophical Quarterly*, 34(2):163–196.
- [Leslie, 2006] Leslie, S.-J. (2006). *Generics, Cognitions, and Comprehension*. Princeton University PhD. Dissertation.
- [Levi, 1987] Levi, I. (1987). The demons of decision. *Monist*, 70:193–211.
- [Lewis, 1986] Lewis, D. (1986). Probability of conditionals and conditional probability. In *Philosophical Papers II*. Oxford UP.
- [Maher, 1992] Maher, P. (1992). Diachronic rationality. *Philosophy of Science*, 59:120–141.
- [Makinson, 1965] Makinson, D. (1965). The paradox of the preface. *Analysis*, 25:205–207.
- [Norton, 1987] Norton, J. (1987). The logical inconsistency of the old quantum theory of black body radiation. *Philosophy of Science*, 54(3):327–350.
- [Pollock, 1986] Pollock, J. L. (1986). *Contemporary Theories of Knowledge*. Rowman & Littlefield.
- [Pollock, 1995] Pollock, J. L. (1995). *Cognitive Carpentry*. MIT Press.
- [Pryor, 2000] Pryor, J. (2000). The skeptic and the dogmatist. *Nous*, 34:517–549.
- [Quine, 1951] Quine, W. (1951). Two dogmas of empiricism. *Philosophical Review*, 60:20–43.
- [Rott, 2000] Rott, H. (2000). Two dogmas of belief revision. *Journal of Philosophy*, 97(9):503–522.

- [Sklar, 1975] Sklar, L. (1975). Methodological conservatism. *Philosophical Review*, 84:374–400.
- [Skyrms, 1987] Skyrms, B. (1987). Coherence. In Rescher, N., editor, *Scientific Inquiry in Philosophical Perspective*. University of Pittsburgh Press.
- [Skyrms, 1993] Skyrms, B. (1993). A mistake in dynamic coherence arguments? *Philosophy of Science*, 60(2):320–328.
- [Smith, 1995] Smith, M. (1995). Internal reasons. *Philosophy and Phenomenological Research*, 55:109–131.
- [Tennant, 1994] Tennant, N. (1994). Changing the theory of theory-change. *British Journal for the Philosophy of Science*, 45:865–897.
- [Unger, 1975] Unger, P. (1975). *Ignorance*. Oxford UP.
- [vanFraassen, 1995] vanFraassen, B. (1995). Fine-grained opinion, probability, and the logic of belief. *Journal of Philosophical Logic*, 24:349–377.
- [Wright, 1985] Wright, C. (1985). Facts and certainty. *Proceedings of the British Academy*, 71:429–472.