UC Irvine
UC Irvine Electronic Theses and Dissertations

Title
Commitments of Rationality

Permalink
https://escholarship.org/uc/item/9c4252h9

Author
Sanders, Adam Christopher

Publication Date
2019

Copyright Information
This work is made available under the terms of a Creative Commons Attribution License, availalbe at https://creativecommons.org/licenses/by/4.0/

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

UNIVERSITY OF CALIFORNIA,
IRVINE

Commitments of Rationality

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Philosophy

by

Adam Christopher Sanders

Dissertation Committee:
Professor Annalisa Coliva, Chair
Professor Margaret Gilbert
Professor Duncan Pritchard

2019

© 2019 Adam Christopher Sanders

DEDICATION

To the memory of my mother, Linda, who was strong enough
to raise and love such a rowdy and wild child.
I could not have started this journey
without her unequaled wisdom and guidance.

To my fiancé, Brooke, who is crazy enough to marry me
and put up with my perpetual nonsense.
I could not have finished this journey
without your immense love and support.

TABLE OF CONTENTS

	Page
LIST OF FIGURES	iv
LIST OF TABLES	v
ACKNOWLEDGMENTS	vi
CURRICULUM VITAE	vii
ABSTRACT OF THE DISSERTATION	xi
INTRODUCTION: Committed Rationality and Understanding	1
CHAPTER 1: A Taxonomy of Cognitive Endorsements	10
CHAPTER 2: Cognitive Endorsements as Commitments	63
CHAPTER 3: Semantic Commitments	109
CHAPTER 4: Doxastic Normative Tension and the Primacy of Cognitive Commitments	142
CHAPTER 5: Robust Epistemic Rationality Through Commitments and Requirements	169
CHAPTER 6: Species of Understanding	202
CONCLUSION: New Foundations of Rationality and Understanding	235
REFERENCES	240
APPENDIX A: A Diagram of Acceptance	246
APPENDIX B: A Diagram of Supposition	247

LIST OF FIGURES

		Page
Figure 1	The Rational Axes of Cognitive Endorsements	62
Figure 2	The Different Weights View	146
Figure 3	A Diagram of Acceptance	246
Figure 4	A Diagram of Supposition	247

LIST OF TABLES

		Page
Table 1	The Fundamental Features of Belief	60
Table 2	The Fundamental Features of Acceptance	60
Table 3	The Fundamental Features of Supposition	60
Table 4	The Fundamental Features of Presupposition	61
Table 5	The Fundamental Features of Hypothesizing	61
Table 6	Cognitive Commitments of Belief	106
Table 7	Cognitive Commitments of Acceptance	106
Table 8	Cognitive Commitments of Supposition	107
Table 9	Commitments of Hypothesizing	107
Table 10	Notes on Acceptance	246
Table 11	Notes on Supposition	247

ACKNOWLEDGMENTS

I want to express the deepest appreciation and gratitude to my committee chair, Professor Annalisa Coliva. Professor Coliva's expertise and insightful guidance was an essential part of my philosophical growth. Her ideas on beliefs as commitments greatly shaped my own views concerning the normativity of thought. Moreover, her dedication and enthusiasm for scholarly research was a source of inspiration. She pushed me to produce the best work possible. Completing this dissertation would not have been possible without her support.

I would also like to thank my committee member, Professor Margaret Gilbert. Professor Gilbert is a brilliant philosopher who introduced me to personal commitments of the will. Her important contributions to the philosophy of social phenomena and the normativity of commitments was instrumental in helping me craft my own theory of commitments governing our cognitive endorsements. Furthermore, my writing benefited greatly from her astute comments, suggestions, and guidance.

I must also recognize and thank my committee member, Professor Duncan Pritchard, for his incredible support. Professor Pritchard is an intelligent and thoughtful epistemologist who introduced me to issues concerning the nature and value of knowledge and understanding. Without his help and wisdom, I would not have been able to adequately develop my own account of understanding and rationality.

In addition to my committee, I want to thank Professor Sven Bernecker for the important role he played in shaping my epistemological views. Many parts of my dissertation benefited from his sage and perceptive advice and suggestions. He helped me become a better teacher and researcher. Moreover, he motivated me to explore new areas in epistemology.

Finally, I want to thank my family: Brooke Hankawa, Linda Sanders, Geoff Sanders, and Chelsey Sanders. I could not have finished this project without their love and support.

CURRICULUM VITAE

ADAM CHRISTOPHER SANDERS

Areas of Specialization
Epistemology and Philosophy of Mind

Areas of Competence
Ethics, Social Ontology, Metaphysics, and History of Analytic Philosophy

EDUCATION

Ph.D. in Philosophy (2019)
University of California, Irvine
Dissertation: *Commitments of Rationality*
Committee: Annalisa Coliva, Margaret Gilbert, Duncan Pritchard

M.A. in Philosophy (2017)
University of California, Irvine

M.A. in Philosophy (2013)
California State University, Los Angeles

B.A. in Philosophy (2012)
California State University, Los Angeles
Cum Laude, Phi Kappa Phi Honor Society, Golden Key International Honor Society

Certificate in Course Design (2019)
Division of Teaching Excellence and Innovation
University of California, Irvine

AWARDS AND GRANTS

Matchette Foundation for Graduate Student in Philosophy (2019 – 2020)
Dean of Graduate Studies Dissertation Fellowship Grant (2018 – 2019)
Kavka Endowed Prize in Philosophy (Fund# 7161) (2016 – 2017)
University of California, Irvine Fellowship Grant (2014 – 2015)

PUBLICATIONS AND PAPERS

Working Papers (unpublished)
"Doxastic Normative Tension"

TEACHING EXPERIENCE

Primary Instructor (Teaching Associate) – UC Irvine
Contemporary Moral Issues (Summer 2018)

Teaching Assistant – UC Irvine

Introduction to Philosophy (Fall 2015, Summer 2016, Winter 2017, and Fall 2017)
Theory of Knowledge (Winter 2016)
Puzzles and Paradoxes (Spring 2016, Spring 2017, and Spring 2018)
Contemporary Moral Problems (Fall 2016)
Introduction to Ethics (Summer 2017 and Winter 2018)

Guest Lecturer – UC Irvine

Introduction to Ethics (Summer 2017)
Introduction to Philosophy (Summer 2017)
Topics in Epistemology (Spring 2017)

Course Designer and Virtual Learning Development (Summer 2017)

Served as a Graduate Student Assistant Researcher that assisted in the creation and development of online course sites, virtual learning environments, and discussion forums for Introduction to Philosophy and Introduction to Ethics
Designed and created quizzes for Introduction to Ethics

Course Grader

Grader for Introduction to Ethics (Summer 2017)
Grader for Topics in Epistemology (Spring 2017)

GRADUATE RESEARCH EXPERIENCE

Graduate Student Researcher (2019)

Edited and prepared papers for publication in the book, *The Selected Writings of Eva Picardi* (by Annalisa Coliva)

Graduate Student Researcher (2017)

Editorial assistant for the book, *Epistemic Pluralism* (by Annalisa Coliva and Nikolaj Pederson)
Edited and formatted contributor papers
Compiled abstracts, book index, and key terms
Collected consent forms and author biographical information

Graduate Student Researcher (2017)

Coordinator for the *Hinge Epistemology Conference* at the University of California, Irvine

Graduate Student Researcher (Spring 2016)

Compiled the index for the book, *Varieties of Self-knowledge* (by Annalisa Coliva)

SERVICE TO THE DEPARTMENT

Graduate Student Representative for the philosophy department – UCI (2016 – 2017)

Served as a liaison between graduate students and the department of philosophy
Developed strategies and solutions that addressed the needs of graduate students

Graduate Student Workshop Coordinator – UCI (2016 – 2018)

Coordinated and managed several workshops for graduate students to present their work

OUTREACH

TH!NK Program Volunteer Instructor (2017)
Provided critical thinking and philosophical instruction to elementary school students

GRADUATE COURSES

Graduate Courses Taken at University of California, Irvine

Courses in Epistemology
Seminar: Medical Epistemology (Sven Bernecker and Annalisa Coliva)
Seminar: Bayesian and Mainstream Epistemology (Sven Bernecker)
Seminar: Epistemology of Conspiracy Theories, Testimony, and Experts (Sven Bernecker)
Seminar: Skepticism (Duncan Pritchard)
Seminar: Wittgenstein and Hinge Epistemology (Annalisa Coliva)
Seminar: Self-knowledge (Annalisa Coliva)
Directed Study: Understanding (Duncan Pritchard)
Directed Study: Research on Epistemic Normativity (Annalisa Coliva)
Directed Study: Testimony, Experts, and Trust (Annalisa Coliva)

Courses in Ethics, Social Ontology, and Political Philosophy
Seminar: Philosophy of Social Phenomena (Margaret Gilbert)
Seminar: Topics in Ethics: Rights (Margaret Gilbert)
Seminar: Evolution and Morality (Kyle Stanford)
Seminar: Rawls and Utility (Brian Skyrms and Aaron James)
Directed Study: Metaethics (Jeffrey Helmreich)

Courses in Metaphysics, Philosophy of Mind, and Free Will
Seminar: Philosophy of Mind (David Smith)
Seminar: Metaphysics (Yuval Avnur)
First Year Seminar: Free Will and Blame (Sean Greenberg)

Courses in Logic and Philosophy of Science
Seminar: Modal Logic (Sean Walsh)
Seminar: Set Theory (Sean Walsh)
Seminar: Philosophy of Biology (Caitlin O'Conner)
Directed Study: Truth (Sean Walsh)

Courses in the History of Philosophy
Seminar: History of Analytic Philosophy I: Frege, Russell, Wittgenstein (Jeremy Heis)
Seminar: History of Analytic Philosophy II: Carnap, Quine, Sellars (Jeremy Heis)
Directed Study: Husserl (David Smith)

Courses on Pedagogy and Professional Development
Seminar: Professional Development (Duncan Pritchard)
Certificate Program: Course Design (Division of Teaching Excellence and Innovation)

Graduate Seminars and Courses Taken at California State University, Los Angeles

Seminar: Logic (Mark Balaguer)
Seminar: Social and Political Philosophy: War and Genocide Studies (Mohammed Abed)
Seminar: Social and Political Philosophy: Karl Marx (Anna Carasthathis)
Seminar: Theory of Knowledge (Michael Shim)
Seminar: History of Modern Philosophy: Berkeley (Talia Bettcher)
Seminar: Metaphysics (Mark Balaguer and David Pitt)
Directed Study: Sartre (Jay Conway)
Course on 19th Century Philosophy (Jay Conway)
Course on Latin American Philosophy (Ricardo Gomez)
Course on Metaphysics (Mark Balaguer)
Graduate Philosophical Writing and Journal Editing & Production (Michael Shim)

REFERENCES

Annalisa Coliva
a.coliva@uci.edu

Margaret Gilbert
mpgilber@uci.edu

Duncan Pritchard
dhpritch@uci.edu

Sven Bernecker
s.bernecker@uci.edu

ABSTRACT OF THE DISSERTATION

Commitments of Rationality

By

Adam Christopher Sanders

Doctor of Philosophy in Philosophy

University of California, Irvine, 2019

Professor Annalisa Coliva, Chair

This dissertation focuses on several questions. How should we demarcate our concepts of cognitive endorsements (where "cognitive endorsements" include belief, acceptance, supposition, presupposition, and hypothesizing)? Are cognitive endorsements inherently normative states? Finally, how do cognitive endorsements factor into rationality and understanding?

In addressing these questions, I argue for several theses. I first show that belief, acceptance, supposition, presupposition, and hypothesizing are unique cognitive endorsements with distinct features concerning (i) their conceptual links to the truth, (ii) their rational basis, and (iii) their connection to an agent's will. Second, I argue that consciously formed cognitive endorsements engender distinct cognitive commitments that normatively constrain how one ought to treat the target propositions of these states. Third, I defend the view that diachronic belief-formation is governed by semantic commitments. *Semantic commitments* are normative constraints to form new beliefs depending on one's current beliefs, cognitive abilities, conceptual resources, and the logical relations that hold between propositions and the contents of one's current beliefs. Fourth, I argue that these commitments factor into rationality. More precisely, I defend a view called, "Robust Rationality." On this view, epistemic rationality consists of external, mind-independent rules and internal, mind-

dependent doxastic commitments that govern beliefs. Finally, I argue that rational cognitive endorsements can facilitate cognitive achievements of understanding. Understanding, I contend, comes in different species: veridical understanding, conjectural understanding, and narrative understanding. The type of understanding an agent possesses will depend on the underlying cognitive endorsement and the degree and quality of the relevant rational support.

This project has important consequences for epistemology and philosophy of mind. It establishes a taxonomy of cognitive endorsements and their conditions of rationality. This taxonomy can be employed for theoretical work mapping the plurality of ways rational agents can endorse and employ propositions in thought. Second, it is shown that both acceptance and supposition play an important role in manifesting rational agency and shaping our epistemic lives. Third, this project illuminates the normative structure of our thinking, reasoning, and processes of inquiry. For it shows that our cognitive endorsements are inherently normative states. Furthermore, this project advances our understanding of epistemic rationality and understanding.

INTRODUCTION: COMMITTED RATIONALITY AND UNDERSTANDING

Adam C. Sanders

Thinking is populated by a plurality of mental states. Among these states are propositional attitudes. These are cognitive states that take propositions as their objects. I might, for example, desire that *p*, hope that *p*, intend that *p*, love that *p*, or regard *p* as obtaining. To regard some proposition as being true is to believe that proposition. An agent's beliefs therefore settle matters concerning the truth or falsity of their target propositions. To believe is to commit oneself to how things stand in the world. In some cases, our beliefs amount to knowledge. In ideal situations, our beliefs facilitate epistemic accomplishments like understanding (see Pritchard 2014). It is therefore no surprise that beliefs occupy a central role in epistemology.

Recently, however, the focus in epistemology has broadened to include other belief-like states. Attitudes of acceptance and supposition are routinely employed in one's thinking. We can accept things we don't fully believe (see, for example, Cohen 1992). In some cases, we momentarily suppose propositions. Moreover, our attitudes of acceptance and supposition can be employed as mental premises in our theoretical and practical reasoning (see Bratman 1999). If this is the case, then our attitudes of acceptance and supposition can also play a crucial role in our epistemic lives.

Belief, acceptance, supposition, presupposition, and hypothesizing fall under a general conceptual category that I refer to as "cognitive endorsements". Cognitive endorsements are a class of positive mental attitude that one can take toward propositions. An attitude or state is a cognitive endorsement if it involves a kind of *mental fixing* or *positing* of a proposition, *p*, as being "given" in thought or reasoning. Mentally fixing *p* as being *given* can involve holding *p* to be true, the case, highly probable, or momentarily obtaining for one's current goals. The precise way in which *p* is fixed as being *given* can vary depending on the type of cognitive endorsement involved. Beliefs, for example, fix their target propositions as being true. Acceptance and supposition, however, can lack this

commitment to the truth of their target propositions. Accepting *p* might only involve a disposition or readiness to regard *p* as being plausible or practically useful. Moreover, a supposition that *p* might only amount to a tentative mental positing of *p*. Presuppositions are typically thought to be tacitly held cognitive endorsements that can affect conscious thinking and behavior. And hypothesizing is usually characterized as a type of inquiry in which one posits or mentally fixes a proposition, *p*, in thought as a possible explanation for some phenomena in order to determine the truth or plausibility of *p*. Mentally fixing a proposition as "given" distinguishes cognitive endorsements from other mental states. For instance, desiring, hoping, loving, and similar states are not appropriately classified as cognitive endorsements. For these attitudes do not necessarily involve positing their target propositions as being given or true in thinking and reasoning. One might, for example, desire *p* while thinking that *p* is neither true nor plausible.

The aim of this dissertation is to address several central questions concerning the nature and normativity of our cognitive endorsements. The first primary question I focus on is: How should we conceive of our individual cognitive endorsements? There are many accounts of belief, acceptance, supposition, presupposition, and hypothesizing. Epistemologists and philosophers of the mind do not always seem to agree on how to properly understand these attitudes. Indeed, it is not clear whether or not these states are distinct. Moreover, if these attitudes are truly distinct, then it is not straightforward how we are to understand their conditions of rationality. These questions deserve answers.

The second central question I address is: What kind of normativity is involved when we hypothesize and form conscious attitudes of belief, acceptance, and supposition? On one view, beliefs are simply dispositional states (see, for example, Quine 2008). That is to say, to believe *p* is to be in a state that can manifest certain causal and psychological dispositions to act or behave in specific ways. On this view, belief need not involve any normativity over and above the relevant dispositions. We might extend this line of thought and also characterize other cognitive endorsements as being purely

dispositional states that lack any normative element. Contrary to this, one might maintain that our cognitive endorsements generate a specific kind of normative element. But what would this normativity look like? What sort of normative constraints would these mental states impose on their possessors?

A related question I aim to address asks: If cognitive endorsements are inherently normative states, then how does this normative element factor into epistemic rationality? Many theorists conceive of epistemic or theoretical rationality as a system of rules (see Broome 2007). These are typically taken to be external requirements governing beliefs. Accordingly, agents ought to satisfy these rules in order to be rational. However, if our conscious attitudes of belief, acceptance, supposition, and hypothesizing are inherently normative states that constrain or govern the way agents ought to think, then how does the normative force of these states relate to our practices of rationality? If our cognitive endorsements generate their own unique normative elements, then our conception of rationality should countenance this fact. It is an important task to investigate the relationship between the normativity of our cognitive endorsements and the external rational requirements that govern them.

The final question that concerns me is: How do rational cognitive endorsements facilitate individual epistemic accomplishments of understanding? Many theorists contend that rational beliefs, in ideal situations, can contribute to understanding (see Pritchard 2014). However, attitudes of acceptance and supposition also influence our thinking and reasoning. Can these belief-like states also generate understanding? If so, how does this type of understanding differ from the understanding manifested by rational belief?

Each of the above questions deserves careful attention. Indeed, these are deep and important questions concerning foundational issues in epistemology, philosophy of mind, and theories of rationality and cognitive achievement. Hence, addressing these questions has theoretical importance. Illuminating the nature and normativity of these states will help to clarify our concepts and will lead

to a unified account of cognitive endorsements. This will further advance our epistemic theories and our understanding of epistemic rationality. For in order to achieve a satisfactory account of the nature and rationality of thinking, reasoning, and understanding, we need to first explicate the nature and normativity of our cognitive endorsements. This dissertation is a contribution to this important task.

In this dissertation I advance and defend the following main, overarching thesis (T):

> **(T)** Beliefs, acceptances, suppositions, presuppositions, and hypothesizing are distinct cognitive endorsements generating unique normative commitments that govern rational thinking and facilitate understanding.

Establishing the main thesis, (T), would provide a substantial contribution to epistemology. First, it facilitates a proper conception of our cognitive endorsements. Beliefs, acceptances, suppositions, presuppositions, and hypothesizing are distinct cognitive endorsements with unique conditions of rationality. Establishing (T) therefore has theoretical value, in that it explicates the way these cognitive endorsements interact in thought and contribute to rational thinking, reasoning, and understanding. Second, it helps explicate the normative structure of thought. Consciously formed cognitive endorsements are shown to engender unique normative *commitments*. I demonstrate that these commitments govern how an agent ought to treat the target propositions of her cognitive endorsements. A corollary of this is that the establishment of (T) advances our understanding of epistemic rationality. Indeed, it is further argued that epistemic rationality is best construed as a system of external, mind-independent requirements and internal, mind-dependent commitments. Finally, the truth of (T) promotes a better understanding of our cognitive achievements. For understanding, I contend, comes in species. And rationally held attitudes of belief, acceptance, and supposition can (in ideal situations) generate unique species of understanding.

This project leads to several surprising results. First, belief is not the only important epistemic state that can manifest cognitive achievements and rational agency. On the contrary, acceptances and suppositions can also contribute to rational agency and understanding. Second, it is a mistake to

conceive of cognitive endorsements as purely dispositional states generating causal or psychological tendencies to behave in certain ways. Rather, our consciously formed cognitive endorsements are inherently normative mental states. Consequently, thinking and understanding are inherently normative processes and states. Third, *commitments* are properly characterized as a broad class of normative constraints that factor into rationality. Moreover, satisfying our epistemic and cognitive commitments is necessary for rational thinking.

This dissertation is composed of several essays. Each essay establishes claims that can function as steps in an overarching argument for my main thesis, (T). The dissertation is structured in the following way.

In the first chapter I establish the claim that belief, acceptance, supposition, presupposition, and hypothesizing are distinct cognitive endorsements with unique conditions of rationality. It is shown that each of these states and processes has different qualities. Specifically, I show that these cognitive endorsements can differ with respects to how they are conceptually connected to the truth, the types of reasons that rationally motivate them, and how an agent brings them about in thought.

The second chapter argues for the claim that cognitive endorsements of belief, acceptance, supposition, and hypothesizing generate cognitive commitments that normatively constrain the way one mentally treats and implements the target propositions of these states. My research here is partly an extension of the important work done by Annalisa Coliva on beliefs as commitments (Coliva 2016, 2019). It is further shown that these commitments differ from the kinds of commitments generated by an agent's decisions and intentions (see Gilbert 2014). For example, the commitments stemming from intentions and decisions are normative constraints to perform some further action or bring about some state-of-affairs. However, cognitive commitments are normative constraints governing how one treats a proposition in thought. If cognitive endorsements generate commitments, then these mental

states are inherently normative. It is incorrect to characterize cognitive endorsements as purely non-normative dispositional states.

The third chapter focuses on diachronic commitments of belief-formation. Some theorists defend the idea that an agent's beliefs generate diachronic normative commitments to believe the propositions that logically follow from her current beliefs (see, for example, Levi 1997). I critique this view. Standard views of diachronic doxastic commitments face several challenges. I then argue for the view that an agent's diachronic belief-formation is constrained by what I call, "semantic commitments". This position avoids the challenges facing rival accounts and provides a more intuitive picture of the normativity governing diachronic belief-formation. According to the semantic commitment view, an agent, *S*, has a normative commitment to believe *p* if and only if (1) *p* follows from *S*'s current belief-set via some inference rule, and (2) *S* is in the cognitive position to competently form a cognitive endorsement with *p* as its content. One's cognitive position is constituted by one's beliefs, understanding, conceptual repertoire, and cognitive abilities. Hence, semantic commitments differ from cognitive commitments. Cognitive commitments are constitutive commitments of the individual attitudes themselves. Semantic commitments are *extraneous* commitments that depend on one's attitudes, cognitive abilities, understanding, the propositional contents of one's beliefs, and the logical relations that hold between those contents and further propositions. The picture that emerges is that doxastic reasoning and belief-formation is a normative process.

In chapter four I address some further issues concerning semantic commitments. Depending on one's current belief-set, one could have conflicting semantic commitments to believe contradictory things. In these situations, one is therefore normatively compelled to believe inconsistent things. I aim to address this issue. I contend that we lack sufficient reason to accept the claim that our semantic commitments have different degrees of known normative weight. If semantic commitments lack different degrees of normative weight, then these commitments differ from other normative

constraints, in that they do not weigh against each other to determine what an agent ought to do. Consequently, we must employ an additional normative procedure to resolve conflicts between semantic commitments. I argue that we should appeal to our cognitive commitments to resolve these conflicts. Successful satisfaction of one's cognitive commitments can normatively override the fulfillment of one's semantic commitments. Finally, I contend that this shows that our cognitive commitments are normatively prior to our semantic commitments. Cognitive commitments provide a fundamental normative guidance that our semantic commitments do not enjoy. Hence, our epistemic lives have a specific normative structure composed of a hierarchy of commitments.

In chapter five I focus on epistemic rationality. Some theorists think that rationality is understood as a system of rules governing the coherence and consistency of mental states. Contrary to this, I advance the "Robust View" of epistemic rationality. According to this account, epistemic rationality is construed as a system of external, mind-independent epistemic rules and internal, mind-dependent doxastic commitments. The Robust View, I contend, is a more tenable theory of epistemic rationality. For countenancing both epistemic rules and doxastic commitments provides a more complete picture of theoretical rationality. The Robust View provides better normative guidance and has the conceptual resources to handle a wide variety of cases. Moreover, the Robust View can resolve some problems facing current theories of rationality.

In chapter six I discuss epistemic accomplishments and their relationship to cognitive endorsements. Many theorists hold that knowledge and understanding require belief. Here I focus on understanding. Understanding is typically thought to be a kind of cognitive accomplishment that our rational beliefs can engender (see Pritchard 2014). I advance the "Multiple Species" view of understanding. Understanding, I contend, is generated by a cognitive endorsement and an ability I refer to as "epistemic discernment". *Epistemic discernment* is (roughly) a cognitive capacity to discern how information and propositions fit together and relate to each other in some epistemically

significant way. On this view, rational belief can generate a strong kind of "veridical understanding". Acceptance can generate what I refer to as "conjectural understanding". This species of understanding is a moderate cognitive achievement. Finally, it is shown that suppositions can generate a weak kind of "narrative understanding". It is important to note that each species of understanding can be achieved when the associated cognitive endorsements are rationally held. To be rationally held requires satisfying the corresponding cognitive commitments. Hence, satisfying cognitive commitments associated with belief, acceptance, and supposition is a fundamental aspect of understanding and cognitive achievements.

I further contend that this alternative view of understanding is a more tenable theory and preferable account. First, it is appropriately robust and provides the conceptual machinery to better handle a wide variety of cases involving understanding. Second, this theory overcomes several challenges facing rival accounts. Finally, by maintaining that understanding comes in different species, it can be shown that much of the disagreement over the nature of understanding is merely semantic.

Each chapter in this project establishes a number of propositions that can function as premises in support of the main thesis, (T). We have good reason to understand belief, acceptance, supposition, presupposition, and hypothesizing as being distinct kinds of cognitive endorsements. Consciously formed cognitive endorsements engender normative commitments that constrain how their target propositions are treated and implemented in thought and reasoning. These are normative commitments that differ from the semantic commitments stemming from individual beliefs and subjective cognitive positions. Furthermore, our cognitive commitments play a crucial role in our theories of rationality and accounts of epistemic accomplishments. Satisfying cognitive commitments is a constitutive part of epistemic rationality. And successful fulfillment of one's cognitive commitments is a necessary condition for understanding.

Finally, I conclude the dissertation by providing some general remarks concerning future topics of research. This dissertation advances a unique and tenable picture of the normativity and nature of our cognitive endorsements. It further advances a plausible account of robust epistemic rationality. Hence, this project lays a foundation for important future work in epistemology and rationality. First, much of my current work focuses on doxastic rationality – i.e., the rationality concerning belief. However, acceptance and supposition are also shown to be inherently normative states. It is an important task to expand on the rationality of these cognitive endorsements. I am specifically interested in the rationality of acceptance. If acceptance is a unique cognitive endorsement, then the rational requirements of this state will differ from the requirements governing belief. I briefly discuss how a complete account of the rationality of acceptance might differ from the rationality of our beliefs. Second, my research provides the conceptual framework needed for further work on understanding. I am particularly interested in collective understanding. Hence, I will provide a rough sketch of how an account of collective understanding might be formulated, and how this type of epistemic achievement might differ from the understanding that individual agents possess.

I. A TAXONOMY OF COGNITIVE ENDORSEMENTS

Adam C. Sanders

0. INTRODUCTION

Agents can take certain cognitive stances toward propositions. We can consider propositions to be true, false, likely, implausible, and so on. When one considers a proposition, *p*, to be true, one typically *believes p*. Our beliefs indicate how we take the world to be. In optimal cases, beliefs can amount to knowledge.[1] Sometimes our beliefs generate understanding.[2] Moreover, beliefs can be employed in reasoning to form new beliefs and extend knowledge and understanding. Beliefs are therefore an important object of study for epistemologists. However, we can hold other, belief-like attitudes toward propositions that play a crucial role in our thinking and reasoning. We often endorse propositions that we may not fully believe or take to be true. Sometimes these belief-like endorsements factor into our reasoning and inquiry. In some cases, our thinking is guided by propositions that we tacitly hold.

Consider, for example, a case in which I am planning a formal party. During my planning process I take my current evidence to be insufficient to establish the proposition, *that there will be more than 100 people in attendance*. I do not consider this proposition to be true. Nor do I consider this proposition to be false. However, I decide to endorse this proposition in my thinking in order to make certain preparations. Consequently, I may assert this proposition to vendors and caterers. Moreover, I affirm this proposition in my thinking to settle further matters concerning whether I should have available seating and food to accommodate more than 100 people. However, while maintaining this positive stance toward the proposition, *that there will be more than 100 people in attendance*, I also momentarily consider the negation of this proposition. This is to say, I also temporarily endorse the

[1] One can distinguish between true beliefs and knowledge. I do not provide an analysis of knowledge here. I assume a fairly standard or mainstream account of knowledge presented by Wayne Davis: an agent *knows* some proposition, *p*, when (1) *p* is true (2) the agent believes *p*, and (3) the agent's belief is sufficiently justified to preclude instances of luck (Davis 1988, p. 169).
[2] I have my own account of understanding that I defend in chapter 6.

proposition, *that it is not the case that there will be more than 100 people in attendance*. I do so to determine what consequences might follow from this proposition if it were true.

In the above case, I mentally endorse several propositions. Several questions present themselves. First, what kind of propositional attitudes are picked out by my mental endorsements? Second, am I rational to hold these attitudes simultaneously during my planning and reasoning? Answering the second question depends on how we answer the first question. If, for example, we take my attitudes to be instances of belief, then it follows that I come to believe inconsistent things. Moreover, I seem to form these beliefs on insufficient evidence. One might claim that I am irrational or that I am not forming my beliefs in an appropriate way. On the face of it, however, my planning and thinking in the above case seems perfectly rational. There are many situations in which one can rationally endorse and put forward propositions in thought that one does not believe. How can this be rational? Many theorists hold that the attitudes in these situations fall under different kinds or categories of propositional attitudes. Indeed, it might be insisted that during my party planning I have a mixture of different propositional attitudes. Let "*p*" refer to the proposition, *that there will be more than 100 people in attendance*. It seems that I do not believe *p*. Nor do I believe not-*p*. However, I still take on an attitude of *acceptance* toward *p* to make further party arrangements.[3] This attitude may not be a type of belief. Moreover, at the time when I consider what might follow from the negation of *p*, I could be appropriately characterized as *supposing* or *assuming* not-*p*. Hence, my thinking includes a mixture of different propositional attitudes. If these attitudes are different, then they might have different conditions of rationality. Consequently, it is possible to rationally hold these attitudes together.

States of belief, acceptance, supposition, presupposition, and hypothesizing fall under a generic conceptual category that I refer to as *cognitive endorsements*.[4] When an agent hypothesizes or

[3] Cohen (1992) and Bratman (1999) have the view that one can accept *p* for pragmatic purposes.
[4] To be clear, this list of states may not exhaust the different kinds of cognitive endorsements.

forms a belief, acceptance, or supposition, then she can be understood as taking a mental stance of endorsement or affirmation toward a proposition. This *endorsement*, I claim, is a broad representational stance: it includes mental stances of "fixing" a target proposition as either being true, likely, plausible, credible, or tentatively "given" in one's thought. How the target proposition is fixed in thought depends on the kind of cognitive endorsement involved. It follows that cognitive endorsements fundamentally differ from other intentional attitudes like desires and hoping. For if I desire *p* or hope that *p* obtains, then I need not fix *p* in my thought as being either true, plausible, or likely.

Cognitive endorsements play a fundamental role in our thinking and reasoning. They are attitudes that can factor into an agent's theoretical and practical reasoning. Hence, cognitive endorsements can be employed to facilitate action and extend our beliefs, knowledge, and understanding. A satisfactory account of the nature and normativity of reasoning, knowledge, and understanding requires a satisfactory account of these cognitive endorsements. Many philosophers have provided different accounts of these intentional attitudes. The problem, however, is that it is not entirely clear how these attitudes relate to each other. Nor is it clear how these attitudes differ in the way they respond to rational considerations.

The central aim of this chapter is to address these worries. I advance a plausible and unified taxonomy of cognitive endorsements. In doing so, I explicate the fundamental features of belief, acceptance, supposition, presupposition, and hypothesizing. I argue that these attitudes differ in three fundamental ways. First, these states differ in how they rationally dispose an agent to regard the truth or plausibility of their target proposition. For example, beliefs are conceptually tied to the truth whereas suppositions and acceptance lack this quality. Second, cognitive endorsements differ in the types of reasons that rationally motivate and shape them. Rationally held beliefs enjoy a high degree of epistemic justification. Rational acceptance is similarly responsive to rational considerations. But one can rationally accept a proposition for pragmatic or practical considerations that one has in a given

context. Suppositions, however, need not be based on any rational support. I argue that suppositions come in two different species: rational suppositions and arational suppositions. *Rational* suppositions are based on reasons (epistemic or pragmatic), whereas *arational* suppositions are not based on reasons. Finally, these attitudes differ in how they are brought about in response to reasons. One's rational belief that *p* is not brought about at will and without regard to one's epistemic reasons in support of the truth or likelihood of *p*. However, other kinds of cognitive endorsements can be brought about at will without having to be sufficiently based on the available evidence for their target propositions.

The project of demarcating cognitive endorsements shares some resemblances with John Searle's theory of "speech acts" (see Searle 1969). According to Searle, a sentence or proposition can be employed in a language by means of a speech act (Searle 1969, p. 16). *Speech acts* are understood to be a linguistic class including demanding, asserting, questioning, and so on (Searle 1969, pp. 16, 22 – 23). However, Searle insists that different speech acts can share the same content: an agent, *S*, can demand that *p*, assert that *p*, question that *p*, and so on (Searle 1969, p. 16). Accordingly, the same content, *p*, can be linguistically presented in different ways depending on the specific kind of speech act involved. (Searle 1969, pp. 16, 22 – 23, 125). In a similar vein, cognitive endorsements are taken to be a broad class of epistemic states an agent can possess. These states involve a mental endorsement and presentation of a proposition in thought. Different cognitive endorsements can share the same content. For instance, one can believe *p*, accept *p*, suppose *p*, and so on. However, the way *p* is mentally treated and put forward in thought will differ depending on the type of cognitive endorsement involved.[5]

[5] I should make two clarificatory points here. First, this is a rough and incomplete sketch of Searle's theory. However, given that my focus is on epistemic states and cognitive endorsements, I will not discuss the finer points of speech acts. Second, Searle also thinks that speech acts are governed by rules (Searle 1969, pp. 22, 33 – 42). This is another point of similarity between Searle's theory and my account of cognitive endorsements. In the next chapter I demonstrate that cognitive endorsements are governed by normative commitments. These commitments are constitutive aspects of cognitive endorsements.

Several caveats are in order. First, I do not aim to provide necessary and jointly-sufficient conditions for our concepts of belief, acceptance, supposition, presupposition, and hypothesizing. Rather, my goal is to identify the most salient features of these attitudes when they are explicitly formed. Second, there are many different conceptions of cognitive endorsements. I will briefly discuss some of these accounts. However, I will primarily focus on *explicitly* formed cognitive attitudes. Agents typically form these attitudes consciously and based on their available evidence or reasons. Third, I will focus on propositional attitudes that mature, cognitively developed agents typically possess.

What emerges from this discussion is a plausible taxonomy of cognitive endorsements. This is important for several reasons. First, many epistemologists seem to understand these cognitive endorsements in different ways. The taxonomy defended here provides one plausible way to incorporate some of these different conceptions into a unified theory. It can be shown that many theorists are merely talking past each other when articulating different accounts of belief, acceptance, and supposition. Second, this taxonomy elucidates the way cognitive endorsements can be responsive and sensitive to different kinds of reasons and epistemic justification. Hence, it explains how one can rationally hold different kinds of cognitive endorsements toward the same proposition at the same time. Finally, the account explicates the ways agents manifest rational agency by taking different cognitive stances towards the truth, plausibility, and usefulness of propositions.

The chapter is structured in the following way. In section 1 I articulate the core features of explicitly held belief. Section 2 advances an account of acceptance that is partly based on Cohen's (1992) view claiming that acceptance involves a mental activity of employing a proposition as a premise in reasoning. On my view, accepting *p* involves a positive mental endorsement and a resolution or disposition to use *p* in thought. In section 3 I turn to suppositions. Suppositions involve weak and tentative endorsements of propositions. I further argue that suppositions come in different species: rational suppositions and arational suppositions. In section 4 I discuss presuppositions. I follow

Stalnaker (1987) in taking presuppositions to be tacitly held (unconsciously held) cognitive endorsements that guide thought and action. Contrary to Stalnaker, I argue that presuppositions are not restricted to tacit beliefs. Tacit acceptances can also function as presuppositions. Section 5 advances an account of hypothesizing. I argue that hypothesizing is a mental process involving an attitude of acceptance in order to bring about certain epistemic ends an agent has in a context. I conclude the chapter in section 6 by situating cognitive endorsements on an axis of rational support.

1. BELIEF

Belief is an important concept for epistemology. First, beliefs play an important role in theoretical and practical reasoning. Moreover, many theorists maintain that knowledge and understanding requires belief. In this section I demarcate some of the core features of one common conception of explicitly formed propositional belief.[6] The main aspects of belief that I focus on here are (1) the connection belief has with the truth, (2) the kind of reasons that rationally motivate belief, and (3) the way in which one comes to rationally form a belief.

Several preliminary remarks are in order. First, beliefs are sometimes characterized as being *full* beliefs or being *degrees* of belief. According to Gilbert Harman, a *full* belief is properly understood as an "all-or-nothing" attitude (Harman 1986, pp. 12-13). Either one believes *p,* or one does not believe *p*.[7] Degrees of belief are different. According to David Christensen, we sometimes feel more confident in the truth of some propositions than others (Christensen 2004, p. 13). This might suggest that beliefs come in *degrees* or *credences* (i.e., subjective probabilities one has toward propositions). According to

[6] There can be different mental states that are picked out by our general concept of *belief*. I will not discuss every conception of belief that has been put forward in epistemology or the philosophy of mind. I will focus on a few central features of explicit belief. Many of the paradigmatic features of this attitude that I discuss here have been endorsed by other theorists.

[7] As David Christensen puts the point, a full belief is understood to be "an attitude that one can either take, or fail to take, with respect to a given proposition" (Christensen 2004, p. 12).

Christensen, an agent's credence toward *p* is defined as a "level of confidence" in the truth of *p* that is assigned a value falling between the numerical values of 0 and 1 (Christensen 2004, pp. 13-16).[8]

While it can be useful to sometimes describe beliefs as credences, for my purposes here I will focus on full beliefs.[9] I take full beliefs to be especially important given their role in propositional reasoning, rationality, and understanding. Indeed, we typically employ full beliefs as premises in conscious processes of reasoning and inquiry. From this point forward, then, when I speak of "belief" I will be referring to an all-or-nothing state.

Second, the kind of beliefs under discussion here are best thought of as mental representations. This naturally leads into certain questions concerning the content of what is represented when one possesses a belief. Fred Dretske maintains (roughly) that for an agent, *S*, to believe *p* is for *S* to represent a state-of-affairs as obtaining, insofar as *S* possesses the information or concepts required to specify the content, *p* (Dretske 2000, pp. 67-71). In what follows, I focus on beliefs as mental representations that take propositions as their objects.[10] To believe *p* is to possess a state that

---

[8] For example, I am highly confident that I have a body. I have good evidence for this proposition, namely, my current perceptual evidence. However, I am less confident that my body locked the front door before I left the house this morning. For I cannot remember the exact time I locked my door. In fact, I admit that it could be possible that I failed to lock the door. I do not share the same degree of confidence concerning the possibility that I do not have a body. Hence, I might have a 0.9 credence toward the proposition *that I have a body*. But I might have a 0.5 credence toward the proposition *that I locked the front door this morning*.

[9] One could adopt the view that full beliefs and degrees of beliefs are related. For example, one could endorse a "threshold" view of credences (Foley 1993, pp. 141-142). According to Foley, an agent, *S*, would fully believe *p* if *S* has some credence in the truth of *p* that surpasses some stipulated threshold between 0 and 1 (Foley 1993, pp. 141-142). On this view, we might take any credence over 0.5 to count as a full belief. Hence, if an agent, *S*, has a credence toward *p* that was greater than 0.5, then *S* would be attributed with a full belief that *p*. Any credence that fell below 0.5 would indicate a lack of a belief that *p*. The so-called "threshold" view is just one way to think of the relation between credences and full beliefs. It is a separate question as to whether the threshold view is plausible. Pursuing these questions will exceed the scope of my project here. I therefore set these concerns aside.

[10] I will not take any position on the specific nature of propositions. I will, however, follow standard accounts of propositions that maintain that a proposition normally has the following characteristics: (1) a proposition is composed of, or constituted by, concepts; (2) A proposition can be the meaning expressible by a declarative sentence; (3) A proposition has truth-conditions – it indicates how things

represents a content specified by the proposition, *p*. This conception of belief is especially important for epistemology. Propositional beliefs can amount to propositional knowledge and understanding. If I am to be credited with knowing *p*, then I must also believe *p*. Moreover, beliefs with propositional contents can factor into propositional reasoning. For example, if I believe *p*, and if I believe that *q* is implied by *p*, then I can reason with these beliefs to form a new belief with the proposition, *q*, as its content. Hence, I will restrict my focus to propositional beliefs.[11]

Third, I primarily focus on consciously formed or *occurrent* beliefs. Occurrent beliefs can be distinguished from *dispositional beliefs*. Gilbert Harman defines an *occurrent* belief as a belief that "is either currently before one's consciousness or in some other way currently operative in guiding what one is thinking or doing" (Harman 1986, p. 14). A *dispositional* belief, however, is not an occurrent mental representation. Harman claims that a dispositional belief need not be consciously put forward in one's thought or reasoning (Harman 1986, p. 14).[12] Rather, it is an unconscious belief one possesses that can be "potentially occurrent" if one recalls it and consciously puts it forward in thought (Harman 1986, p. 14). I will take it for granted that beliefs can persist and be held in different ways. An agent can consciously hold a belief. In this case, her belief is an occurrent one. However, this belief can persist as a dispositional belief (i.e., an unconscious, background belief) even when the agent is no longer entertaining the belief in her conscious thought.

---

must be (or what sort of facts or states-of-affairs must obtain) in order for it to be true (i.e., "*p*" is true if and only if *p*); and (4) A proposition is typically true or false (but not both) at a particular time.

[11] To be clear, this is not the only way to construe the notion of *belief*. According to Leslie Stevenson, we can distinguish between different classes of belief that differ in their contents (Stevenson 2002). Stevenson identifies a "lower" kind of "non-linguistic" belief that can be attributed to creatures who lack the concepts required to specify the content of their beliefs and perceptual states (Stevenson 2002, p. 119). These beliefs represent non-conceptual or non-propositional contents. Considering my epistemological interests concerning knowledge and understanding, I will only focus on propositional beliefs. These beliefs are related to what Stevenson refers to as "linguistic beliefs" that are "reasoned"-based: one possesses a linguistic and reasoned-based belief that *p* when one has the concepts specifying *p*, and one judges that *p* obtains by considering the evidence in favor of *p* (Stevenson 2002, p. 120).

[12] Annalisa Coliva makes a similar point. She notes that a dispositional belief need not always be the product of consciously judging that some proposition is true (Coliva 2019, p. 4).

Fourth, we can distinguish between *explicit* beliefs and *implicit* beliefs. Harman defines an *explicit* belief that *p* as a "mental representation" that some agent actually holds (Harman 1986, p. 13). Occurrent and dispositional beliefs are types of explicit beliefs. But an *implicit* belief that *p* is, according to Harman, not a mental state that an agent currently possesses (perhaps by explicitly judging that some *p* is true) (Harman 1986, p. 13). Rather, an implicit belief, he claims, is a belief that is either "easily inferable or implied by explicit beliefs" that an agent currently has (Harman 1986, p. 13).[13]

Consider a case where I come to explicitly believe the proposition, *that a Hydrogen atom has one electron*. Call this proposition, "*p*". Suppose that I also believe *that werewolves do not exist*. My explicit belief that *p* logically implies many other propositions. According to introduction rules in classical logic, one such implication is the disjunction, *p* or *q*, where *q* can be any proposition. If I have the relevant concepts, including the concept of *disjunction*, then I can be attributed with the *implicit* belief *that a Hydrogen atom has one electron or werewolves exist*, even though I have never explicitly held this belief.

I do not doubt that there can be a useful distinction between implicit and explicit beliefs. Moreover, I would go further and suggest that this distinction applies to other attitudes like acceptance. However, explicit beliefs play a crucial role in our conscious reasoning. Moreover, knowledge seems to involve explicitly held beliefs. For it would not seem right to attribute knowledge that *p* to an individual that only implicitly believes *p*. This individual does not currently believe *p*, even if she is in a position to form this belief if certain conditions were satisfied (e.g., if she justifiably comes to judge and believe *p* to be true). Hence, I will set aside any further discussion of implicit beliefs.

[13] Implicit beliefs, as I describe them here, closely resemble Robert Audi's notion of *dispositions to believe* (Audi 1994). According to Audi, "a disposition to believe is a readiness to form a belief" (Audi 1994, p. 424). At time, *t*, *S* might not explicitly believe *p*. But *S* could form a belief that *p* at *t*. Depending on *S*'s situation, conceptual repertoire, and stock of explicitly held attitudes (including beliefs), *S* might have a disposition to believe *p*. I take implicit beliefs to be synonymous with dispositions to believe.

I now turn to the salient features of belief.[14] Frist, an explicitly held belief does not simply represent its content in any non-specific way. Rather, to believe *p*, and thus to mentally represent a state-of-affairs specified by the propositional content, *p*, is to be in a state that portrays the world to be a certain way. One's belief typically represents a certain fact or state-of-affairs as *obtaining*. This view is shared by many philosophers. Bernard Williams, for example, claims that, "beliefs aim at truth" (Williams 1973, p. 136). Duncan Pritchard similarly asserts that there is a tight conceptual "connection between belief and truth" (Pritchard 2016, p. 90). To believe *p*, he insists, is to possess an attitude that is "incompatible with an attitude of agnosticism about [*p*'s] truth" (Pritchard 2016, p. 101). Moreover, both Timothy Williamson and John Searle characterize belief as having a "mind-to-world direction of fit" (Williamson 2000, p. 2; Searle 2001, p. 37).[15] Accordingly, to believe *p* is to mentally represent the state-of-affairs specified by *p* as obtaining. This seems right. Beliefs are concerned with what is true or what is the case.[16] We don't believe things that we think are false or highly unlikely. If an agent, *S*, believes *p*, and therefore takes *p* to be the case, then it would be incoherent for *S* to simultaneously have an attitude in which she was also committed to being unsure of *p*'s truth.

Even if we take beliefs to aim at truth, this does not imply that their target propositions are always true. The so-called "fit" of an agent's belief need not always be correct. As Williamson claims, a "belief is true if it fits the world, false otherwise" (Williamson 2000, p. 2).[17] For example, I might form the belief that there are four hyenas in front of me. My belief represents a certain state-of-affairs as being the case. Hence, my belief is correct if and only if it is indeed the case that the world is the

[14] As noted earlier, I focus on the kind of propositional attitudes and beliefs that mature agents can possess.
[15] See also Wayne Davis (1988). Davis insists that, "[t]o believe a proposition is to believe that it is true" (Davis 1988, p. 170).
[16] I say more about the distinction between belief's aim at *truth* and *what is the case* below (see p. 20).
[17] This point seems obvious. However, I explicitly note it to sufficiently characterize the representational nature of beliefs.

way my belief represents it to be. More precisely, my belief in this case is correct if an only if there are four hyenas in front of me. Thus, a belief that *p* at some time, *t*, is correct if and only if *p* is true at *t*.

I take it that a central feature of our notion of *belief* is that beliefs are concerned with the truth. Typically, to believe *p* involves being disposed to regard *p* as being true. However, this conviction of *p*'s *truth*, I submit, can be understood broadly. According to William Alston, small children lacking the concept of *truth* can still be credited with beliefs (Alston 2007, p. 132). These individuals, he claims, might believe *p* by merely considering or "feeling" *p* to "be the case" (Alston 2007, p. 132). Perhaps this is so.[18] What is important for my purposes here is that to believe *p* is to regard *p* as *obtaining* in some broad sense of the term. If an agent believes *p*, then she takes a non-agnostic stance toward *p*, such that she either considers *p* to be true, or takes *p* to be the case, or regards *p* to be highly probable, or thinks that *p* is more likely than not-*p*.[19] Belief therefore has the following feature:

> **(B1)** For an agent, *S*, to believe *p* at some time, *t*, is for *S* to be disposed to regard *p* as being true, or being the case, or being more likely than not-*p* at *t*, in such a way that *S* cannot be agnostic toward the truth of *p* at *t*.

Now, L. Jonathan Cohen characterizes belief as involving certain feelings (Cohen 1992, pp. 4-5). He claims that a "belief that *p* is a disposition…normally to feel it true that *p* and false that *not-p*, whether or not one is willing to act, speak, or reason accordingly" (Cohen 1992, p. 4). It might be tempting, then, to modify feature (B1) in the following way:

> **(B1*)** For an agent, *S*, to believe *p* at some time, *t*, is for *S* to have a disposition to feel that *p* is true at *t*, such that *S* cannot simultaneously feel that not-*p* is true, or that not-*p* is just as likely as *p*.

I admit that it is possible for one to have a disposition to manifest a feeling that *p* is true when one believes *p*. However, this does not seem to be a necessary feature of belief. One might follow

[18] It could be that merely having the conceptual abilities to think or feel that something is "the case" is sufficient for having the capacity to think that something is true.

[19] This also allows for my conception of belief to be compatible with some accounts of degrees of belief. One might be credited with a belief that *p* if one's credence or degree of belief is such that one is highly confident that *p*, or if one is more confident that *p* is the case than not-*p*.

Alston here and hold that small children might lack the concept of *truth* and nevertheless be credited with believing *p* when they are only disposed to *feel* that *p* is *the case* (Alston 2007, p. 132). I would argue, however, that beliefs might lack a necessary phenomenal character of *feeling* that something is the case. It is conceivable that one can believe *p* without being disposed to *feel* that *p* is true or feel that *p* is likely. Moreover, it seems that the mode of *regarding* or *thinking* that a proposition, *p*, is true or likely is more fundamental than having a disposition to feel that *p* is true or likely. Intuitively, when one believes *p*, then one would regard *p* as being the case or more likely than not-*p*. And it seems (at least to me) that to regard *p* as either being the case or being likely is an antecedent condition for one to feel that *p* is true (or that *p* is the case). My intuition is that (B1), rather than (B1*), is a more plausible formulation that identifies a more basic quality of belief.

It is important to note that (B1) distinguishes beliefs from other propositional attitudes. One can desire *p*, or suppose *p*, or hate *p*. But possessing these attitudes toward *p* need not require one to think (or be disposed to think) that the target proposition, *p*, is true or likely. Beliefs are a specific kind of cognitive endorsement: to believe *p* requires one to regard or affirm *p* as being the case or true.[20]

At this point, we can turn our attention to the ways in which belief acquisition and revision occurs. These considerations illuminate several further features of belief. First, when an individual comes to believe *p*, she does not typically do so for no reason whatsoever. It is commonly thought that rational belief is based on evidence. For example, Hume claimed that an agent ought to "proportion his [or her] belief to the evidence" (Hume 1975, p. 110). Pritchard also argues that a belief that *p* is appropriately responsive to "rational considerations" and evidence for the truth of *p* (Pritchard 2016, p. 90). In a similar vein, Richard Foley notes that it is only evidence and epistemic reasons for *p*

[20] Notice that this conception of belief differs from other conceptions of belief. As noted earlier, some theorists claim that there are non-propositional beliefs (see Stevenson 2002, p. 119). If there are non-propositional beliefs, then these beliefs cannot involve a disposition to regard a proposition as being true. Moreover, an implicit belief is not an explicitly held belief that an agent possesses. Hence, one does not regard a proposition, *p*, as being the case if one only implicitly believes *p*.

that count as the appropriate *kind* of reasons for one to rationally regard *p* as being true (Foley 1993, p. 16). Nonevidential considerations (e.g., purely practical considerations), he claims, are "reasons [that] do not aim at truth" (Foley 1993, p. 16). Hence, in normal situations, one rationally believes *p* based on evidence or epistemic reasons in support of the truth of *p*.[21]

Similar considerations apply to typical cases of belief-revision. As Michael Bratman notes, one can revise one's belief when there is a "change of mind" (Bratman 1999, p. 19). Here I understand an agent's *change of mind* as consisting in either (1) a new belief that *p* (or not-*p*) is true or is the case, or (2) suspending belief regarding the truth or falsity of *p*. An explicitly held belief is normally altered or revised in the light of new information that counts as evidence affecting the truth or falsity of the propositional content of that belief.

Consider, for example, a case in which I believe that my house is haunted by ghosts. Suppose that my belief is based on evidential support. Perhaps I perceived spooky, ghost-like sounds and apparitions. In this case, I take such sounds and apparitions to count as evidence for the truth of my belief. My belief is therefore rationally motivated by my available evidence. But beliefs are revisable. I might suspend my belief about my haunted house if, at some later time, I were to find out that, unbeknownst to me, I had been given a hallucinogenic drug, or that my neighbor was playing a trick on me by entering my house each night wearing a white sheet. My belief that my house is haunted can be subject to rational revision if it is appropriately based on my available evidence or epistemic reasons.

An interrelated point is that belief is a relatively stable and persistent doxastic state. According to Bratman, rational beliefs are not dependent on pragmatic considerations that one has in a specific situation (Bratman 1999, p. 18). He claims that one's belief is "context-independent" – that is, either one believes a proposition, or one fails to believe that proposition, regardless of the practical

[21] I understand "practical reasons" in favor of a proposition, *p*, to include prudential and pragmatic considerations that would make one "better off", in the broadest sense of the term, if one cognitively endorsed *p*. "Evidential reasons" are taken to be epistemic considerations that support the truth of *p*.

considerations one has within a context (Bratman 1999, p. 18). I do not believe that my house has a roof depending on the time of day, or because holding this belief has certain practical benefits. Rational beliefs are based on sufficient *epistemic* considerations that warrant the truth of their target propositions. Hence, I rationally believe that my house has a roof given my evidence in support of the truth of this proposition. Moreover, I can rationally continue to hold this belief consciously (an occurrent belief) or unconsciously (a dispositional belief) in different contexts given my available evidence. I do not normally abandon this belief even if I have good pragmatic reasons to do so.

Belief therefore has the following feature:

(B2) An agent's rational belief that *p* is responsive and sensitive to evidence and epistemic considerations for or against the truth of *p*, regardless of nonevidential or pragmatic reasons that agent has in a specific context.

To be clear, (B2) does not imply that agents only form beliefs on evidential grounds. In some cases, an agent might form an irrational belief. Irrational beliefs are not appropriately sensitive to epistemic considerations. A belief is rational if and only if it is based on good evidence or epistemic reasons that supports the truth of its target proposition (see Pritchard 2016, p. 90).[22] Moreover, (B2) does not suggest that an agent will indefinitely hold on to her beliefs in all contexts. An occurrently formed belief at time, t_1, can persist occurrently or dispositionally at some future time, t_2. But one can rationally revise, alter, or abandon one's belief that *p* in light of new evidence and epistemic considerations affecting the truth or falsity of *p*. Moreover, it is possible to simply lose a belief by merely forgetting what one believes. The crucial point is that rational beliefs are held and revised on the basis of good epistemic reasons.

Feature (B2) leads to certain issues concerning the connection between an agent's will and her ability to form or revise beliefs. Many theorists hold that beliefs are involuntarily acquired. Hume

[22] Here I am concerned with the rationality of an individual state of belief. I will discuss the rationality of systems of beliefs later in the dissertation.

understood a belief to involve "a certain feeling or sentiment" that something was the case, such that this conviction and feeling was not dependent "on the will" (Hume 1978, p. 624). Cohen also asserts that one cannot possess a belief by merely "deciding to believe" (Cohen 1992, p. 21). And Coliva argues that the process of believing *p* based on sufficient evidence in support of *p*'s truth is not a voluntary affair (Coliva 2016, pp. 33-34).[23] The point is that voluntarily willing oneself to believe *p* without considering the relevant evidence is not an appropriate route to belief-formation. Typically, if an agent, *S*, considers her evidence for or against some proposition, *p*, and takes this evidence to be sufficient to support the truth of *p*, then *S* will *ipso facto* believe *p*. In normal circumstances, *S* cannot voluntarily will herself to believe *p* for purely pragmatic reasons. Rather, rational beliefs are the products of taking into consideration the available evidence for the truth of their target propositions. One cannot rationally believe *p* by voluntarily *ignoring* the epistemic reasons for the target proposition.

Of course, one can will oneself to seek out evidence for or against propositions. According to Foley, an agent might have nonevidential and pragmatic reasons to put herself "into an evidential situation in which belief will be possible" (Foley 1993, p. 16). Nevertheless, rational belief-formation is not dependent on one's willing to believe for pragmatic considerations.

Even if a rational belief that *p* involves taking into consideration the epistemic reasons for *p*, it is not clear how this belief is brought about. One plausible view concerns the relationship between belief-formation and mental acts of *judgment*. In his paper, "Negation," Frege characterized a *judgment* as an "act" of "acknowledging the truth" of a so-called "thought" or proposition (Frege 1997, pp. 354, fn. D).[24] Several theorists claim that beliefs are intimately connected to these individual acts of

[23] See also Quine (1987) and Swinburne (2001). According to Quine, a belief is involuntarily formed given one's evidence (Quine 1987, p. 19). Swinburne insists that one cannot believe at will (Swinburne 2001, p. 39).

[24] In his "Über Sinn und Bedeutung," Frege notes that a judgment "is not the mere grasping of a thought, but the admission [*Anerkennung*] of its truth" (Frege 1997, p. 158, fn. G). I take this admission or acknowledgment to be a sort of conviction, such that one cannot help but regard the target proposition as being true or the case (according to feature (B1) above).

judgment (Peacocke 1998; Coliva 2016). For example, Christopher Peacocke notes that a typical way in which an agent comes to believe *p* consists in that agent first judging *p* to be the case (Peacocke 1998, p. 88). But a judgment that *p* is not synonymous with believing *p*. According to Coliva, a belief that *p* is a mental state, whereas an individual judgment that *p* is a "mental action" that an agent performs when considering the evidence for or against *p* (Coliva 2016, pp. 32-33). On this view, we can think of belief-formation as a *function* that takes mental judgments as arguments that produce explicit belief states as values. A belief that *p* results from a mental act of judging that *p* is true.

To illustrate these points, consider a case in which I possess good evidence supporting the truth of the proposition, *that the ocean tides are rising*. Call this proposition, "*p*". My evidence might consist of observations, empirical measurements, testimony, and so on. If I take my evidence to be sufficient to support the truth of *p*, then, other things being equal, I can come to judge that that *p* is true. My judgment results in a mental state, namely, a belief. Consequently, I represent a certain state-of-affairs as obtaining. Specifically, I consider it to be true that *the ocean tides are rising*. The associated judgment is *not* something I rationally make by basing it solely on my pragmatic considerations. Rather, my mental judgment concerning the truth of *p* rationally occurs in such a way that I cannot voluntarily ignore my available evidence for *p*. It is this judgment that automatically results in a belief that *p*.

One final clarification is in order. Judging that *p* is true based on one's evidence is typically sufficient to bring about a belief that *p*. However, judging that *p* is *true* might not be necessary for belief. Given feature (B1), an agent can come to believe *p* by regarding *p* as being likely or more probable than not-*p*.[25] Hence, belief-formation can involve a wider range of judgments. One could

[25] On my conception of *belief*, judgments are typically associated with belief-formation. This, however, does not suggest that there can be other ways to characterize how agents acquire beliefs. Quassim Cassam, for instance, argues that performing the mental act of judging that *p* is true is not a necessary condition for producing a belief that *p* (Cassam 2010, pp. 81-82). Cassam argues for the possibility that one could judge a proposition, *p*, to be true, and yet this mental act can fail to generate a corresponding belief due to "non-rational factors such as self-deception, prejudice and phobias" (Cassam 2010, p. 82). What is important to note, however, is that Cassam is operating with a different

form a belief that *p* by either judging that *p* is true, or that *p* is the case, or that *p* is highly probable, or that *p* is more likely than its negation.[26] Consequently, we can identify the following feature of belief:

> **(B3)** An agent's consciously formed rational belief that *p* is typically the product of a mental act of judging that *p* is true, or that *p* is the case, or that *p* is more likely than not-*p*, such that she cannot voluntarily will herself to believe *p* by disregarding the evidence and epistemic reasons for the truth of *p*.

I contend that features (B1) – (B3) are the most salient features of explicitly held belief. These features distinguish rational belief from other cognitive endorsements that agents can possess. Of course, this is not to say that (B1) – (B3) provide necessary and sufficient conditions for our concept of *belief*. There may be other necessary features of this species of cognitive endorsement.

Before concluding this section, I briefly discuss some of the important ways in which beliefs shape individual inquiry and reasoning. To believe *p* is to take a stance on the truth of *p*. But as Bratman notes, one can "agglomerate one's various beliefs into a larger, overall view" (Bratman 1999, p. 19). Doxastic agglomeration enables agents to form a deeper understanding of how pieces of believed information can relate to one another. Moreover, most theorists agree that our beliefs can function as

---

conception of *belief*. On my conception, occurrent beliefs are *typically* the products of judgments that something is true, or that something is likely or more probable than not. Moreover, it is possible to explain cases in which one judges that *p* is true while failing to manifest the relevant dispositions and behavioral traits by appealing to a distinction between explicitly formed beliefs and unconscious dispositional beliefs. According to Coliva, we can distinguish between full beliefs that one consciously forms via judgments that certain propositions are true (she refers to these as "beliefs as commitments") and unconscious, dispositional beliefs that an agent merely possesses (Coliva 2019, pp. 4-7). If one fully believes *p* by consciously judging *p* to be true, then this belief can sometimes conflict with one's dispositional beliefs (Coliva 2019, p. 5). One might, for example, act in ways that are inconsistent with one's full belief that *p* due to having dispositional beliefs that are inconsistent with the truth of *p*. This conflict, however, does not suggest that one lacks the full belief that *p* (Coliva 2019, p. 5). Rather, it suggests that one can possess different *kinds* of beliefs at a given time that can conflict with each other. At any rate, we can maintain that explicitly formed beliefs are typically the products of judgments.

[26] This quality of belief, like (B1), could apply *mutatis mutandis* to some accounts of degrees of belief. An advocate of a "threshold" view of degrees of beliefs could understand (B3) as claiming that one believes *p* when one regards or judges *p* to be more likely than not-*p* in such a way that one has a credence towards *p* above some stipulated threshold (e.g., a credence above 0.5).

premises in theoretical and practical reasoning.[27] When one consciously forms a belief that *p* based on one's available evidence, one is essentially settling matters about the truth or likelihood of *p*. Hence, one puts oneself into the cognitive position to utilize *p* in future inquiry and reasoning.[28] That is to say, one is normally able to posit *p* as a premise in appropriate cases of reasoning.[29] These "appropriate" cases typically involve inquiry and reasoning aimed at forming new beliefs, knowledge, or understanding. This does not mean that one will always feel compelled to employ one's beliefs in thought and reasoning.[30] Rather, the point is that if one believes *p* and is put into a situation where one needs to inquire into matters involving *p*, then one will (other things being equal) be cognitively able to utilize *p* as a premise.[31] Beliefs are therefore an important part of inquiry and reasoning.

2. ACCEPTANCE

Belief is one kind of cognitive endorsement. Many theorists insist that we should countenance another belief-like attitude called "acceptance". To accept *p* is to be in a mental state of endorsing *p*. There is, however, little consensus among philosophers on how to properly characterize this mental

[27] See, for example, Harman (1986).

[28] Note that other kinds of beliefs may not put one into a cognitive position to employ the contents of those beliefs in conscious reasoning and inquiry. Dispositional beliefs might only guide our unconscious reasoning and behavior.

[29] This idea is inspired by Catherine Elgin's (2017) view on acceptance and Alan Millar's (2004) view on belief. Elgin adopts a modified version of Cohen's notion of *acceptance*, which she distinguishes from *belief* (Elgin 2017, p. 19). She claims, "[t]o accept that *p* involves being willing to take *p* as a premise, as a basis for action, or, I add, as an epistemic norm or rule of inference, when one's ends are cognitive" (Elgin 2017, p. 19). According to Millar, an agent's belief will engender a "psychological commitment" to posit the target proposition of that belief "as an assumption in [her] thinking, should the need to do so arise" (Millar 2004, pp. 122-123). On the view advanced here, both beliefs and acceptance can be employed as premises in reasoning. I say more about acceptance in the next section. I will discuss Millar's view regarding beliefs and commitments in chapter 2.

[30] One might not be willing to employ one's beliefs to make trivial inferences. As Harman notes, we should avoid filling our minds with trivial information (Harman 1986, p. 12).

[31] This does not mean that a belief will cause its possessor to *correctly* employ that belief as a premise in her reasoning. One can reason incorrectly.

endorsement. Many philosophers employ different conceptions of acceptance for different purposes.[32] Moreover, the relationship between acceptance and belief is not entirely clear. According to one view, belief is a kind of acceptance (Harman 1986; Stalnaker 1987; Velleman 2000).[33] Stalnaker, for instance, claims that the notion of *acceptance* "is a broader concept than belief; it is a generic propositional attitude concept with such notions as presupposing, presuming, postulating, positing, assuming and supposing as well as believing falling under it" (Stalnaker 1987, p. 79). On the other hand, some views take acceptance to be a kind of belief (Lehrer 1990).[34]

---

[32] I will discuss some of these different conceptions in passing. However, I will not focus on the account of acceptance within scientific inquiry that is advanced by van Fraassen. According to van Fraassen, if an agent, *S*, accepts a (scientific) theory, *T*, then *S*'s attitude toward *T* is a belief that *T* is "empirically adequate," and not necessarily that *T* is true (van Fraassen 1980, pp. 4, 12, 46). And, according to van Fraassen, for a theory, *T*, to be *empirically adequate* simply means that *T* accurately describes only the *observable* objects and states-of-affairs that *T* is supposed to capture (van Fraassen 1980, pp. 12). Now, van Fraassen's view is a bit more complicated than this. But as I understand it, if empirical adequacy of a theory, *T*, only requires that *T* accurately describes the relevant observable objects, then this does not entail that the entire theory is true. As it currently stands, van Fraassen's account does not adequately describe our ordinary concept of *acceptance*. Outside of our scientific practices, paradigmatic attitudes of acceptance typically concern ordinary propositions rather than empirical theories. Moreover, a typical acceptance that *p* is not primarily based on a belief or judgment that *p* is *empirically adequate*. One can accept many types of propositions – that is, one can accept empirical propositions, moral propositions, mathematical and logical propositions, and so on. One might not judge these propositions to be empirically adequate. Indeed, it will be shown that one can accept a proposition for nonevidential or pragmatic reasons.

[33] Harman, for instance, has used "belief" and "full acceptance" interchangeably (Harman 1986, pp. 46-47). He claims that one's "tentative acceptance of a working hypothesis" or proposition can sometimes turn out to be "fruitful" or well-supported, such that one might be justified in "fully accepting" it, and thus believing it (Harman 1986, pp. 46-47). This suggests that belief is, on Harman's view, a kind of acceptance that is sufficiently justified or evidentially supported. Velleman has also argued that belief is a kind of acceptance. He claims that "to regard a proposition as true…is to accept the proposition" (Velleman 2000, p. 184). He thinks that both beliefs and assumptions are states with "two-tiers:" an underlying first tier of acceptance with a second tier incorporating a specific "aim" (Velleman 2000, p. 184). For Velleman, believing that *p* is a form of acceptance that incorporates a particular aim toward *p*'s truth, whereas assuming that *p* is a form of acceptance that incorporates an aim that is "indifferent to [*p*'s] truth" (Velleman 2000, pp. 184-185). I am sympathetic to the idea that these states have different aims. But I will ultimately defend a different view that distinguishes belief, acceptance, and assumption.

[34] It is not entirely clear if Lehrer thinks that acceptance in general is ultimately reducible to belief. But there are some passages that suggest that he takes acceptance (or at least some forms of acceptance) to be a kind of belief. He claims that the notion of *acceptance* can be characterized as a state that aims at "attaining truth and avoiding error" (Lehrer 1990, pp. 10-11). Lehrer then goes on to say that we

In this section I address these concerns. I articulate a congenial account that correctly describes explicitly formed attitudes of acceptance. Acceptance, on this view, is a cognitive endorsement that differs from belief. The concept of *acceptance* that I focus on here is a technical notion that may not correspond to the way we typically employ this term in ordinary language. But the notion of *acceptance* that I defend picks out a genuine phenomenon of endorsing propositions in thought and reasoning. This view is partly based on the account of acceptance advanced by Cohen (1989, 1992). I argue, however, that my account accurately captures a wider range of phenomena that can be correctly characterized as attitudes of acceptance. Acceptance is a broad conceptual category that picks out different species of explicit mental attitudes we form in conscious deliberation and reasoning. At a first glance, accepting *p* involves a positive stance of endorsing *p*, and a mental resolve or commitment to implement and manipulate *p* in either thought or reasoning. Accepting *p* can sometimes accompany a belief that *p*. However, this need not always be the case.[35] It is plausible that much of the debate over the nature of acceptance and its relationship to belief is merely apparent or semantic. Many other accounts of acceptance might plausibly fall under the concept I advance here.

At its core, our ordinary notion of *acceptance* seems to involve a kind of mental endorsement. Larry Laudan, for instance, claims that to accept a proposition is to "*treat it as if it were true*" (Laudan 1977, pp. 108-109).[36] But what does it mean to treat a proposition *as if it were true*? To believe *p* is to regard *p* as being true. However, treating *p as if* it were true need not require one to think that *p* is, in

---

can understand this attitude of acceptance "to be a special kind of belief" (Lehrer 1990, p. 11). In what follows, acceptance will be characterized as a distinct state that is not reducible to belief.

[35] Many theorists hold that acceptance can come apart from belief (see, for example, Cohen 1992).

[36] Laudan is primarily concerned with the attitude of acceptance that scientists can have toward theories, which he distinguishes from an attitude in a "*context of pursuit*" (Laudan 1977, pp. 108-113). Unlike scientific theories that are accepted, a theory that is merely pursued is often "less acceptable, less worthy of belief" (Laudan 1977, p. 110). Laudan claims that when one pursues a theory, one typically does so for pragmatic reasons (e.g., the theory might, according to Laudan, "be capable of generating new solutions to problems at an impressive rate") (Laudan 1977, p. 111). In what follows, I focus primarily on our explicit, ordinary attitudes of acceptance that we can possess in many domains of inquiry (i.e., I will not be primarily concerned with acceptance of scientific theories).

fact, true. The type of endorsement essential to acceptance is different. At a first approximation, if accepting *p* is distinct from believing *p*, then to accept *p* would seem to involve a mental positing of *p* in thought, such that this state is compatible with being agnostic toward the truth of *p*.

What sort of mental endorsement and positing is sufficient for acceptance? We might follow Cohen here and maintain that acceptance involves employing a proposition in a line of reasoning and inquiry (Cohen 1989, 1992). He defines this attitude as follows:

> …to accept that *p* is to have or adopt a policy of deeming, positing, or postulating that *p* – i.e., of including that proposition or rule among one's premisses for deciding what to do or think in a particular context, whether or not one feels it to be true that *p* (Cohen 1992, p. 4).

By employing *p* as a premise in one's reasoning one treats *p* as if it were true.

Several points are worth noting. On Cohen's view, accepting *p* necessarily involves a mental policy of employing *p* in one's reasoning and inquiry. Accepting *p* is understood to be a temporally extended mental *activity* of reasoning with *p*. As Keith Frankish puts the point, Cohen's notion of *acceptance* is more akin to a "*behavioral* state" (Frankish 1998, p. 251, italics original). If an agent, *S*, accepts *p*, then *S* will possess a disposition or psychological compulsion to act in a certain way. Specifically, *S* is disposed to mentally posit *p* as a premise in her theoretical and practical reasoning.[37] Many theorists adopt some version of this view. For instance, Catherine Elgin claims that when an agent accepts *p*, her acceptance "involves being willing to take *p* as a premise, as a basis for action, or…as an epistemic norm or rule of inference, when one's ends are cognitive" (Elgin 2017, p. 19).

Consequently, Cohen's view of acceptance fundamentally differs from the conception of *belief* defended in section 1. Belief is a representational *state*. It is not simply a mental activity of employing

[37] Moreover, *S* takes on a disposition or commitment to deduce and accept those propositions that *S* recognizes as following from *p* (Cohen 1992, pp. 28-29). I discuss the commitments associated with acceptance in chapter 2.

a proposition in reasoning. Moreover, acceptance typically involves a voluntary *decision* to use the target proposition as a premise in reasoning (Cohen 1992, p. 22). Cohen asserts the following:

> Acceptance, in contrast with belief, occurs at will, whether by an immediate decision or through a gradually informed intention. This is because at bottom it executes a choice – the accepter's choice of which propositions to take as his premisses (Cohen 1992, p. 22).

Thus, if an agent, *S*, accepts *p*, then *S* decides to employ *p* in her reasoning. This mental act is a decision (or intention) to bring about a certain state-of-affairs, namely, a mental process of reasoning with *p*. Consequently, acceptance is intimately connected to an agent's *will*.[38] Finally, an attitude of acceptance, as Cohen defines it, can be sensitive to a wider range of reasons and rational support. Cohen argues that acceptance need not be sensitive to *only* epistemic or evidential reasons. Rather, an acceptance that *p* is responsive to both *practical* and *evidential* reasons in order to contribute to our practical and theoretical reasoning (Cohen 1992, pp. 4-5, 12).[39] Many theorists correctly point out that, contrary to belief, an agent's acceptance that *p* can be restricted to certain contexts in which she has certain practical considerations to posit *p* as a premise in her reasoning (Stalnaker 1987, pp. 80-81; Bratman 1999, pp. 20-27). Bratman, for instance, argues that "practical pressures" within some context plays an important role in motivating one to accept a proposition that functions as a premise in one's practical reasoning and deliberation (Bratman 1999, pp. 20, 27-28).

Cohen's account provides a good starting point for characterizing acceptance. Mental policies of premising propositions might be sufficient to bring about attitudes of acceptance. However, it is not clear if adopting these mental policies sufficiently captures all the ways that agents can *accept*

---

[38] Other philosophers defend similar views of acceptance. See, for example, Laudan (1977) and Bratman (1999). Larry Laudan argues that, within our scientific practices, an individual scientist can make a *choice* to accept a given theory or proposition, such that she will "*treat it as if it were true*" (Laudan 1977, pp. 108-109). According to Bratman, acceptance "can be subject to our direct voluntary control" (Bratman 1999, p. 27). I take it that acceptance is typically the result of a voluntary decision or choice. However, it might be possible for involuntary decisions to accept a proposition.

[39] Bratman endorses a similar view (see Bratman 1999, pp. 20-21).

propositions. My contention is that his account is too narrow and fails to adequately capture all cases of acceptance. For it seems possible for one to accept *p* without explicitly taking on any mental policy of actively postulating *p* as a premise in reasoning. Perhaps one lacks interest in what follows from *p*. Or perhaps one is not in an appropriate cognitive or epistemic situation to use *p* in one's inferences. Nevertheless, it is possible for one to take on some cognitive attitude of endorsing a proposition, *p*, that (1) is not a state of believing *p*, and (2) does not involve a mental policy or commitment to employ *p* as a premise in reasoning. This kind of attitude would seem to count as a state of acceptance.

Consider, for example, the following case.

> ***Sheep:*** While visiting a national park, Brooke sees a sign she believes to be authored by the local rangers stating, "Reduction in the Bighorn population is probably due to pneumonia". Brooke, however, only has a minimal understanding of the phenomenon of pneumonia. At most, she thinks that pneumonia is probably an illness. Nevertheless, she recognizes that the sign is asserting something that might be true. Moreover, she recognizes that the statement comes from a reliable source. Brooke decides to endorse the proposition, *that the reduction in the Bighorn population is due to pneumonia*. She mentally puts this proposition forward in her thinking for further inquiry. However, she is not *committed* to regard this proposition as being actually true. She proceeds to ask the rangers for more information about the Bighorn population.

In the *Sheep* case, Brooke does not believe the proposition *that the reduction in the Bighorn population is due to pneumonia*. Call this proposition, "*p*". Brooke fails to believe *p* due to the fact that the authors of the statement, who she takes to be reliable, are not *convinced* that pneumonia is the cause of the reduction in the Bighorn population. Her available evidence does not warrant a belief that *p*. She could, however, come to believe a different proposition, namely, *that the reduction in the Bighorn population is probably due to pneumonia*. Call this proposition, "*q*". In this case, *q* has a different content: pneumonia is only characterized as a *probable* cause of the population's decline. Nevertheless, Brooke still takes some positive stance of endorsement toward *p*. For she could still posit *p* as being *given* in her thinking.

My suggestion is that Brooke's attitude toward *p* is one of *acceptance*.[40] First, she takes on some non-belief attitude of *endorsing p*. Second, this attitude is based on her available reasons. These reasons can be both epistemic and pragmatic. She has epistemic reasons for *p*'s truth, namely, testimonial evidence. Moreover, she has practical reasons to endorse *p*, namely, a desire or goal to know more about the content of *p* and a goal to make an accurate judgment concerning *p*'s truth.

However, it is incorrect to say that Brooke explicitly adopts a mental *policy* of employing *p* as a *premise* in a line of reasoning. First, she is unclear as to what pneumonia is and how it affects sheep. Second, it is not clear that Brooke has any reason or interest to employ *p* as a *premise* in an explicit line of reasoning. Rather, Brooke seems to merely endorse *p* by deciding to enter into a *mental state* of taking a positive stance toward *p*. This stance could involve putting *p* forward in her thought in such a way that she is (1) mentally fixing *p* as given, (2) disposed to be receptive to further information and evidence concerning the truth of *p*, and (3) willing to utilize *p* or act on *p* when she is able and interested in doing so. By fixing *p* in her thought, she is prepared to enter into a cognitive position in which she would be willing to either incorporate *p* into her system of beliefs, or simply manipulate *p* in her reasoning (either theoretical or practical). This, however, does not imply that she is currently in any position to adopt a mental policy of employing *p* as a premise in her reasoning. If she eventually comes to adequately understand what pneumonia is, and if she has a goal of reasoning about the reduction in Bighorn sheep, then she could adopt a mental policy of premising *p* in her reasoning.

[40] It might be argued that if an agent, *S*, consciously considers a proposition, *p*, that she does not understand, then the only cognitive endorsement that *S* can possess concerning *p* is one of acceptance. That is to say, *S* cannot believe propositions that she does not fully understand. Edna Ullmann-Margalit and Avishai Margalit seem to defend a similar view (Ullmann-Margalit and Margalit 1992, pp. 167-168). I do not share this view (I will say more about their view of acceptance). The account of belief defended in section 1 allows for agents to believe propositions that they partially understand.

It stands to reason, then, that mental policies of premising propositions are not necessary for acceptance. Acceptance is a broader conceptual category picking out a wide range of mental states and processes. To accommodate this intuition, I propose that we modify Cohen's view as follows:

> **Acceptance:** To accept *p* involves taking a positive mental stance toward *p*, along with a mental resoluteness and disposition to posit and implement *p* in thought or reasoning, given one's available reasons for *p*.[41]

Several clarifications are in order. First, a token attitude of acceptance is fundamentally a mental *state*, and not merely a behavioral state of reasoning. Following Alston, acceptance is a state involving "the *taking on* of a positive attitude toward a proposition" (Alston 2007, pp. 132).[42] On my account, this positive stance is one of *endorsement*: one accepts *p* by endorsing *p* in one's thought. This endorsement can be understood broadly as a type of mental affirmation or positing of the target proposition. Moreover, this mental stance can involve a range of dispositions: by accepting *p*, an agent will typically possess a disposition to consider *p* as either plausible, credible, approximate, useful, or beneficial for her when she contemplates *p* and her available reasons for positing *p*.[43]

To be clear, these dispositions need not always be manifested when accepting *p*. Moreover, dispositions to consider a proposition, *p*, as plausible, approximate, useful, or beneficial might not be

---

[41] It might be objected that possessing a mental resoluteness is too strong for acceptance. That is to say, acceptance could simply involve having a disposition to posit and implement a proposition in thought or reasoning. However, I am focusing on a type of explicit and consciously formed acceptance that mature agents can possess. In these cases, one will consciously take on some kind of mental resoluteness or determination to employ the target proposition in thought for one's reasons. Moreover, merely having a disposition to use a proposition in thought seems insufficient for genuine acceptance. One might simply momentarily assume or suppose *p* by having a tentative disposition to use *p* in thought. But accepting *p* seems to require more. For one to accept *p* means that one will commit oneself to mentally employ *p* for one's reasons in the appropriate contexts.

[42] There are significant differences between Alston's conception of acceptance and my view. I will address these differences below.

[43] Intuitively, it seems that to accept *p* involves being disposed to consider *p* as being plausible (see Alston 2007). Accepting *p* therefore differs from merely acting as if *p* were the case (or acting as if *p* were plausible or useful). Merely behaving as if *p* obtained might be better characterized as *behaving as if p* or *taking p for granted*. But this behavioral state does not capture our ordinary intuition that acceptance involves a positive mental stance toward a proposition.

manifested *solely* by the attitude of accepting *p*. Accepting *p* can be accompanied by corresponding beliefs about the truth, plausibility, or usefulness of *p*. For example, an agent, *S*, might fail to believe *p* and yet still come to accept *p* given that *p* is practically beneficial. In this case, belief and acceptance come apart. However, *S*'s reasons for accepting *p* might prompt her to form a corresponding belief that *q*, where the content of "*q*" is *that p is practically beneficial*. Hence, *S*'s acceptance that *p* and belief that *q* will both contribute to *S*'s disposition to take some positive mental stance toward *p*.

We can generalize this point. By accepting *p*, one might also form a belief concerning the rational appropriateness of *p* given one's current situation and reasons.[44] More precisely, while an agent might not take her available evidence to support the truth of *p*, she may nevertheless take her reasons to warrant a belief that *p* is rational or appropriate to endorse. An agent can believe that *p* is rationally appropriate to endorse for a variety of ways that does not subsequently imply that she also regards *p* as being true. One might judge *p* to be practically useful, plausible, appropriate, credible, and so on. These judgments can terminate in corresponding beliefs – i.e., a belief *that p is useful*, or a belief *that p is plausible*, or a belief *that p is appropriate*, or *that p is credible*, and so on. The propositional contents of these beliefs differ from the proposition, *p*, that is accepted and implemented in thought. Furthermore, the associated beliefs one forms when accepting propositions can partly explain the positive mental stance an agent takes when she consciously considers the target proposition of her acceptance.

Second, it is also important to keep in mind that, on Cohen's account, acceptance involves a "commitment to a pattern, system, or policy – whether long or short term – of premising" a proposition in one's reasoning or "as a basis for a decision" (Cohen 1992, p. 12). *Pace* Cohen, I have argued that this is not a necessary feature of acceptance. A commitment to employ *p* as a premise is only one way to accept *p*. It is best to say that acceptance can put one into the cognitive position to

[44] This is similar to Bas van Fraassen's (1980) view of the acceptance of scientific theories, in which an acceptance of a theory, *T*, typically involves believing *T* to be "empirically adequate" (van Fraassen 1980, pp. 12, 46). See footnote 32.

employ the target proposition in reasoning and inquiry, depending on one's current goals, interests, beliefs, and understanding. On my account, the relevant commitment or resolution involved is much broader: acceptance results in a mental determination to implement or use the target proposition in thought, reasoning, or inquiry. To *implement* a proposition, *p*, in thought is to be in a cognitive state of postulating and using *p* in one's thinking, reasoning, or inquiry. This mental postulating is typically accompanied by a disposition to take a positive mental stance toward *p*. Postulating and using a proposition is a cognitive ability that can be carried out in different ways. Hence, implementing *p* is to undergo a certain kind of mental resolution: one takes on a commitment or mental disposition to execute, invoke, manipulate, or uphold *p* in one's thinking. This can involve either (1) a willingness to endorse and uphold *p* in one's thought; (2) a disposition to be sensitive and receptive to new evidence, information, or practical considerations concerning *p*; (3) a deliberate mental integration of *p* into a body of information; or (4) a readiness to either act on *p*, come to form a belief that *p* (if her available evidence warrants the truth of *p*), or employ *p* as a premise in one's theoretical or practical reasoning.

This broader account provides a plausible characterization of acceptance. First, it coheres with our ordinary intuitions concerning how agents come to accept propositions. Accepting a proposition, *p*, does not merely involve contemplating *p*. Rather, one takes some sort of positive stance toward *p* – i.e., one endorses *p*. Moreover, accepting *p* would involve some kind of resolution to implement *p* in one's thinking based on one's available reasons. Indeed, we typically take acceptance that *p* to involve more than a momentary endorsement of *p*. Acceptance plays a more substantive role in our thinking and reasoning. When one accepts *p,* one executes *p* in thought as being epistemically plausible, effective, or practically beneficial for oneself in a given situation. Acceptance therefore plays a crucial role in reasoning, belief-formation, and action. Second, it seems correct to say that an attitude of acceptance need not always involve an explicit mental commitment to posit a proposition in reasoning and inference. To say otherwise would be overly strict with the notion of *acceptance*. In the *Sheep* case

above, the agent, Brooke, seems to accept a proposition, *p*, without being committed to explicitly using *p* as a premise in come line of reasoning. Instead, she accepts *p* by (1) taking a positive mental stance toward *p* given her available reasons, and (2) fixing and using *p* in her thinking. She undergoes a commitment to uphold *p* as being credible. Moreover, she integrates *p* into her background information. Finally, she is disposed to be receptive to new information and evidence concerning *p*.

Consequently, accepting *p* is fundamentally different from believing *p*. First, to believe *p* requires its possessor to regard *p* as being either true, the case, or more likely than not-*p*. Accepting *p* involves *using p* in thought, and being disposed to take a positive stance toward *p*. However, this positive stance can be compatible with an attitude of agnosticism toward the truth of *p*. One might accept *p* and merely consider *p* to be pragmatically useful. Hence, acceptance has the following feature:

(A1) To accept *p* is to implement *p* in thought and to be disposed to adopt a positive stance toward *p*, such that one can still be agnostic toward the truth of *p*.

Second, an agent's rational acceptance, like a rational belief, is based on her available reasons. The crucial difference between states of acceptance and attitudes of belief is that acceptance is rationally sensitive to both epistemic and practical considerations. For instance, one might possess some epistemic reasons for accepting a proposition, *p*. These epistemic reasons may not be sufficient to warrant or cause a belief that *p*. However, one could have some further pragmatic considerations to rationally motivate one to accept *p*. Indeed, one could accept *p* for *predominately* pragmatic reasons. In these cases, the pragmatic considerations for accepting *p* will outweigh the available epistemic reasons in favor of *p*. Perhaps one accepts *p* because doing so is a way to achieve some desired goal. In these cases, the nonevidential reasons are not indicators of the truth or falsity of *p* (see Foley 1993, p. 16). Hence, one might only be in the position to rationally regard *p* as being practically useful for achieving one's ends.

To illustrate this idea, consider the following case.

> ***Bridge:*** Christine needs to cross the river. She can either take the bridge or she can take the ferry. Christine becomes aware of a rumor from an unknown source that the bridge was blown up. She does not believe that the rumor is true. Nor does she believe that the rumor is false. However, she decides to employ the proposition, *that the bridge is blown up* as a premise in her practical reasoning. Given her desire to both cross the river and remain safe, her resulting acceptance informs her process of reasoning to determine what she ought to do. In this case, she decides to take the ferry.[45]

Let "*p*" refer to the proposition, *that the bridge is blown up*. In the *Bridge* case, Christine lacks a belief that *p*.[46] However, given that she endorses *p* and commits herself to utilize *p* in her reasoning, Christine comes to accept *p*. Her acceptance is based on epistemic and pragmatic reasons. Her epistemic reasons consist of the rumor and her background beliefs. And her pragmatic reasons are based on her goal to cross the river in a safe manner. However, given Christine's situation, her pragmatic considerations for accepting *p* could outweigh her epistemic reasons for accepting *p*.[47] If this is the case, then her acceptance is primarily grounded in her pragmatic reasons.

Christine's acceptance in the *Bridge* case seems entirely appropriate. It follows that attitudes of acceptance are rationally sensitive to a broad range of reasons. Both epistemic and pragmatic considerations can rationally motivate and ground a given acceptance that *p*. This coheres with other mainstream views of acceptance.[48]

Considering that acceptance is rationally sensitive to both epistemic and pragmatic reasons, it seems correct to say that there can be different *functions* of acceptance. That is to say, one can accept and utilize *p* in one's thinking in different ways and for different purposes, depending on one's reasons,

---

[45] This is a modified version of a case that Bratman provides (Bratman 1999). Bratman's example concerns a situation where one accepts an overestimation of building costs to determine if one can afford to complete the building project (Bratman 1999, pp. 22-23). From this example, Bratman argues that acceptance and belief come apart (Bratman 1999, pp. 23-24).

[46] Of course, her situation and available reasons can also prompt her to form a corresponding belief *that the bridge might be blown up*. For further discussion of this possibility see pp. 33 – 34 in this chapter.

[47] I am appealing to Joseph Raz's view that reasons can be "stronger or more weighty than others" (Raz 1975, p. 25).

[48] See, for example, Cohen 1992, pp. 4 – 5, 12.

motivations, and intended goals. We can therefore draw a distinction between *epistemic acceptance* and *pragmatic acceptance*. I discuss each of these species in turn.

An *epistemic acceptance* is based on epistemic or evidential considerations for its target proposition. Moreover, it is paradigmatically formed in conjunction with an intended epistemic goal. Specifically, it is typically employed in inquiry, investigation, and hypothesizing. In some cases, an agent's epistemic acceptance that *p* can coincide with a belief that *p*. To illustrate this species of acceptance, consider a case where a detective is investigating a murder. Suppose that the detective discovers that the murder victim's clothes was imbedded with silt and strands of fishing line. The detective might fail to believe *that the murderer works at the wharf*. Call this proposition, "*p*". In this case, her available evidence is not sufficient for believing *p*. Nevertheless, she comes to accept *p*. Accordingly, she endorses *p* and considers *p* to be credible. Moreover, she commits herself to use *p* in her long-term thinking and reasoning for her epistemic goal of investigating and settling matters concerning *p*'s truth. In this situation, the detective has an epistemic acceptance.

On the other hand, one might *pragmatically accept* a proposition, *p*. A *pragmatic acceptance* is fundamentally based on practical considerations. In these cases, rational acceptance that *p* is grounded on pragmatic considerations in favor of *p* that outweigh the epistemic reasons one has for or against the truth of *p*.[49] That is to say, endorsing *p* would make one better off in some sense. Furthermore, a pragmatic acceptance is not conceptually tied to any specific epistemic goal. One might simply endorse and implement a proposition in thinking in order to guide one's behavior and achieve some practical end. On this view, a pragmatic acceptance that *p* is similar to *taking p for granted*.[50]

---

[49] Epistemic reasons can still partially ground pragmatic acceptances. The important point is that this type of acceptance is based on pragmatic reasons that outweigh the relevant epistemic reasons.

[50] My account of pragmatic acceptance is similar in some respects to Ullmann-Margalit and Margalit's account of "holding true" (Ullmann-Margalit and Margalit 1992, pp. 167-187). For them, holding a proposition, *p*, as true is primarily a practical affair involving one to act or behave as if *p* were true (Ullmann-Margalit and Margalit 1992, p. 171). Similarly, pragmatic acceptance will typically guide one's

For example, consider a case where a janitor is told that many of the second-floor rooms in a hotel have rats.[51] The janitor enters room 237 and notices that there are no signs indicating that the room has a rat problem. This, however, does not suggest that the room is in fact free of rats. Consequently, the custodian fails to believe the proposition, *that room 237 has a rat problem*. Call this proposition "*q*". Nevertheless, the custodian has strong practical reasons for endorsing *q*. For the hotel has a policy of delivering excellent customer service. Allowing any rats in the rooms would conflict with this policy. These pragmatic considerations outweigh her epistemic reasons supporting *q*. At any rate, the custodian still comes to accept *q*. She mentally endorses *q* and commits herself to implement and use this proposition in her long-term thinking. Her acceptance is employed to guide her behavior, and it enables her to provide the best possible customer service. In this situation, the custodian has a *pragmatic acceptance*.[52] She accepts *q* (or mentally takes *q* for granted) for pragmatic considerations in order to obtain a practical goal that she has.[53]

This broad sensitivity to different kinds of reasons is a distinguishing mark of acceptance. The type of reasons and goals one has will dictate the way in which one will rationally accept a proposition. One can pragmatically accept *p* for practical reasons and goals. On the other hand, one can epistemically accept *p* for evidential considerations and epistemic goals. The specific type of acceptance involved will partly determine how one implements the target proposition in thought.

---

behavior to reach some practical goal. On my account, however, it is conceivable that pragmatic acceptance can, in some cases, play some theoretical role.

[51] The following case is similar to cases that other philosophers have described. Cohen, for example, presents a case in which an attorney accepts that her client is innocent to perform her duties even though she does not believe that her client is innocent (Cohen 1992, p. 25). Ullmann-Margalit and Margalit describe a situation in which, contrary to what one actually believes, one can still "hold true" the proposition, *that the customer is always right*, so as to deliver quality service (Ullmann-Margalit and Margalit 1992, pp. 170 – 171).

[52] As noted earlier, epistemic reasons can play a role. However, pragmatic acceptances are primarily based on pragmatic reasons that outweigh the epistemic considerations for the target proposition.

[53] This is consistent with other accounts that hold that an agent can accept *p* for pragmatic reasons even when she fails to believe *p*. See, for example, Cohen 1992.

Acceptance therefore has the following feature:

> **(A2)** A rationally held attitude of accepting *p* is based on reasons. These reasons can be either epistemic or practical considerations in favor of *p*. A *epistemic acceptance* is primarily based on evidential reasons an agent has, and it is formed in pursuit of an epistemic goal. A *pragmatic acceptance* is primarily based on practical reasons an agent has, and it is formed in pursuit of some pragmatic goal.[54]

The final fundamental feature of acceptance that I will discuss here involves the connection between acceptance and the will. Many theorists hold that acceptance is voluntarily brought about.[55] Hence, acceptance differs from beliefs. For one cannot rationally believe *p* by voluntarily ignoring one's epistemic considerations for or against *p*. Contrary to this, acceptance is rationally motivated by both epistemic and practical reasons. One's motivations for accepting *p* need not be directed at truth. Hence, accepting *p* need not require its possessor to take a positive *epistemic* stance toward the truth of *p*. It follows from this that an agent, *S*, can accept *p* by voluntarily choosing or intending to mentally fix *p* in her thought, such that she need not take into consideration her available epistemic considerations for or against *p*. Acceptance therefore has the following feature:

> **(A3)** An acceptance that *p* can be the result of a voluntary choice or decision: an agent can will herself to accept *p* in such a way that she ignores the epistemic reasons for or against *p*, or she does not base her acceptance on good evidence.

I take it that features (A1) – (A3) are sufficient to show how acceptance can differ from belief. Consequently, acceptance is a distinct cognitive endorsement.

One benefit of this view is that it is appropriately robust. First, it does not restrict acceptance to mental policies of premising propositions in reasoning. Consequently, this account can adequately accommodate a wide variety of cases. Second, this account does not restrict attitudes of acceptance to cases of *practical* reasoning. According to Edna Ullmann-Margalit and Avishai Margalit, acceptance

---

[54] In most cases one will have a mixture of both epistemic and practical reasons for one's acceptance that *p*. For example, if an agent, *S*, takes her available evidence to be insufficient to warrant belief, then she might decide to accept *p* to achieve some desired end. The end might be epistemic or practical. In these cases, *S* will have some practical reason that accompanies her epistemic reasons.
[55] See, for example, Cohen (1992), Bratman (1999), and Alston (2007).

primarily informs one's practical deliberations on how to act (Ullmann-Margalit and Margalit 1992, pp. 175-177). On the present account, however, acceptance can play an important role in theoretical reasoning. One can accept *p* for epistemic considerations to determine what one should think.

Finally, the view adequately accommodates a wide range of "positive" mental stances one can take toward the propositions that we accept. Some theorists maintain that accepting *p* requires taking a positive *epistemic* stance toward *p*. For example, Alston insists that accepting *p* requires "*taking* [*p*] to be true", such that one cannot simultaneously take *p* to be false (Alston 2007, p. 132). He claims that for an agent to accept *p*, she must consider *p* to be supported by some epistemic considerations (Alston 2007, p. 132). On the contrary, the view I am defending here allows for one to accept *p* while rationally believing not-*p*. In these cases, one would pragmatically accept *p* in a given context for pragmatic considerations. Hence, my view of acceptance is more inclusive and robust.

Before concluding this section, I want to briefly canvas a few further details concerning acceptance, its relation to belief, and its role in our epistemic lives. First, while acceptance is a distinct kind of cognitive endorsement, it can be related to belief in several important ways. As noted earlier, one might accept *p* and also believe that *p is plausible*. Hence, accepting propositions can occur within a network of beliefs with similar propositional contents. Moreover, one might accept *p* to include that proposition as a postulate (perhaps tentatively) into one's current system of beliefs and acceptances. Acceptances can therefore supplement our world pictures and our systems of cognitive endorsements.

It is also worth noting an important similarity between belief and acceptance. Beliefs can be occurrent or dispositional. The account defended here allows for dispositional acceptances. Typically, one will occurrently hold an attitude of acceptance. However, it is plausible to continue attributing an agent with an acceptance that *p* in situations where that agent is no longer consciously entertaining *p*. Consider a case where an agent accepts many propositions. This agent can hold these attitudes without having to continuously contemplate each of the target propositions. It seems correct to say that the

non-occurrent attitudes of acceptance could persist *dispositionally* insofar as the agent did not change her mind or rescind her decision to cognitively endorse the relevant propositions. Moreover, it is possible for these dispositional acceptances to play a tacit role in directing behavior and reasoning.

Of course, acceptance that *p* need not enjoy the same stability as a belief that *p*. If one accepts *p* in one context, then this acceptance need not persist in all other contexts. Typically, one will accept *p* given some pragmatic considerations one has for *p* in a certain situation. Hence, one's acceptance can be restricted to those contexts in which one has similar pragmatic considerations in favor of *p*. The crucial point is that acceptance can persist tacitly or dispositionally in certain situations and contexts. These dispositional acceptances, it seems, can help guide reasoning and behavior.[56]

This account provides the conceptual machinery needed to accommodate a wide range of mental states that qualify as being attitudes of acceptance. Moreover, this view elucidates the way acceptance is tied to the manifestation of rational agency. This is not a novel point. According to Elgin, "acceptance involves agency": if an agent accepts *p*, then she can employ *p* as a "premise or rule" to extend her "cognitive objectives" (Elgin 2017, p. 22). This seems correct. On my view, if an agent decides to accept *p* based on her available reasons, then she is willfully determining a cognitive position, stance, or mental course of action. Her acceptance is based on what she takes to be sufficient reasons for endorsing a proposition, *p*. Furthermore, her rational endorsement of *p* can, in optimal situations, put her into the cognitive position to use *p* in her reasoning. In some cases, accepting *p* can lead to new beliefs. Acceptance therefore plays an important role in our thinking and reasoning.

3. SUPPOSITION

Supposition is another kind of cognitive endorsement. On the face of it, supposing *p* involves a weak endorsement of *p*, such that one *tentatively* posits *p* in thought and treats *p* as if it were given. In this

[56] I will say more about dispositional acceptance when discussing presuppositions in the next section.

section I characterize the fundamental features of supposition. Suppositions, I argue, can be distinguished from belief and acceptance. I further argue that suppositions come in two species: rational and arational. *Rational* suppositions are based on, and sensitive to, rational considerations for their target propositions. However, *arational* suppositions can lack a rational basis.

In ordinary discourse, we typically find statements and attributions of the form, "*S* assumes that *p*", and "*S* supposes that *p*." Moreover, one often makes claims of the form, "Let us assume that *p*", or "Let us suppose that *p*." But not all linguistic expressions containing "suppose" or "assume" will refer to the same attitude one employs in thought and reasoning that I am interested in here. For example, expressions like, "I suppose that is true" might be taken to mean that a speaker is indicating to a listener that she has come to agree with some claim or view, or that she has *judged* that some proposition is likely or plausible (such judgments might actually indicate belief or acceptance). Thus, to be clear, I am interested in a type of representational state picked out by our concepts of *assumption* and *supposition* that involves a certain explicit cognitive endorsement of a proposition, where this attitude is not equated to either belief or acceptance. In what follows, I understand this notion of *assumption* to be synonymous with the notion of *supposition*. Hence, a statement of the form, "*S* Φ-that *p*" would have the same meaning if either "supposes" or "assumes" was substituted for "Φ".[57] From this point forward, then, I will use the term "supposition" to refer to this unique epistemic attitude.

Suppositions routinely factor into thinking, reasoning, and belief-revision. To suppose *p* is to be in a representational state. If suppositions are cognitive attitudes in this sense, then their propositional contents can be true or false. If I suppose that koalas typically have marsupiums (i.e., external pouches), and it is a fact that koalas typically have marsupiums, then my supposition is true.

[57] It might be possible to distinguish some types of assumption from suppositions. For instance, in some cases an agent might be credited with a tacit assumption that guides long-term behavior. This kind of assumption is more akin to a dispositional belief (i.e., there's a more entrenched commitment to tacitly treat the target proposition as if it were true). My focus, however, is on consciously formed and held suppositions.

Hence, like belief and acceptance, suppositions can be characterized as states that are correct when their contents are true (see Stalnaker 1987, pp. 79 – 80). That is to say, one correctly supposes *p* if and only if *p* is true.[58]

Moreover, many theorists take these attitudes to play a special role in suppositional reasoning, argumentation, and the construction of proofs (Rescher 1964; Fisher 1989; Levi 1996; Jackson and Jackson 2013; Murphy 2013). For example, Levi claims that an agent, *S*, can *suppose p*, "*for the sake of argument*," such that *S* momentarily treats *p* as if it were a true premise in her reasoning, to see what might follow from *p* (Levi 1996, p. 3).[59] But while many theorists agree on the role our suppositions can play in our thinking and reasoning, there is some disagreement on how we ought to properly conceive of the nature of these attitudes. Some philosophers like Robert Stalnaker and David Velleman take suppositions (or assumptions) to be a kind of *acceptance* toward a proposition (Stalnaker 1987, p. 79; Velleman 2000, p. 112). On the other hand, L. Jonathan Cohen and Michael Bratman take suppositions to be psychological states that are distinguished from both belief and acceptance (Cohen 1992, p. 12; Bratman 1999, p. 27).[60]

---

[58] Stalnaker argues that suppositions are types of acceptances, and that these states are "correct" if they are true (Stalnaker 1987, pp. 79-80). It might be best to resist the view that suppositions can be normatively evaluated as *only* being correct when their target propositions are true. In chapter 2 I show that a supposition generates normative commitments that determine its standards of correctness. Moreover, I show in this section that suppositions can be formed without any rational support, or for purely pragmatic reasons, with no consideration for the truth. Consequently, normatively evaluating all instances of suppositions according to their truth or falsity might not always be appropriate.

[59] On the face of it, this is similar to acceptance. However, I will show that, unlike acceptance, a supposition is a provisional state lacking a commitment to utilize a target proposition in thought.

[60] According to Velleman, a supposition or assumption shares a similar "direction of fit" associated with attitudes like belief, in the sense that when one supposes that *p* one *accepts p* and treats *p* as if it were true (Velleman 2000, pp. 110-112). Of course, Velleman does not think that this attitude of endorsement is the same between belief and supposition: one who supposes *p* need not consider *p* as actually obtaining (Velleman 2000, pp. 111-112). Similarly, Stalnaker claims that supposing *p* is to form a type of *acceptance* in which one treats *p* as being true (Stalnaker 1987, p. 79). Note that my view is different. I will demonstrate that belief, acceptance, and supposition are distinct states.

Following Cohen and Bratman, my aim is to provide a satisfactory account of the *distinct* cognitive attitude of *supposition* that fits within the taxonomy of cognitive endorsements.

Before identifying the salient features of supposition, consider the following case.

> ***Vacation:*** Sue is considering what it would be like to travel to Asia. She considers traveling to Indonesia and Japan. She first supposes that she travels to Indonesia. She knows that Indonesia is a great location for surfing. And she deduces from this that, if she visited Indonesia, then she would go surfing. She then supposes that she travels to Japan. She knows that Japan has great sushi. And she deduces from this that, if she travels to Japan, then she would eat sushi.[61]

In the *Vacation* case, Sue comes to suppose several propositions, namely, the proposition, *that I visit Indonesia*, and the proposition, *that I visit Japan*. Her suppositions act as premises in her reasoning concerning what sort of things and activities she would do if she traveled to either Indonesia or Japan. As Bratman correctly points out, suppositions like Sue's can factor into "contingency planning," which can, in some cases, affect how one acts and behaves (Bratman 1999, pp. 27-28).[62]

Now, it is clear that the mental states picked out by Sue's "suppositions" are not beliefs. First, Sue need not consider the target propositions to be true, likely, or the case. Suppositions permit one to be agnostic toward the truth of the target propositions. Second, Sue's endorsements of these propositions need not be classified as *irrational* if they are not appropriately based on epistemic considerations. It follows, then, that her attitudes are not beliefs.

On the other hand, it is not obvious that the attitudes that Sue forms in the *Vacation* case are different from attitudes of acceptance. Sue does not believe that she will travel to Japan. Nor does she believe that she will travel to Indonesia. Nevertheless, she decides to employ these propositions as

---

[61] I take it that the order in which the agent, Sue, forms her suppositions does not matter.

[62] Bratman, however, claims that suppositions do not "directly shape [one's] action" in the same way that acceptances do: acceptance of a proposition in a context is "tied more directly to practical reasoning than" supposition, and typically influences how one acts in that context (Bratman 1999, p. 28). The claim that suppositions have less influence over action might appear plausible. However, it seems that there are many cases of acceptance that are only tied to theoretical reasoning and not practical decision-making. For my purpose, I will remain agnostic about this feature.

premises in her reasoning. Indeed, she mentally fixes and implements these propositions in her thinking. It is therefore tempting to characterize Sue as simply *accepting* these propositions.

However, my contention is that it is correct to say that Sue's cognitive endorsements are *not* genuine attitudes of acceptance. Acceptance involves a rational resoluteness to endorse and implement a proposition for one's available reasons in a given context. Sue would lack this mental determination. Her mental positing of these propositions is *tentative* and *exploratory*. Her attitudes fall under a different class of cognitive endorsement, namely, the class of *supposition*. Following Cohen, we can maintain that supposing *p* (or assuming *p*) is usually "an inherently makeshift manoeuvre" or "an inherently temporary act of imagination" (Cohen 1992, pp. 12-13).[63] Indeed, suppositions involve a weak and tentative resoluteness to posit a proposition in short-term thought or reasoning.[64]

We thus have good reason to think that there is a notion of *supposition* that can pick out a distinct kind of mental state that differs from both belief and acceptance. To suppose *p* involves tentatively (or weakly) endorsing *p* in one's thought. On this conception, suppositions are closely related to acceptance. There are, however, several features of suppositions that distinguishes this attitudes from acceptance and other cognitive endorsements. Let me explain this in more detail.

Recall that accepting *p* has a distinct direction of fit. To accept *p* usually involves a positive stance toward *p* in a given context, *C*. One can either think that *p* is plausible, credible, or useful for one in *C*. Moreover, acceptance involves a mental resoluteness to implement *p* in one's long-term thought or reasoning. Suppositions, however, lack this resoluteness and direction of fit.

---

[63] On Cohen's view, suppositions need not involve any "commitment to a pattern, system, or policy – whether long or short term – of premising" a proposition in one's reasoning or "as a basis for a decision" (Cohen 1992, p. 12). My view of acceptance is different. However, I share the view that supposition need not involve any commitment or determination to *implement* a proposition in thought (including a willingness to continue fixing that proposition in thought while being willing to be receptive to new information and evidence concerning that proposition).

[64] Indeed, this also applies to cases involving suppositional reasoning and *reductio* arguments. One might, for example, articulate a *reductio* argument where one first supposes a proposition one does not believe in order to derive a contradiction.

First, suppositions do not require their possessors to be disposed to take a "positive" mental stance toward their target propositions. Unlike belief, one can suppose a proposition, *p*, while simultaneously possessing an attitude of agnosticism or even doubt toward the truth of *p*. One might explicitly form a supposition that conflicts with one's current beliefs. According to Nicholas Rescher, we can refer to these suppositions as "belief-contravening" suppositions: a supposition is *belief-contravening* if the target proposition is "of unknown truth-status, known to be false, or believed to be false" (Rescher 1964, pp. 3-4).[65] This is not implausible. In the *Vacation* case above, Sue can suppose the relevant propositions while believing that their negations are true. She need not consider it to be true that she will travel to either Japan or Indonesia. Moreover, one might suppose *p* while regarding *p* as lacking any epistemic plausibility, instrumental value, or practical usefulness. Accordingly, supposing *p* functions as a weak and tentative cognitive endorsement, such that one only puts *p* forward momentarily to either think about *p* or what might follow from *p*. This highlights the main features of suppositions that distinguishes these states from acceptance. *Supposing p need not involve any disposition to take a positive mental stance toward p, nor need it involve a determination or commitment to implement p in thought or reasoning within a given context or range of contexts*. Suppositions have the following feature:

> **(S1)** For an agent, *S*, to suppose a proposition, *p*, involves a weak and tentative positing of *p* as being momentarily *given* in thought or reasoning, such that (1) *S* need not commit herself to implement *p* in her planning, thinking, and reasoning, and (2) *S* need not commit herself to take a positive mental stance toward *p* – i.e., *S* need not regard *p* as being true, likely, possible, credible, useful, or practically beneficial for her decisions and plans.

At minimum, suppositions involve putting a proposition forward in thought. One need not be disposed to continue endorsing the target proposition within a given context or situation.

---

[65] We can extend Rescher's view to acceptance. A supposition that *p* can be acceptance-contravening: *S*'s supposition that *p* is *acceptance-contravening* at time, *t*, if *S* accepts the *negation* of *p* for some reason at *t*. If suppositions can conflict with beliefs and acceptances, then suppositions are distinct cognitive states. Suppositions need not coincide with belief or acceptance. It is possible for one to possess inconsistent attitudes of belief, acceptance, and supposition toward a single proposition, *p*.

Second, suppositions can be based on rational considerations. Like acceptance, suppositions can be rationally motivated and sensitive to both epistemic and practical reasons. One might have some evidential considerations for supposing *p*. But one can also suppose *p* for certain pragmatic considerations within a given context. For example, Sue might have practical considerations to plan a trip. She might therefore base her supposition that she will go to Japan on these practical reasons. Hence, an agent can explicitly form a supposition that *p* for either evidential reasons for *p*, pragmatic reasons for *p*, or a mixture of both evidential and pragmatic reasons.[66]

At this point it is important to note one unique feature of suppositions. Some suppositions need not be based on any available reasons. We can distinguish two types or species of suppositions: *rational* suppositions and *arational* suppositions. A *rational* supposition that *p* is based on, and sensitive to, rational considerations in support of *p*. These rational considerations can be evidential or pragmatic. On the other hand, an *arational* supposition that *p* is neither based on, nor sensitive to, rational considerations or support for *p*. In light of Cohen's remarks above, we can think of arational suppositions as temporary *imaginings* that one comes to possess. One simply postulates a proposition in thought for no obvious or determinate reason for doing so.

A corollary of this is that the warrant or justificatory support for suppositions, like the support for both belief and acceptance, can vary and fluctuate. One might have a *rational* supposition that *p*, such that this attitude is supported by either evidential or pragmatic reasons. Another agent's *rational* supposition that *p* might only be minimally supported by her available reasons. And at the far end of

[66] We must be careful here when demarcating the kind of cognitive endorsement associated with supposition. For example, if one supposes a proposition, *p*, that one knows or believes to be false, then this supposition will not be based on evidential reasons supporting the truth of *p*. As discussed earlier, a belief precludes being agnostic toward the truth of the target proposition. Hence, if I rationally believe not-*p*, and if I suppose *p* momentarily, my supposition will most likely not be based on evidential reasons supporting the truth of *p*. In this case I might only have pragmatic reasons for *p*. Even if a supposition that *p* is based on epistemic reasons, it nevertheless involves a *weak* form of endorsement: one tentatively puts *p* forward in thought by positing or weakly affirming *p*, without any conviction or disposition to think that *p* is true, plausible, credible, or useful.

the spectrum, one might have an *arational* supposition that *p* that lacks a rational basis. If this is correct, then suppositions have the following feature:

> **(S2)** Suppositions are sensitive and responsive to evidential and practical considerations. Rational suppositions can be based on either evidential or pragmatic reasons. Arational suppositions need not be based on any reasons.

From this it is clear that suppositions are rationally evaluated in a different way. One might suppose *p* in thought without any compelling reasons for doing so. Moreover, supposing *p* could be inconsistent with one's current, justified beliefs. Unlike other cognitive endorsements, one need not be *irrational* for holding this unjustified attitude of supposing *p*.

Another important aspect of supposition concerns how this attitude is formed. Some *arational* suppositions might be involuntarily formed without appeal to one's reasons. That is to say, one might simply find oneself supposing a proposition without a prior decision to do so. However, most *rational* suppositions are formed in a voluntary manner.[67] In paradigm cases of *explicitly* supposing *p*, one will consciously take it upon oneself to tentatively *adopt* or *posit p* in thought. Like acceptance, an agent can *will* herself to postulate *p* as being tentatively "fixed" in her thinking given her available reasons. The reasons need not be evidential. Indeed, this mental positing of *p* can be performed for purely pragmatic reasons without any consideration for the available evidence concerning the truth of *p*.

A related point is that, like acceptances, suppositions lack the temporal stability that belief typically enjoys. Suppositions can, according to Cohen, "be relative to a particular context" (Cohen 1992, pp. 13). This seems correct. Suppositions are not *necessarily* context-independent. I might, for example, not know whether my dog has enough food. But, in this context, I could suppose for pragmatic reasons that my dog lacks food. My supposition is constrained to this particular context. I need not indefinitely hold this attitude. This suggests that the following is true of suppositions:

[67] I grant that it could be possible to form *tacit* suppositions. Perhaps these attitudes are brought about at the sub-personal level. I will set this issue aside. I am concerned with conscious suppositions.

(S3) An agent's supposition that *p* can be the product of a willful and voluntary decision (or intention) to weakly and tentatively adopt, posit, or fix *p* in thought within a specific context, such that she need not have to take into consideration the available epistemic reasons for or against *p*.[68]

I have sketched the core conception of consciously formed suppositions. Features (S1) – (S3) specify the most salient characteristics of this type of cognitive endorsement. Suppositions differ from belief and acceptance. They are a unique kind of cognitive endorsement factoring into our thinking.

4. PRESUPPOSITION

Agents are sometimes described as having presuppositions. Presuppositions also factor into thinking, reasoning, and behavior. Accordingly, they can be classified as cognitive endorsements. I will briefly sketch an account of the nature of presuppositions. The account I develop here is primarily based on Stalnaker's work. However, there are some important divergences between our two views.

Let us first consider the following case:

Breakfast: Jack desires scrambled eggs for breakfast. He goes to the refrigerator and discovers that he is out of eggs. He forms the belief that if he wants scrambled eggs, then he should go to the store to buy more eggs. He then believes that he should go to the store to buy more eggs. Consequently, Jack decides to go to the store.

In the *Breakfast* case, Jack performs a line of reasoning terminating in a belief that he should go to the store to buy more eggs. His beliefs, along with his desire, play a role in his decision to buy more eggs at the store. Jack's thinking and behavior is therefore partly guided by his occurrent beliefs. However, it also seems that Jack's thinking and behavior is partly guided by what we might call "presuppositions". First, in the *Breakfast* case, Jack does not consciously form the distinct belief *that*

[68] I say that suppositions "can be" voluntarily brought about in most cases. I take it that (S3) is not a necessary feature of suppositions. Arational suppositions might occur involuntarily and without conscious appeal to any reasons for the target proposition. These suppositions are more akin to involuntary imaginings that simply happen to agents. The important point here is that a supposition that *p* is different from a belief that *p*, in that it can be voluntarily brought about by an act of the will, such that one need not seriously consider the available evidence for or against the target proposition.

the store has eggs. Moreover, he does not explicitly employ this proposition in his reasoning. Instead, Jack would seem to tacitly take this proposition for granted. We can plausibly say that Jack *presupposes* that the store has eggs. His occurrent beliefs and reasoning are partly based on this presupposition.

The question that concerns us now is: What are presuppositions? On one view that Stalnaker discusses, presuppositions are "beliefs ascribed [to a believer] even though the believer has never expressed or consciously thought about them" (Stalnaker 1987, p. 68). Stalnaker claims that an agent's presupposition that *p* will involve a "tacit" acceptance (i.e., belief) that *p*, such that this tacit cognitive endorsement can influence her behavior, even if she has never consciously considered *p* (Stalnaker 1987, p. 80).[69] If we connect this idea with the taxonomy of cognitive endorsements defended here, then we might conclude that presuppositions are appropriately identified as *dispositional beliefs* – i.e., unconscious beliefs fixed in the background of one's current system of beliefs (see section 1).

Stalnaker is working with his own conceptual framework. However, my account of cognitive endorsements provides a different picture of presuppositions. I admit that it is possible for one's acceptances to be tacitly held. An agent might consciously accept a proposition *p* at time t_1. Insofar as she does not revise her acceptance that *p*, it is plausible that she could retain this acceptance at some future time, t_2, even when she is no longer consciously considering *p*. We can plausibly take this acceptance at t_2 to be a tacit acceptance that *can* guide thinking and behavior. Tacit acceptances can therefore function like the presuppositions that we might identify as dispositional beliefs.

It follows that presuppositions are best understood as a class of *tacitly* held cognitive endorsements that can include dispositional beliefs and dispositional acceptances (and perhaps

[69] Note that Stalnaker takes *acceptance* to be a conceptual category containing other cognitive endorsements like beliefs, assumptions, and presuppositions (Stalnaker 1987, p. 79). On my view, acceptance is a distinct kind of cognitive endorsement that differs from both belief and supposition (or assumption).

dispositional suppositions).[70] These dispositional cognitive endorsements involve a tacit positing or fixing of a proposition, such that one tacitly takes that proposition as given. In some cases, these attitudes underpin explicit and conscious process of reasoning. In other cases, they can provide a partial basis for action. Hence, dispositional states of acceptance can influence thinking and behavior. The difference is that, unlike rational beliefs as presuppositions, our presuppositional acceptances need not be rationally based on good *epistemic* reasons. Nor do they dispose their possessors to consider their propositional contents as being true when (or if) these attitudes are occurrently held. Moreover, acceptances that function as presuppositions can be confined to certain contexts for practical considerations. At any rate, insofar as presuppositions are taken to be tacitly held cognitive endorsements that can affect thought and behavior, we have good reason to think that presuppositions are not restricted to dispositional beliefs. Presuppositions are a conceptual class of cognitive endorsements picking out dispositional attitudes of belief and acceptance that can influence thinking and behavior.

5. HYPOTHESIZING

It is typically held that, in ordinary cases of hypothesizing, an agent will consciously consider and put forward a proposition as either a possible explanation for some phenomena, or a possible solution to a problem (Harman 1986, p. 46). Hypothesizing can thus be characterized as a type of cognitive endorsement. Several questions can be raised. First, we can ask: What kind of mental *processes* constitute hypothesizing? A second question we might ask is: What kind of cognitive *states* factor into hypothesizing? These questions deserve answers if we are to have a proper grasp of the nature and

[70] Typically, suppositions are occurrent states. However, it might be possible for there to be cases of suppositional thinking based on other, unconscious suppositions. If such cases are possible, then tacit suppositions can function as presuppositions. I will not address this issue. My account is not held hostage to the truth of any view positing the existence of tacit suppositions. At most, I take presuppositions to be a class of tacit cognitive endorsements, namely, tacit beliefs and acceptances.

rationality of hypothesizing. I advance an account of hypothesizing that fits with my taxonomy of cognitive endorsements. I show that hypothesizing is a mental process involving inference and evidence gathering. I further argue that acceptance is the attitude underlying hypothesizing.

One view concerning the nature of hypothesizing can be taken from Gilbert Harman's writings on "tentative acceptance" (Harman 1986, pp. 46-48). According to Harman, "tentative acceptance" that *p* typically involves treating *p* "as a working hypothesis" (Harman 1986, p. 46). Harman claims that if one treats a proposition as a working hypothesis, then one will "see where one gets by accepting it, to see what further things such acceptance leads to" (Harman 1986, p. 46). For Harman, one's tentative acceptance that *p* as a working hypothesis will be "fruitful" if *p* "allows one to make sense of various phenomena; if it leads to solutions of problems, particularly when there are independent checks on these solutions; and if it leads naturally to other similarly fruitful hypotheses" (Harman 1986, p. 46). Furthermore, he asserts that a fruitful working hypothesis can eventually justify or warrant one in coming to "fully accept" or believe that hypothesis (Harman 1986, p. 46).

Harman's remarks suggest that we can understand hypothesizing in several ways. First, hypothesizing involves a mental process of employing a proposition that one weakly endorses to see what follows from that proposition. A weakly endorsed proposition lacks sufficient justification to warrant belief in that proposition. However, by employing these weakly supported propositions, one can come to mentally explore their consequences (logical or inductive). Moreover, one might even determine whether the target proposition is tenable depending on what further propositions it would entail if it were in fact true. Suppose, for example, that an agent, *S*, lacks justification for *p*, and is therefore agnostic toward the truth of *p*. Suppose that *S* employs *p* in her reasoning to determine what propositions are implied by *p*. *S* might come to recognize that some proposition, *q*, logically follows from *p*. But it could be the case that *S* correctly believes on good evidential grounds that not-*p*. *S* could then take this to count against the truth of *p*. In this way, an agent's hypothesizing functions as a form

of reasoning. An agent can employ a weakly endorsed proposition, *p*, as a premise in reasoning in order to determine the consequences of *p* or the tenability of *p*. Call this "inferential hypothesizing".

A second mode of hypothesizing involves evidence gathering or rational scrutiny. Hypothesizing with some proposition, *p*, might be constituted by a process of gathering evidence for *p*. Alternatively, it might involve gathering evidence against the truth of *p*. In these cases, one hypothesizes with a proposition, *p*, by reviewing and gathering the available evidence for or against *p*. One does not simply posit *p* as a premise and make inferences to determine which propositions follow from *p*. Rather, one is engaged in a process of rational scrutiny: one determines whether the available evidence supports the truth of a hypothesis. Call this "rationalized hypothesizing".

These modes of reasoning pick out genuine ways to hypothesize. Hence, hypothesizing is best construed as a broad mental activity of reasoning. In what follow, I will not focus on any particular process of hypothesizing. I admit that there can be other forms of reasoning that count as hypothesizing. The task at hand is to identify the cognitive endorsements factoring into hypothesizing.

Hypothesizing requires some kind of cognitive state. If an agent hypothesizes about the explanation or cause of some event, *E*, then she will put some possible explanation of *E* forward in thought. In this way, one mentally endorses this explanation. However, genuine hypothesizing is not compatible with every kind of cognitive endorsement. If an agent hypothesizes by postulating a proposition, *p*, as a possible explanation for some phenomenon, then that agent will not believe *p*. Following Harman's remarks above, hypothesizing typically involves propositions that are weakly justified. If *S* takes *p* to be a hypothesis that explains some event, then *S* will not *rationally* judge *p* to be true given her evidence. For to treat a proposition as a working hypothesis would entail that one is open to the possibility that *p* is false. And when one forms a belief, one judges that the target proposition is true (or is the case). One does not simultaneously doubt or hold an attitude of agnosticism toward a proposition that one rationally believes to be true. One would not hypothesize

about the explanation or cause of some state-of-affairs if one already possessed a belief about the relevant cause or explanation. Hypothesizing therefore requires some other cognitive endorsement.

The two most obvious candidates for the kind of cognitive endorsement involved in hypothesizing are acceptance and supposition. Some theorists restrict mental acts of hypothesizing to one of these attitudes. Harman, for instance, takes hypothesizing to involve acceptance (Harman 1986, p. 46). On the other hand, Cohen takes hypothesizing to be a state that is distinct from acceptance (Cohen 1992, p. 12). He claims that hypothesizing "always begins as a mere assumption or supposition" (Cohen 1992, p. 13). And if one's hypothesis, *h*, is sufficiently confirmed, then Cohen argues that one's supposition that *h* can lead one to eventually accepting *h* (Cohen 1992, p. 13).

Each of these theorists are working with their own views of *acceptance* and *supposition*. For my purposes, I will disregard these alternative views. Given my taxonomy of cognitive endorsements, we have the conceptual resources to articulate an intuitive picture of hypothesizing based on *acceptance*.[71]

To see why, consider a case involving a mental process of hypothesizing:

> ***Raccoons:*** Oscar believes that the neighborhood raccoon population has drastically increased. He has also witnessed raccoons tipping over his trash cans at night. One day, Oscar comes home to find that someone (or some thing) has ransacked his kitchen. Oscar considers the possibility that the neighborhood raccoons are to blame for the mess. Moreover, he considers this possibility to be a plausible explanation for the ransacked kitchen. He then collects evidence to see if the explanation is true.

In *Raccoons*, Oscar is hypothesizing. First, he mentally endorses a proposition as being a possible explanation for the occurrence of another state-of-affairs. In this case, we can stipulate that the relevant proposition is *that raccoons have ransacked the kitchen.* Call this proposition, "*p*". By considering *p* as a plausible explanation for the mess in his kitchen, Oscar treats *p* as a working hypothesis. Finally, Oscar examines and collects evidence to see if *p* is true.[72]

---

[71] It is possible that hypothesizing can sometimes employ both acceptance or supposition. However, I am inclined to think that paradigmatic cases of hypothesizing are restricted to attitudes of acceptance.
[72] Oscar's hypothesizing can include both inferential and rationalized modes of hypothesizing. Oscar can gather and scrutinize evidence for his working hypothesis, *p*. He can posit *p* in his reasoning and

The most plausible candidate for Oscar's cognitive endorsement of *p* is acceptance. Specifically, Oscar has an *epistemic acceptance*. First, he takes a positive mental stance toward *p* as a possible explanation for a state-of-affairs. Second, he mentally postulates *p* in such a way that he commits or dedicates himself to implement *p* in his thinking and reasoning. More precisely, he commits himself to use *p* as a basis for action, evidence gathering, inquiry, and rational scrutiny. Third, he basis his acceptance on his epistemic reasons. Fourth, his acceptance is formed with an epistemic goal in mind, namely, to perform inquiry. Finally, Oscar likely takes on a disposition to be rationally sensitive to new information concerning the truth or falsity of *p*. Hence, Oscar would appear to possess an epistemic acceptance toward the hypothesis, *p*. We can generalize this idea. Epistemic acceptance is therefore the cognitive endorsement that factors into mental practices of hypothesizing.

Now, one might object here and claim that Oscar could still hypothesize about *p* by merely supposing *p*. Let us "suppose", for *reductio*, that this is the case. Accordingly, Oscar might momentarily fix *p* in his thinking and conduct some line of reasoning to see if *p* is a good explanation for his ransacked kitchen. If Oscar's hypothesizing is based on a supposition, then he need not take any positive mental stance toward *p* – i.e., Oscar need not be disposed to regard *p* as being a plausible explanation for his ransacked kitchen.[73] Nor would Oscar be committed to fix and implement *p* in his reasoning and inquiry for the reasons he has in his specific situation. Finally, if Oscar's hypothesizing is based on his supposition that *p*, then he would not automatically be disposed to be receptive and sensitive to new evidence concerning *p*'s truth. If this is correct, then it is possible for Oscar to simultaneously believe that not-*p* is true. Moreover, it follows that he need not be sensitive to new evidence concerning *p*. But to genuinely hypothesize about *p* precludes having an attitude that *p* is not

---

inquiry to see what propositions follow, and to determine whether *p* is a true explanation for the ransacked kitchen.

[73] An advocate of this view would hold that suppositions factoring into hypothesizing must be *rational* suppositions. An *arational* supposition is not based on reasons. But hypothesizing is a rational affair. One hypothesizes about *p* when one has some reason to think that *p* explains some phenomenon.

true. Moreover, hypothesizing seems to require being receptive to evidential considerations for the target hypothesis. Indeed, one is committed to determine whether some hypothesis is an adequate or correct explanation of some phenomenon. Merely supposing *p* would not necessarily dispose one to be receptive and sensitive to rational considerations for or against *p*. It follows, then, that Oscar's hypothesizing does not involve an attitude of supposition. Hypothesizing involves attitudes of acceptance.

This does not suggest that one cannot reason about the explanations of things with suppositions. I can reason about the cause of some event, *E*, by supposing that *p* explains *E*. I need not be disposed to regard *p* as being a plausible explanation. Nor must I be willing to be sensitive to evidence concerning the explanatory connection between *p* and *E*. In fact, my supposition could conflict with my currently held belief that not-*p*. In this case, I am engaged in a line of reasoning about the explanatory relation between *p* and *E*. But it seems that I am merely reasoning *as if* I am hypothesizing about *p*. This, however, is not a genuine process of hypothesizing. Rather, I am taking something for granted. This is best characterized as a case of hypothetical reasoning or suppositional reasoning.

I have argued that hypothesizing requires an attitude of acceptance. An agent, *S*, genuinely hypothesizes about *p* when *S* accepts *p* and considers *p* to be a plausible explanation for some phenomenon, *E*. Accordingly, *S* will be disposed to take a positive mental stance toward *p* when *S* contemplates the plausibility of *p* as an explanation for *E*. Moreover, *S* will take on a commitment to use *p* in reasoning and inquiry about *E*. Finally, hypothesizing is accompanied by a willingness to be sensitive to rational considerations concerning the truth of *p*. If this is correct, then attitudes of acceptance play an important role in our theorizing, inquiry, and rational agency. Hypothesizing can lead to new beliefs. It can expand our knowledge. Thus, if hypothesizing paradigmatically involves

acceptance, then attitudes of acceptance can play a bigger role in shaping and extending an individual's knowledge.

6. Conclusion: The Rational Axis of Cognitive Endorsements

I have advanced a taxonomy of cognitive endorsements. Explicating this taxonomy is important for several reasons. First, it provides an intuitive and unified conceptual framework for understanding belief, acceptance, supposition, presupposition, and hypothesizing. Second, the taxonomy explicates how these cognitive endorsements manifest rational agency and contribute to our epistemic lives. These attitudes can be implemented in explicit cases of reasoning and inquiry. In some situations, an agent's cognitive endorsements function as tacit presuppositions guiding her thought and behavior. Moreover, acceptance can factor into processes of hypothesizing that can extend knowledge and understanding. Consequently, cognitive endorsements shape the way we view and consider the world. If we are to inquire into the nature and epistemology of cognitive endorsements, then we must have a suitable understanding of their features and how they rationally function in thought, inquiry, and reasoning.

In establishing this taxonomy, I argued that beliefs, acceptances, suppositions, and hypothesizing are distinct cognitive endorsements. What determines the type of cognitive endorsement an agent forms toward a proposition, *p*, crucially depends on her available reasons, aims, and the type of mental act she performs when considering *p*. I have shown that these states and processes differ with regards to three things: (1) the conceptual connection to the truth of their target propositions, (2) the type of reasons and aims that rationally motivate these states, and (3) the way in which one consciously brings about these attitudes.

The following tables specify the fundamental features of consciously formed cognitive endorsements.

Belief	
B1	For an agent, *S*, to believe *p* at some time, *t*, is for *S* to be disposed to regard *p* as being true, or being the case, or being more likely than not-*p* at *t*, in such a way that *S* cannot be agnostic toward the truth of *p* at *t*.
B2	An agent's rational belief that *p* is responsive and sensitive to evidence and epistemic considerations for or against the truth of *p*, regardless of nonevidential or pragmatic reasons that agent has in a specific context.
B3	An agent's consciously formed rational belief that *p* is typically the product of a mental act of judging that *p* is true, or that *p* is the case, or that *p* is more likely than not-*p*, such that she cannot voluntarily will herself to believe *p* by disregarding the evidence and epistemic reasons for the truth of *p*.

Table 1: The Fundamental Features of Belief

Acceptance	
A1	To accept *p* is to take on a commitment to implement *p* in thought and to be disposed to adopt some positive stance toward *p*, such that one can still be agnostic toward the truth or falsity of *p*.
A2	A rationally held attitude of accepting *p* is based on reasons. These reasons can be either epistemic or practical considerations in favor of *p*. A *epistemic acceptance* is primarily based on evidential reasons an agent has, and it is formed in pursuit of an epistemic goal. A *pragmatic acceptance* is primarily based on practical considerations an agent has, and it is formed in pursuit of some pragmatic goal.
A3	An acceptance that *p* can be the result of a voluntary choice or decision: an agent can will herself to accept *p* in such a way that she ignores the epistemic reasons for or against *p*, or she does not base her acceptance on good evidence.

Table 2: The Fundamental Features of Acceptance

Supposition	
S1	For an agent, *S*, to suppose a proposition, *p*, involves a weak and tentative fixing of *p* as being momentarily *given* in her thought or reasoning, such that (1) *S* need not commit herself to a long-term implementation of *p* in her planning, thinking, and reasoning, and (2) *S* need not commit herself or be disposed to take any positive mental stance toward *p* – i.e., *S* need not regard *p* as being true, likely, possible, credible, useful, or practically beneficial for her decisions and plans.
S2	Suppositions are sensitive and responsive to evidential and practical considerations. *Rational* suppositions can be based on either evidential or pragmatic reasons. *Arational* suppositions need not be based on any reasons.
S3	An agent's supposition that *p* can be the product of a willful and voluntary decision (or intention) to weakly and tentatively adopt, posit, or fix *p* in thought within a specific context, such that she need not have to take into consideration the available epistemic reasons for or against *p*.

Table 3: The Fundamental Features of Supposition

Presupposition	
P1	A *presupposition* is a broad conceptual category that picks out dispositional or tacitly held cognitive endorsements. This typically includes dispositional beliefs and dispositional acceptances. To presuppose *p* is to tacitly posit or fix *p* as given in one's unconscious thinking, such that this tacit endorsement can (in some cases) affect one's conscious thinking and overt behavior.
P2	A presupposition that *p* can be based on both epistemic or pragmatic reasons. Its conditions of rationality will depend on the specific kind of cognitive endorsement that is involved.
P3	A presupposition that *p* need not always trace its origins to some consciously formed attitude. In some cases, one might come to possess a presupposition tacitly and without willingly forming the underlying cognitive endorsement by considering reasons and evidence for the target proposition.

Table 4: The Fundamental Features of Presupposition

Hypothesizing	
H1	Hypothesizing is a mental activity that comes in multiple species. *Inferential hypothesizing* involves weakly endorsing a proposition, *p*, as a premise in reasoning with the aim of determining what follows from *p*. *Rationalized hypothesizing* involves evidence gathering and rational scrutiny directed at some weakly endorsed proposition, *p*.
H2	Hypothesizing about some proposition, *p*, requires an underlying attitude of *epistemically* accepting *p*. To hypothesize about *p* requires taking a positive mental stance toward *p* as a possible explanation for some state-of-affairs. It further requires some epistemic goal of inquiry or investigation. Hence, genuine hypothesizing precludes having a corresponding belief concerning the truth or falsity of *p*.
H3	Hypothesizing about a proposition, *p*, involves a commitment to use *p* in reasoning and inquiry. Hence, hypothesizing precludes tentative or momentary suppositions.
H4	Hypothesizing about a proposition, *p*, involves a resolution to be sensitive and receptive to evidence and rational considerations concerning the truth or plausibility of *p*.

Table 5: The Fundamental Features of Hypothesizing

The types of judgments, mental implementations, and postulating required to form these cognitive endorsements depend on the kind and quality of rational considerations for or against the target propositions. Furthermore, this shows that cognitive endorsements have distinct conditions of rationality. The figure below illuminates the rationality of our consciously held cognitive endorsements.

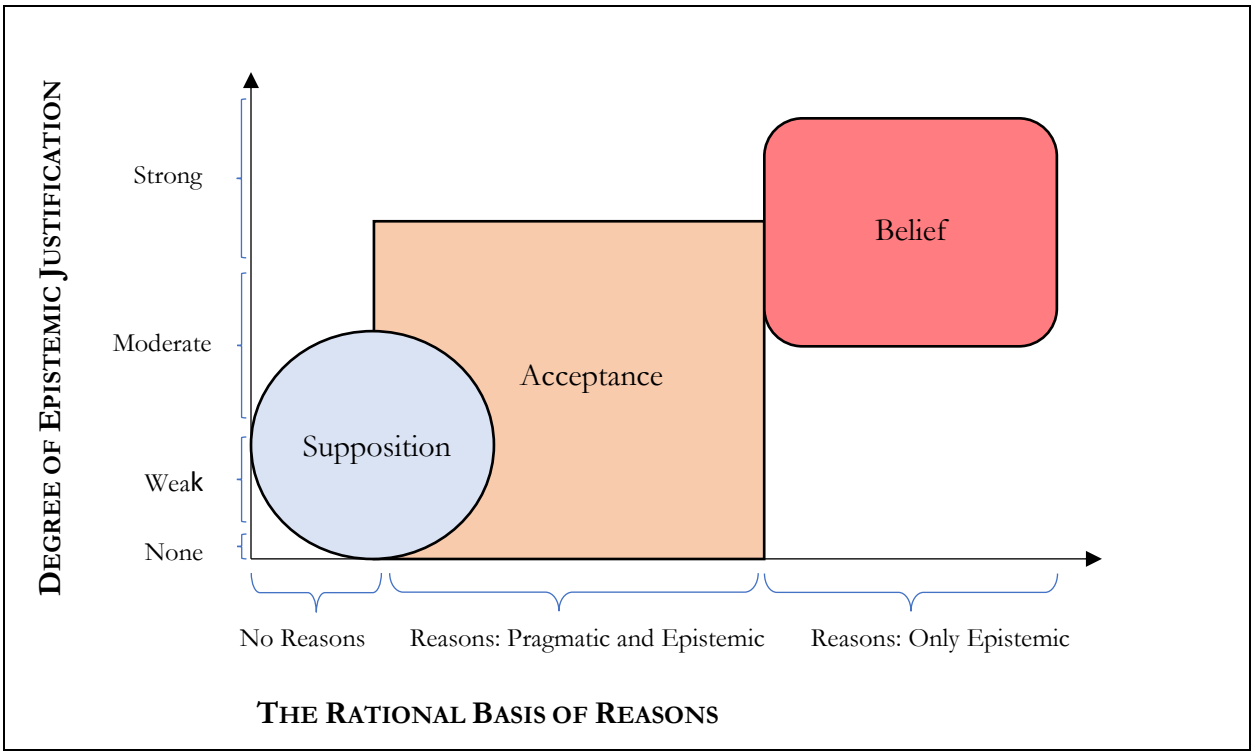

Figure 1: The Rational Axes of Cognitive Endorsements

Rational beliefs and acceptances require rational support. On the other hand, some suppositions need not be based on any rational considerations. While *rational* suppositions can be based on reasons, an *arational* supposition lacks rational support. Moreover, justification for rational supposition can significantly overlap with the kind of justification for acceptance. In some cases, acceptance can enjoy a high degree of evidential support or warrant that overlaps with the support one might have for a belief. In other cases, one's acceptance that *p* can factor into one's process of hypothesizing which can (in optimal situations) lead to a belief that *p*. What determines whether an agent rationally accepts, believes, or supposes *p* is dependent on the reasons available to her, the way she regards the truth, plausibility, or usefulness of *p*, and the method in which she comes to mentally endorse *p*.

II. COGNITIVE ENDORSEMENTS AS COMMITMENTS

Adam C. Sanders

0. INTRODUCTION

Cognitive endorsement is a broad conceptual category that includes attitudes and states of belief, acceptance, supposition, presupposition, and hypothesizing.[74] These states are cognitive endorsements precisely because they "fix" their propositional contents in thought as either being true, the case, plausible, credible, approximate, beneficial, useful, or tentatively "given". Conscious states of belief, acceptance, supposition, and hypothesizing significantly differ in their properties. Moreover, they differ in how they dispose their possessors to mentally treat and fix their target propositions. However, they all share a common feature: cognitive endorsements can affect the way their possessors behave and reason. If, for example, an agent believes *p*, then she is in a cognitive position to employ *p* in her reasoning, to testify that *p* is true, and act in accordance with her beliefs. Similar considerations apply to acceptance and supposition. Acceptance and supposition can factor into reasoning and, in some cases, decision-making and planning (see Bratman 1999). Hence, one might think that we can construe these cognitive endorsements in purely dispositional terms. According to W. V. Quine, a mental state as a *disposition* picks out a "physical trait" to act or behave in a certain way in appropriate situations (Quine 2008, pp. 322-325). He claims, for example, that to believe a proposition is to be disposed to "affirm" that proposition (Quine 2008, p. 321). If this is right, then perhaps other cognitive endorsements of acceptance, supposition, and hypothesizing should be understood as purely dispositional states that affect and direct behavior in certain circumstances.

[74] Recall that presuppositions typically include tacit cognitive endorsements of belief and acceptance. In some cases, a presupposition might not trace its origin to a consciously formed attitude. In this chapter I only discuss the cognitive commitments associated with consciously formed cognitive endorsements. If a presupposition was originally formed in a conscious way by either accepting *p* or believing *p*, then any commitments it has will be associated with the commitments of belief or acceptance. Hence, I will focus on the commitments of belief, acceptance, supposition, and hypothesizing.

One question we can ask is the following: Are cognitive endorsements to be understood in purely non-normative, dispositional terms? Let us call this the "normativity question". The primary aim of this chapter is to provide a satisfactory answer to the normativity question. I contend that if we think that our cognitive endorsements are inherently non-normative dispositions, then we are left with an unintuitive picture concerning the relationship we bear to our own mental states of belief, acceptance, supposition, and hypothesizing.[75] Instead, I argue that the correct answer to the normativity question is "no". I defend the claim that cognitive endorsements generate *cognitive commitments*. Cognitive commitments are normative commitments: they normatively prescribe how one ought to present and treat the propositional content of one's cognitive endorsements.[76]

This is not a novel position. Other philosophers similarly think that beliefs involve normative commitments (Levi 1997; Millar 2004; Bilgrami 2006; Shpall 2014; Coliva 2016, 2019; Liberman and Schroeder 2016). Most of these philosophers hold that an agent's beliefs engender commitments to believe the propositions that are implied by her beliefs (Levi 1997; Millar 2004: Bilgrami 2006; Shpall 2014; Liberman and Schroeder 2016). Cognitive commitments, however, are different.[77] They are not commitments to form new beliefs. Second, they extend to states of acceptance, supposition, and

[75] Alan Millar argues for a similar point (Millar 2004, ch. 4). I say more about his view later in the chapter.

[76] The notion of *cognitive commitment* is not to be confused with what Michael Lynch refers to as an "epistemic commitment" (Lynch 2013). Lynch defines *epistemic commitments* as "commitments to *epistemic principles*, which broadly speaking, concern what we ought to believe, or are justified in believing" (Lynch 2013, p. 344). These principles are external rules concerning how one ought to infer, revise, and form beliefs. On the present account, *cognitive commitments* are understood to be normative commitments that are constitutive elements of cognitive endorsements themselves. These commitments do not concern how one ought to reason and form new attitudes. Rather, they normatively constrain how one treats the target proposition of one's cognitive endorsement. More precisely, cognitive commitments prescribe (i) how one should regard the truth, plausibility, or usefulness of a target proposition, *p*; (ii) how one consciously puts *p* forward in thought and reasoning; and (iii) how one should be sensitive and responsive to different reasons for or against *p*.

[77] I will focus on the alternative view of doxastic commitments endorsed by Levi, Millar, Bilgrami, Shpall, and Liberman and Schroeder in chapter 3.

hypothesizing. Cognitive commitments are normative elements of the attitudes themselves that govern the way their target propositions are treated or regarded in thought and reasoning.

The account I defend here is developed out of the important work done by Margaret Gilbert and Annalisa Coliva. According to Gilbert, an agent's personal decisions and intentions generate normative commitments – i.e., commitments that normatively constrain or govern her behavior (Gilbert 2013; 2014). In a similar vein, Coliva maintains that when an agent forms a belief at time, *t*, her belief engenders a commitment to regard the target proposition as being true in such a way that she cannot simultaneously believe the conjunction, *p* & not-*p*, at *t* (Coliva 2019).[78] I develop this idea further. Cognitive endorsements generate a plurality of normative commitments. A consequence of this is that our cognitive endorsements are inherently normative states.

This account of cognitive commitments is beneficial for two reasons. First, it provides a plausible account of the normative structure of our thinking and reasoning. Moreover, it provides a satisfactory view of the kind of relationship agents have toward explicitly formed attitudes of belief, acceptance, supposition, and hypothesizing. Second, cognitive commitments help explain how we rationally evaluate and license criticism for our cognitive endorsements. Finally, this account provides a plausible basis for both the normativity of belief-formation and achievements of understanding. I discuss these further issues in chapters 3 through 6.

The chapter is structured in the following way. In section 1 I briefly discuss the worry with construing our cognitive endorsements in purely non-normative, dispositional terms. In section 2 I present Gilbert's account of *commitments of the will* which will serve as a background for the discussion on the normative commitments engendered by our cognitive endorsements. In section 3 I outline the normative commitments of belief. In section 4, I discuss the normative commitments of acceptance. In section 5 I articulate the commitments associated with suppositions. In section 6 I present the

[78] As I understand Coliva, this commitment cannot be violated. I will discuss her view in more detail.

normative commitments engendered by hypothesizing. Section 7 compares and contrasts cognitive commitments with commitments of the will. I conclude the chapter in section 8.

1. THE WORRY WITH NON-NORMATIVE ACCOUNTS OF COGNITIVE ENDORSEMENTS

Beliefs and other cognitive endorsements guide behavior. We employ these attitudes in theoretical and practical reasoning. We sometimes act in accordance with our cognitive endorsements. In this sense, cognitive endorsements have a *dispositional* element: they can causally factor into our behavior and produce certain psychological tendencies. But can we sufficiently analyze these cognitive endorsements in purely non-normative, dispositional terms? This is the normativity question.

Trying to conceive of cognitive endorsements without any normative element seems to be a mistake. Alan Millar, for instance, claims that our beliefs and intentions generate normative commitments (Millar 2004, chapters. 3-4). He argues that if an agent were to take her beliefs and intentions to be merely dispositional in nature, then she could simply "adopt a purely *contemplative* stance" towards these attitudes (Millar 2004, pp. 110-111). The so-called "contemplative stance" is, according to Millar, a stance in which one "simply register[s]" that one has a particular intention or belief, without thinking that one's intention or belief normatively constrains them in any way (Millar 2004, p. 111). But this seems incorrect. The contemplative stance misconstrues our relationship to our consciously formed attitudes and intentions (Millar 2004, p. 112). For instance, Millar correctly notes that if you intend to perform some action, Φ, then it seems that your intention normatively constrains you to carry out Φ (Millar 2004, p. 111). Beliefs also seem to generate normative commitments. Millar claims that our beliefs generate normative commitments to believe what follows from our beliefs (Millar 2004, pp. 73-74, 116-117).

I think Millar is right to worry about a purely dispositional and non-normative understanding of belief and intention. Indeed, I agree that our consciously formed beliefs and intentions engender

certain normative commitments. I will argue, however, that acceptance, supposition, and hypothesizing can also generate commitments. The commitments that I discuss here differ from the commitments that Millar identifies. My contention is that our cognitive endorsements generate a plurality of unique *cognitive commitments*. *Pace* Millar (and other theorists), these cognitive commitments do not normatively constrain their possessors to form further attitudes. Rather, they normatively constrain the way one treats the propositions that one cognitively endorses. Countenancing cognitive commitments, I contend, accurately portrays the normative relationship we bear to our attitudes and mental states of belief, acceptance, supposition, and hypothesizing. It follows that cognitive endorsements are inherently normative.

2. COMMITMENTS

Before sketching my account of cognitive commitments, I first need to discuss the normativity of commitments. Cognitive commitments, on my view, can be contrasted with the non-epistemic commitments that can be argued to issue from an agent's decisions and intentions (on the latter, see Gilbert 2014). In this section, I briefly discuss Gilbert's account of *personal commitments of the will* (Gilbert 2013; 2014). Taking her account into consideration provides a useful framework for developing an account of cognitive commitments. I then proceed to identify some salient features of commitments in general. Specifically, many theorists maintain that a commitment is a normative constraint placed on agents (Bilgrami 2006; Gilbert 2013; Coliva 2016). All things being equal, commitments provide standards of correctness and a rational basis for exercising criticism when one fails to live up to one's standing commitment (Gilbert 2014; Gilbert 2018; Coliva 2016). Moreover, if one possesses a commitment, then that commitment involves a disposition or readiness to be critical of oneself when one fails to live up to their commitment (Bilgrami 2006; Gilbert 2013; Coliva 2016). Cognitive commitments possess these features.

Typically, one's voluntary and intentional actions are causally preceded by corresponding intentions or decisions. Consider, for example, the following case. I might consider what meal I will make for dinner tonight. After checking the pantry and refrigerator, I come to the decision that I will make pasta. In a sense, my decision governs my future behavior: it prescribes a particular course of action, namely, the making of pasta for dinner tonight. Following Gilbert, we think of decisions and intentions as imposing normative constraints on us: one ought to do what one intends or decides (insofar as one does not change one's mind) (Gilbert 2013; 2014). Let me explain this in more detail.

Gilbert contends that a person's decisions and intentions engender what she calls, "personal commitments of the will" (Gilbert 2013; 2014). She claims that an agent, *S*, accrues a personal commitment of the will when *S* forms a *decision* or *intends* to achieve some goal or to perform some action (Gilbert 2013, pp. 899-905; 2014, p. 6). She claims that, if an agent *decides* to do some action, Φ, then her mental act is "an exercise of that person's will" that personally commits her to perform, execute, or carry out Φ (Gilbert 2014, p. 6). Similarly, individual intentions *qua* states of the will, according to Gilbert, produce personal commitments: If *S* intends to do Φ, then *S*'s intention personally commits *S* to execute or perform Φ (Gilbert 2013, pp. 899-905).[79]

[79] The relationship between individual intentions and decisions is a contentious matter. One might be tempted to say that a personal decision, like a personal intention, could be analyzed in terms of desires and beliefs. Robert Audi, for instance, gives an account of intentions, which can be stated (roughly) as follows: An agent, *S*, *intends* to do *x* if, and only if, (1) *S* has a belief that she will do *x*; (2) *S* desires to bring about *x*; and (3) *S* has no other "equally strong or stronger incompatible want" (Audi 1991, p. 362). Gilbert Harman characterizes "*simple decisions*" as complex states partly constructed by intentions: An agent, *S*, forms a *simple decision* to do *x* when (1) *S* intends to bring about *x*; (2) *S* is aware of some means, *m*, for bringing about *x*; and (3) *S* intends to do *m* (Harman 1983, p. 133). John Broome similarly claims that many decisions an agent makes can be understood as intentions that are formed through some line of reasoning (Broome 2002, p. 86). Moreover, personal decisions share many other features associated with intentions. Both aim at satisfying some goal, and both would seem to normatively constrain one's behavior. In his discussion of individual intentions, Michael Bratman claims that intentions provide a sort of "stability" for one's "plans over time" (Bratman 2009, p. 227). Personal decisions would also have this stabilizing feature. On the other hand, one might argue that personal decisions and intentions are distinct. Gilbert is inclined to hold this view. She takes individual decisions to be a kind of act, whereas intentions are states (see Gilbert 2013). Moreover, she claims that a personal decision has a "*trans-temporal reach* in terms of the persistence of the commitment" that

According to Gilbert, a personal commitment is a "normative matter": an agent's personal commitment to do Φ normatively constrains the behavior of that agent (Gilbert 2014, p. 6). For she claims that the commitment would rationally motivate the agent to act in a way that conforms with her commitment – otherwise she would be acting "in error if [she] fails to" satisfy her commitment (Gilbert 2014, p. 6). In Gilbert's terms, a personal commitment stemming from a decision or intention gives its possessor sufficient reason to carry out the content of that decision or intention (Gilbert 2014, p. 64). My decision to make pasta, for example, generates a normative commitment to make pasta. If there are no competing reasons that override my decision to make pasta, then I *ought* to make pasta given my prior decision.

Several features of this account are worth noting. First, Gilbert's notion of individual *commitments of the will* is a normative concept. The normative commitments that issue from an agent's decision or intention can be distinguished from the "psychological commitments" associated with mental states – i.e., the psychological motivations and causal dispositions to behave in certain ways (Gilbert 2013, p. 899).[80] Second, she claims that an agent's personal commitment of the will is "unilaterally" acquired by the agent when she forms a decision or intention (Gilbert 2013, p. 902). Third, an agent's personal commitment stemming from her personal *decision* is terminated prior to the action decided upon only if she consciously "rescinded [her] decision" (Gilbert 2013, p. 901). But a

---

outstrips the reach of those commitments generated by intentions (Gilbert 2013, p. 902). Hugh McCann similarly holds that decisions are not intentions (McCann 1986, p. 254-255). He claims that, "[d]ecisions are events rather than states" – that is, they are "acts of intention formation" (McCann 1986, p. 254-255). In what follows, I take decisions to be distinct mental acts that differ from intentions. However, whatever view one might have on this matter will not significantly impact my overall aim here. It is only important to note that intentions and decisions can plausibly be thought to generate normative commitments of the will.

[80] For example, Cheshire Calhoun, claims that "[c]ommitment is a species of intention" (Calhoun 2009, p. 615). Calhoun seems to characterize intentions as a psychological commitment or disposition to carry out the content of what one intends. Alan Millar makes a similar point about belief. He claims that an agent's belief that *p* "is a psychological commitment to using [*p*] as an assumption in [her] thinking" (Millar 2004, p. 122). In this chapter I will focus on *normative* commitments associated with cognitive endorsements.

personal commitment generated by an agent's *intention* can be terminated without that agent having to consciously rescind that intention (Gilbert 2013, p. 902). In particular, it can simply lapse. Fourth, she holds that a personal commitment (such as one stemming from a decision) can be stable over time: it will persist and set rational standards for one to behave in a way that aims at the fulfillment of the goal or content of the commitment until it is satisfied or rescinded (Gilbert 2013, pp. 902-903).

Another important element of commitments centers on the employment of criticism and correction when one fails to live up to one's commitment. Recall that, other things being equal, a commitment to Φ provides sufficient reason to Φ (following Gilbert 2014, p. 64). Gilbert correctly claims that an "un-rescinded" or "standing" commitment to Φ provides a rational basis or grounding for carrying out Φ: one "ought" to Φ (Gilbert 2014, p. 114). Failure to live up to a persisting commitment to Φ, she claims, is indicative of going against what rationality requires, and thus one becomes susceptible to criticism (Gilbert 2014, pp. 64, 114; Gilbert 2018, pp. 38, 162, 184). Following Gilbert, I take it that commitments provide a rational basis for criticism when one fails to live up to a standing commitment.

An interrelated point to note here is that a commitment typically involves a willingness or disposition to be critical of oneself when one fails to live up to a commitment. Gilbert claims that if an agent, *S*, formed a decision to do Φ, and *S* failed to perform Φ, then it would be natural or intelligible for *S* to reprimand herself for failing to live up to her commitment (Gilbert 2013, p. 901). A commitment typically involves (or goes along with) a disposition to be self-critical when one fails to live up to the relative normative standards. As Bilgrami puts the point:

> *To have a commitment, one must be prepared to have certain reactive attitudes, minimally to be self-critical or to be accepting of criticism from another, if one fails to live up to the commitment or if one lacks the disposition to do what it takes to live up to it*… (Bilgrami 2006, p. 226, italics original).

This seems correct. If I am committed to make pasta because I decided to do so, then my failure to make pasta without changing my mind indicates that something is wrong, and that I have acted in

error. For I have failed to live up to a standing course of action or plan that I set out for myself. In this case, I would be prepared to accept criticism (including self-criticism) for not satisfying the goal of my decision. Commitments, as I understand them, involve a preparedness to accept criticism when they are violated.[81]

The foregoing discussion provides a useful background for explicating the kinds of *cognitive* commitments essential to our cognitive endorsements. However, it would be a mistake to equate commitments of the will with cognitive commitments. In what follows, I demonstrate that cognitive commitments have a normativity that differs fundamentally from the normativity of commitments of the will. Consequently, we ought to think of commitments as being a wider class of normative constraints instantiated by a wider variety of mental states, events, and processes.

3. The Cognitive Commitments of Belief

Belief has several unique features. First, if an agent believes *p*, then she possesses a kind of mental representation in which she treats *p* as being true or the case. Second, belief is rationally sensitive to epistemic reasons in support of the target proposition. Finally, one cannot rationally believe a proposition, *p*, at will for pragmatic considerations within a context, while simultaneously ignoring one's evidence for or against *p*. However, beliefs also have an inherent normative element. In this section I defend the view that beliefs engender a cluster of normative commitments. To believe *p* is to be bound by normative constraints that govern the way that *p* is treated in thought and reasoning.

The view developed here is built on the account of doxastic commitment defended by Annalisa Coliva. According to Coliva, there are two species of belief: belief as a disposition and belief

[81] Coliva similarly claims that a belief that *p* "as a commitment" involves "dispositional elements" to employ *p* in reasoning "and accept criticism or be self-critical if it were shown that [*p*] is not supported by sufficient evidence" (Coliva 2016, pp. 32-33). I will say more about this in section 3.

as a *commitment* (Coliva 2016; 2019).[82] On her view, when one explicitly and consciously comes to believe *p* by judging that *p* is true, then one accrues a commitment to regard *p* as being true (Coliva 2019, p. 5). She contends that to have a belief as a commitment that *p* is to take on an epistemic commitment to represent the world in a particular way (Coliva 2019, pp. 5-8). She goes on to say that a belief as a commitment at some time, *t*, is "*constitutive* of epistemic rationality:" one cannot simultaneously consciously believe *p* and not-*p* at *t* (Coliva 2019, pp. 6-7). More precisely, if an agent, *S*, consciously judges that *p* is the case at time, *t*, then *S*'s judgment terminates in a belief as a commitment that *p*. This belief at *t* commits *S* to regard the target proposition, *p*, as obtaining. Accordingly, *S*'s doxastic commitment makes it impossible for *S* to believe both *p* and not-*p* at *t*.[83]

But to form a belief as a commitment that *p* is not to normatively bind oneself to indefinitely regard *p* as being true. Rather, one can, according to Coliva, escape one's doxastic commitment to treat a proposition, *p*, as being true by simply revising one's beliefs as commitments: (i) one can either judge that one lacks sufficient evidence for *p*, and thus suspend one's belief as commitment that *p*; or (ii) one can judge that there is sufficient evidence to support the truth of not-*p*, and thus form a new belief as a commitment that not-*p* (Coliva 2019, pp. 6-7).

We can thus formulate the following doxastic commitment that Coliva identifies:

---

[82] More precisely, Coliva holds that beliefs as commitments have the following five features: (BC1) They are the products of an agent's conscious "mental action of *judging* that some proposition obtains or is true"; (BC2) Generating the attitude is within the agent's "control:" the possession of sufficient evidence that either supports or counts against the target proposition will lead the agent to form the doxastic attitude or revise her current attitude; (BC3) One necessarily has first-personal self-knowledge of these attitudes; (BC4) They are the kinds of doxastic attitudes for which an agent can be "held directly rationally responsible for"; (BC5) They are constituted by commitments to regard their propositional contents as being true (Coliva 2019, p. 5). I take it that the account of belief I sketch in chapter 1 would be compatible with features (BC1) – (BC4) of Coliva's account. Moreover, beliefs as dispositions, as Coliva characterizes them, fits with standard accounts of dispositional beliefs as unconscious, background beliefs (see chapter 1).

[83] The idea that one cannot simultaneously judge and believe *p* and not-*p* at the same time has been endorsed by other philosophers. See, for example, Davidson (1985) p. 353.

> **Doxastic Veridicality:** If an agent, *S*, at time, *t*, believes *p* by judging that *p* obtains, then *S* is committed to regard *p* as being true at *t* (or obtaining at *t*), such that *S* cannot also consciously believe not-*p* at *t* without revising her initial belief that *p*.[84]

Doxastic Veridicality is closely related to what John Searle calls a "constitutive" rule in language (Searle 1970, pp. 33-34). For Searle, a *constitutive rule*, R_C, of some act, *F*, is understood to be a kind of rule that "constitutes (and also regulate[s])" *F*, where the "existence [of *F*] is logically dependent on" R_C (Searle 1970, pp. 33-34). These rules, he claims, differ fundamentally from what he refers to as "*regulative*" rules, which merely "regulate antecedently or independently existing forms of behavior" (Searle 1970, pp. 33-34). For example, an astrology club might adopt the following rule: *All members must wear their clothes inside out on Fridays*. This rule is a *regulative rule* placed on the members of the club. The rule normatively guides the behavior of the astrology club members. However, members can still choose to ignore the rule. In these cases, the members will have violated a norm. But the existence of the club and its members need not be logically dependent on the satisfaction of this rule. The rule is not a *constitutive* rule of the astrology club.

Following Searle's terminology, Doxastic Veridicality is a *constitutive*, normative rule engendered by the attitude of belief itself. This normative commitment is an essential element of a consciously held belief. To believe *p* necessarily requires regarding *p* to be true such that one cannot simultaneously think that not-*p* is true. Consequently, conscious belief is logically dependent on Doxastic Veridicality.

It is important to note that constitutive rules, according to Searle, need not be *genuinely* normative (Searle 1970, p. 41). That is to say, they can have a normative element that cannot be violated. Hence, the normativity associated with constitutive rules will differ from the normativity of

[84] This is similar to some of Duncan Pritchard's (2016) remarks. Pritchard notes that belief is conceptually tied to thinking that the proposition believed is true, such that one cannot be agnostic toward the truth of that proposition (Pritchard 2016, p. 90). Similarly, Coliva (2019) claims that Doxastic Veridicality constrains a believer in such a way that she is incapable of regarding both *p* and not-*p* as being true at the same time.

most regulative rules and norms. For in many cases, regulative rules can be ignored or violated. This idea can be made clear when we focus on beliefs as commitments. According to Coliva, the constitutive normativity of a belief as a commitment is not "purely" or genuinely normative in the same way that conventions and rules are purely normative: if an agent believes as a commitment that *p* at some time, *t*, then that agent cannot simultaneously believe both *p* and not-*p* at *t* (Coliva 2019, p. 7). This aspect of belief cannot be violated. It is a constitutive norm of belief that is not *purely* or *genuinely* normative. For an agent to possess a belief that *p* requires that the she lives up to the normative standard of regarding *p* as true such that she cannot simultaneously think that not-*p* is true. In order to believe not-*p*, the agent must first revise her initial judgment that *p* is true.

I think that Coliva's account is correct. Doxastic Veridicality is a cognitive commitment constitutive of belief.[85] At this point, I want to expand on this notion of cognitive commitment. I demonstrate that belief engenders a plurality or cluster of commitments. These commitments are constitutive normative elements of belief. More precisely, I argue that a consciously formed belief that *p* engenders two additional cognitive commitments. The first deals with belief's sensitivity to rational considerations. The second normatively binds one to continue believing in accordance with one's

[85] One might object and claim that, if Doxastic Veridicality is not *purely* or *genuinely* normative, then it is not a real normative constraint. According to Akeel Bilgrami, a commitment or rule is not normative if one is incapable of violating that commitment or rule (Bilgrami 2017, pp. 77-78). The possibility of violation, I take it, must at least be a logical possibility. But one might consider a rule to be normative if it is also physically possible to violate it. Hence, one could argue that, if an agent's belief that *p* at time, *t*, engenders a commitment that makes it impossible for her to consciously believe not-*p* at *t*, then this commitment is not normative. It is not normative because it cannot be violated. This, however, does not mean that Doxastic Veridicality lacks normativity. For this commitment prescribes how one ought to treat *p* if one believes *p*. If one believes p, then it is incorrect to be agnostic toward *p* or to also believe not-*p*. This is indicative of normativity. The commitment references correctness. Moreover, the mere fact that this commitment cannot be violated may not conclusively show that it lacks normative force. The possibility of violation might apply to some commitments like those stemming from decisions and intentions. But cognitive commitments are different. It is possible that some cognitive commitments cannot be violated. Hence, I take a more inclusive view concerning the normative failures of commitments: some cognitive commitments cannot be violated. Doxastic Veridicality is one of them.

prior judgment and evidence concerning the truth of *p*. These additional commitments are intimately connected with Doxastic Veridicality. All three commitments explain the normative relation we bear to our beliefs. Moreover, they partly explain some salient features and associated dispositions of belief. Let us discuss each commitment in turn.

The first cognitive commitment concerns the epistemic grounds of rational believing. To believe *p* is to regard *p* as being true or the case. As Bernard Williams puts the point, "beliefs aim at truth" (Williams 1973, p. 136). Following Richard Foley, I take it that the available evidence and epistemic considerations for a proposition, *p*, are an indication of the truth or falsity of *p* (Foley 1993, p. 16). Hence, if an agent, *S*, deems that she has sufficient evidence for the truth of *p*, then *S* can, all things being equal, come to judge that *p* is true on the basis of that evidence. Her resulting belief (in normal situations) is rationally grounded and sensitive to epistemic considerations.

Beliefs seem to generate a *psychological* commitment or disposition to be responsive to epistemic considerations. As I noted in section 2, psychological commitments are typically understood to be causal tendencies and dispositions to behave in certain ways. Upon reflection, consciously held beliefs typically disposes agents to form and revise their beliefs on evidential grounds. A rational agent does not typically come to believe a proposition, *p*, based on weak or absent evidence that she recognizes as being insufficient to support the truth of *p*. However, this tendency does not hold in all cases. Sometimes we do form beliefs on bad or insufficient evidence. Other times an agent might believe *p* and be unwilling to alter or revise her belief in light of strong countervailing evidence against the truth of *p*. In these cases, the beliefs in question would be irrational. The commonly shared intuition is that irrational beliefs are not correct. Indeed, as Coliva correctly points out, an agent's beliefs as commitments will dispose her to be self-critical when she recognizes that she fails to form her beliefs on sufficient evidence (Coliva 2016, p. 33). This suggests that there is an additional corresponding *normative* commitment constraining how we ought to form and maintain our beliefs. Beliefs ought to

be rationally formed and sensitive to epistemic considerations. One ought to believe *p* for good epistemic reasons. Moreover, one ought to alter or drop a belief given new evidence that supports belief-revision.

Furthermore, an agent with an explicit belief that *p* can normally be in a position to cite her reasons for believing *p*.[86] According to Coliva, to possess a belief as a commitment that *p* requires the capacity to provide those reasons and evidence that one has for the truth of *p* (Coliva 2019, p. 12). Citing one's reasons for a belief could be thought of as a disposition. But manifesting this disposition is typically due to a normative pressure stemming from the relevant belief. By giving reasons for why one holds a belief, one is indicating that one is living up to the normative demands of believing in accordance with having sufficient evidence. In some situations, however, one might lack one's original evidence for their current belief that *p*. Perhaps one forgot this evidence and yet continues to believe *p*. In these situations, one's belief could still be originally based on what one once considered to be sufficient evidence for the truth of *p*. Insofar as this agent is not aware of any evidence to the contrary, the agent need not revise her belief. She could, however, suspend her belief if her available evidence did not warrant the truth of either *p* or not-*p*.

This supports the following:

> **Doxastic Rationality:** If an agent, *S*, believes *p*, then *S*'s belief commits her to base that belief on sufficient epistemic grounds supporting the truth of *p*, and to be sensitive to new information and evidence concerning the truth or falsity of *p*.

Doxastic Rationality describes a normative commitment governing the attitude of belief. One satisfies this commitment by believing on good evidence and revising one's beliefs in light of new available evidence that defeats one's original evidence. One fails to live up to this commitment when one's

[86] The ability to cite one's reasons for believing *p* need not always be exercised. Nor is it the case that *all* agents are able to adequately or competently exercise this ability. However, our own experiences and intuitions suggest that, other things being equal, agents typically are able to cite some of their reasons for a belief they explicitly and consciously form.

belief is not rationally motivated or sensitive to evidence and epistemic considerations. In these cases, one believes in error.

Moreover, this commitment partly explains the normative relationship we have toward our beliefs. For we typically think that our beliefs ought to be initially grounded on what we take to be sufficient evidence. Moreover, our beliefs are normatively constrained to be sensitive to new, incoming evidence. Doxastic Rationality helps to explain the normative link between evidence and belief. And when one believes that *p*, Doxastic Rationality partly affords one with the capacity to identify her epistemic reasons in favor of *p*. If an agent's explicitly formed beliefs are normatively linked to some degree of justification or evidence, then that agent has (in normal cases) the resources to cite her rational basis for *p*. If belief was not normatively constrained in this way, then an agent could rationally believe propositions on what she takes to be weak or insufficient evidence without considering herself to be in error. But this diagnosis is wrong. It misconstrues (i) how we commonly think of rational belief, and (ii) how we should rationally treat or regard the propositions we come to fully believe.[87]

Doxastic Rationality is therefore a distinct cognitive commitment. The content of the commitment differs from the content of Doxastic Veridicality. Furthermore, one can fail to live up to Doxastic Rationality. That is to say, one can form and possess irrational beliefs. Hence, Doxastic Rationality is purely or genuinely normative. However, Doxastic Rationality is intimately connected to Doxastic Veridicality. To believe *p* is to be committed to regard *p* as being true. And one regards *p* as being true in a rational way when one judges that *p* is sufficiently supported by one's available evidence.

[87] It is possible for agents to base beliefs on insufficient rational grounds. An agent can recognize that *p* is not epistemically supported and still come to believe *p* due to a feeling or instinct she has. The agent might simply feel that *p* is true regardless of the fact that she recognizes that there is inadequate rational support for *p*. She might simultaneously wish that *p* is true or hope for *p*. In these cases, the agent's belief that *p* is irrational. She believes *p* while also thinking or believing that *p* is not well supported. A consciously formed belief that *p* engenders a commitment to believe in accordance with good epistemic reasons for *p*. When one's rational support for a belief is undermined or inadequate, then one's belief constrains him or her to revise the target belief. Belief involves Doxastic Rationality.

The final cognitive commitment I will discuss concerns the persistence of belief and its context-independency. In order to get clear on this kind of commitment, consider the following case:

> ***Linda:*** At time, t_1, Linda is aware of some evidence for the fact that there are rabbits in her garden. Even though she has never seen rabbits in her garden, she has observed signs of digging, small droppings around the flower beds, tufts of fur, and partially eaten leaves. She takes this evidence to sufficiently warrant the proposition, *that there are rabbits in my garden*. Call this proposition, "*p*". Linda thus judges at time, t_2, that *p* obtains which, *ipso facto* generates a belief that *p* at t_2. At some future time, t_3, Linda no longer is aware of her original evidence for the truth of *p*. Moreover, Linda is not consciously entertaining *p* at t_3.

It seems that in the *Linda* case, even though the agent at t_3 is no longer consciously aware of *p* or her evidence for *p*, she likely retains her belief that *p*. That is to say, her non-occurrent (dispositional) belief that *p* at t_3 would be counted amongst her background beliefs. Moreover, if Linda were asked at t_3 whether she believed *p*, then she would likely respond in the affirmative. Indeed, the posed question at some time after t_3 would likely cause her to recall her original evidence, which would then put her in a position to communicate her reasons for holding her belief. This suggests that, all things being equal, Linda would still be credited with a belief that *p*.

Given the *Linda* case, it seems that what partly explains the persistence of her belief that *p* is that her doxastic attitude generates commitments to continue to believe in accordance with her initial, corresponding judgment. As epistemic agents, we are exposed to an enormous amount of information, and we form an enormous number of attitudes based on that information. But a belief that *p* is usually context-independent: it persists throughout contexts and throughout changes in information, insofar as one does not judge that not-*p* (or that not-*p* and *p* are both likely to be true). Given these considerations, a belief that *p* seems to generate a disposition to continue believing in accordance with one's judgment that *p* is true. This helps explain the context-independency of belief (see chapter 1).

Now, one might believe *p* at time, t_1, and fail to continue believing *p* at some future time, t_2, without rationally revising her initial belief. In these cases, one might simply drop one's belief. However, we generally seem to think that, in normal situations, merely dropping one's belief without

any rational basis for doing so is inappropriate.[88] By rationally forming an explicit belief that *p*, one judges that *p* is true based on sufficient reasons for *p*'s truth. If one simply dropped this belief for no reason, then one comes to no longer judge or think that *p* is true. But this epistemic revision is rationally unsupported. Absent special circumstances, an agent with normal cognitive abilities would be in error if she simply abandoned her consciously formed belief that *p* without first making a rational judgment that *p* is no longer true, or that not-*p* is just as likely as *p*. This suggests that there is a corresponding *normative commitment*. Indeed, we think we *ought* to retain our beliefs if we do not revise them. It would be cognitively abnormal if one simply dropped one's beliefs without having good reason to do so.[89] Hence, belief seems to generate a normative commitment to continue believing what we judge to be true until we judge otherwise.[90] This particular normative commitment accurately describes the relationship we have toward our consciously formed beliefs.

Similar considerations apply to cases in which pragmatic considerations weigh against one's rationally held belief that *p*. If an agent believes, *p*, for good epistemic reasons, then she ought to continue believing *p*, even if she has strong pragmatic reasons within some context to *abandon* this

---

[88] There are cases where one simply forgets a belief. Moreover, one might undergo some cognitive stress and unknowingly abandon a belief. These cases might not always be indicative of irrationality or normative failure. However, we tend to think that these cases show that something has gone wrong. Hence, I take it that this suggests that our beliefs should persist unless we rationally revise them.

[89] This idea has also been endorsed by John Broome. According to Broome, beliefs seem to be governed by the following "*persistence*" requirement: "If t_1 is earlier than t_2, rationality requires of *N* that, if *N* believes at t_1 that *p*, and no cancelling event occurs between t_1 and t_2, and if *N* cares at t_2 whether *p*, then either *N* believes at t_2 *p*, or *N* considers at t_2 whether *p*" (Broome 2013, pp. 185-186). My account, however, is slightly different. Broome is concerned with requirements of rationality that govern belief states. I am concerned with normative commitments generated by our attitudes of belief.

[90] This cognitive commitment could persist in those cases where an agent cannot recall her original evidence for her belief that *p*. If an agent in this case was presented with the question as to whether she believed the target proposition, it is possible that she could truthfully assert that she did in fact believe that *p*. However, it is likely that in failing to remember her original evidence, the agent could come to realize that she lacks sufficient evidence for *p*. Such a realization could cause her to judge that not-*p* is just as likely to be true as *p* which, *ipso facto*, would cause her to suspend her belief that *p*. The normative pressure in this case would be credited to Doxastic Rationality.

belief. Beliefs are not rationally motivated by pragmatic reasons within a context. If one believes *p*, then one ought to continue believing *p* insofar as one has sufficient evidence for the truth of *p*.

Furthermore, this commitment helps explain the normative pressure to continue to hold one's belief that *p* even when one momentarily adopts inconsistent cognitive endorsements. To illustrate this last point, consider a modified version of the *Linda* case:

> ***Linda*:*** At time, t_1, Linda* is aware of some evidence for the fact that there are rabbits in her garden. She has observed signs of digging, small droppings around the flower beds, tufts of fur, and partially eaten leaves. She takes this evidence to sufficiently warrant the proposition, *that there are rabbits in my garden*. Call this proposition, "*p*". Linda* thus judges at time, t_2, that *p* obtains which, *ipso facto* generates a belief that *p* at t_2. At some future time, t_3, Linda* is no longer consciously entertaining *p* or her evidence for *p*. Rather, Linda* is engaged in a process of suppositional reasoning and is explicitly supposing that not-*p* holds at t_3.

It is reasonable to attribute the belief that *p* to Linda at t_3 even when she is consciously supposing or assuming that the negation of *p* obtains. Other things being equal, one typically maintains one's beliefs even when one takes on different kinds of mental attitudes with contents that are inconsistent with the target propositions of those beliefs. Indeed, we think that we ought to continue believing in these cases. For our beliefs are rationally motivated and revised on evidential grounds.

This supports the claim that our beliefs engender an additional, normative commitment: one is normatively committed to persistently believe in accordance with one's evidence and corresponding judgment regarding the truth of *p*. This further commitment is intimately connected to Doxastic Rationality. I contend, then, that belief engenders the following cognitive commitment:

> **Doxastic Integrity:** If an agent, *S*, forms a belief that *p* by judging that *p* is true, then *S* accrues a normative commitment to maintain *p* as being true in some doxastic form, such that, *S*'s belief can only be rationally revised or terminated via another judgment concerning the truth of *p* (e.g., a judgment that not-*p* obtains), or that, given the available evidence, not-*p* is just as likely as *p* to be true.

Countenancing this further commitment provides an explanation of the kind of normative pressures one faces when one forms a belief. It also partly explains some of the features of belief, especially its context-independency. An agent's rationally held belief will persist (either consciously or

unconsciously) over time and throughout different contexts until she rationally revises this belief. And this agent will be normatively bound to retain this doxastic state (either consciously or unconsciously) regardless of practical considerations within a context. Finally, if one is normatively constrained to retain one's belief, one can be in a position to employ that belief when one intends to conduct future inquiry bearing on the target proposition.

Moreover, this commitment, in conjunction with Doxastic Veridicality, explains the quality of belief to regard its target proposition as being true over time. For these two cognitive commitments work together to normatively dispose one to be non-agnostic toward *p* over time, and across contexts, insofar as one does not judge that not-*p* obtains (or that both *p* and not-*p* are equally likely to be true). Consequently, rational beliefs will typically persist in normal situations where one forms other, belief-like attitudes like acceptances or suppositions with contradictory or inconsistent contents.

Several clarifications are in order. First, Doxastic Integrity is purely normative. For it is possible to violate this commitment. One could, in some situations, be incapable of rationally maintaining one's beliefs. Hence, Doxastic Integrity is not understood to be a descriptive claim about the way our beliefs always function and persist. Rather, it characterizes a normative commitment that one acquires when one explicitly judges on rational grounds that a proposition, *p*, is true. One *ought* to maintain one's belief that *p* until one judges that *p* is either false or that not-*p* is just as likely to be true.

Second, while the commitment expressed in Doxastic Integrity normatively interconnects with Doxastic Veridicality, these two cognitive commitments are not equally foundational elements of belief. My contention is that Doxastic Integrity is derived from Doxastic Veridicality and Doxastic Rationality. According to Doxastic Veridicality, to judge that *p* is true will normatively bind one to regard *p* as being true. On the other hand, Doxastic Rationality commits one to believe in accordance with one's evidence. From this we can derive a commitment to continue representing *p* as being true insofar as one does not rationally revise one's initial belief-producing judgement.

Third, countenancing this kind of cognitive commitment does not necessarily mean that in order for an agent to form a new belief that *p*, she must first consciously judge that all of her prior beliefs with inconsistent propositions are either false or evidentially unsupported. As Coliva argues, unlike individual decisions that need to be consciously rescinded by the agent who makes them, we often do not consciously rescind our beliefs as commitments (or full beliefs) (Coliva 2019, p. 6). Rather, revising doxastic commitments is accomplished simply by performing new mental judgments. In this sense, she claims that one's doxastic attitude is "more like an intention:" to revise one's beliefs need not require one to consciously rescind any of her doxastic attitudes (Coliva 2019, p. 6). And, according to the additional cognitive commitment that I am advancing here, one's belief will persist according to one's *judgments concerning the truth of the target proposition*. I might, for example, judge that my dog has gotten into the garbage bins outside. Hence, I would be normatively constrained to retain this doxastic state over time in accordance with my judgment. But if, at some later time, I saw raccoons on my garbage bins, I might form the belief that raccoons got into my garbage bins. In this case, I am no longer committed to retain my prior belief concerning my dog – that belief no longer persists given a new judgment that indirectly affects the truth of that proposition.

Finally, Doxastic Integrity does not force us to adopt some form of epistemic conservativism. According to Christensen, *epistemic conservativism* is the view that simply having a belief that *p* provides some justification for one to persist in believing that *p* (Christensen 1994, p. 69).[91] The commitment of Doxastic Integrity is not meant to be understood as a kind of warrant for *p*. It is rather a normative commitment that one accrues when one forms a belief. An agent might be committed to continue

[91] Epistemic conservativism has been endorsed by several theorists. Harman seems to endorse a version of it when discussing what he calls the "*Principle of Conservatism*," which states that we are justified to maintain our beliefs "in the absence of a special reason not to" (Harman 1986, p. 46). Jonathan Adler seems to similarly defend a form of epistemic conservativism where holding a belief that *p* provides a type of *non-evidence-based reason* that warrants *p* (Adler 1996, p. 80). I remain agnostic toward epistemic conservativism. But see Christensen (1994) for some objections to this view.

believing *p*, and yet when she no longer is aware of sufficient evidence for *p*, then her commitment could cause her to continue to treat *p* as being true. But this is not to say that her commitment *justifies* her in holding this belief. If she notices her lack of evidence, and she judges that *p* is not supported, then she ought to relinquish her belief that *p* and the corresponding cognitive commitments. A cognitive commitment to continue believing *p* does not provide evidential support for *p*.

Doxastic Integrity, I submit, is a cognitive commitment constitutive of belief. It accounts for the normative fact that our beliefs ought to persistence and be stable. Moreover, it accurately describes the normative relationship agents have with their beliefs.

I articulated three commitments constitutive of belief. These commitments are normative: they govern how agents ought to mentally treat the target propositions of their beliefs. The argument for this claim proceeds in steps. We can first articulate the following transcendental argument.

> **(P1)** An agent's belief that *p* involves constitutive cognitive constraints: (1) she ought to regard *p* as being true, such that she cannot be agnostic toward *p* (or simultaneously think not-*p* is true); (2) she should base her belief on what she would take to be good epistemic reasons, and she should be sensitive to new evidence concerning the truth or falsity of *p*; and (3) she ought to continue believing *p* until she rationally revises her belief in light of good epistemic considerations to do so.

I argued that (P1) is true. Indeed, (P1) accurately describes our relationship to beliefs. From this we can posit the following premise:

> **(P2)** The cognitive constraints specified in (P1) can be true if and only if belief engenders a normative element.

If (P1) and (P2) are true, then we can conclude the following:

> **(C1)** Therefore, belief engenders a normative element.

Now, I have been insisting that the normative element generated by belief is a cluster of commitments. Why think that this is true? My contention is that commitments provide the best explanation for the normative force inherent in belief. Indeed, other normative constraints are not plausible candidates. The cognitive constraints and normative pressures associated with belief are not

the products of promises and agreements. Moreover, it seems false to think that we are bound by some moral obligation to treat the target propositions of our beliefs in a certain way. For instance, I am not morally obligated to believe *p* on good evidential grounds. My normative pressures associated with my beliefs are epistemic. The best explanation for these normative constraints is that they stem from commitments generated by the attitude of belief. Hence, we have an inference to the best explanation for the claim that belief engenders normative commitments.

The commitments of belief are similar to individual commitments of the will stemming from decisions and intentions. For both types of commitments are the products of mental phenomena unilaterally brought about by agents. However, the commitments associated with belief are *cognitive*: they normatively govern how agents treat propositions in thinking and reasoning. I will say more about the differences between cognitive commitments and commitments of the will toward the end of the chapter. I conclude that belief engenders the cognitive commitments of Doxastic Veridicality, Doxastic Rationality, and Doxastic Integrity. These are normative commitments constitutive of the attitude of belief. Hence, belief is inherently normative.

4. THE COGNITIVE COMMITMENTS OF ACCEPTANCE

In this section I argue that acceptance generates unique cognitive commitments. The commitments of acceptance are normative: they normatively constrain their possessors to treat the target propositions of attitudes of acceptances in a specific way. Moreover, these commitments are constitutive elements of the attitude of acceptance: they neither depend on, nor do they owe their existence to, any other preceding mental act, state, or judgment. The argument for this view is similar to the argument given in section 3. I first establish a normative element of attitudes of acceptance. I show that the normativity of acceptance is best understood to be a cluster of cognitive commitments.

My view of acceptance is as follows. To accept *p* involves fixing *p* in thought and taking on a mental resolution to implement *p* in thinking and reasoning. To fix *p* in thought is to posit and endorse *p* in one's thinking. This mental postulation is typically accompanied by a disposition to take a positive mental stance toward *p*: one either considers *p* to be plausible, credible, useful, or beneficial, when one contemplates *p*.[92] To implement *p* in thinking can be understood broadly. One can implement *p* by treating *p* as given, by introducing *p* into one's system of information or beliefs, by considering *p* as a possible and credible explanation for some phenomena, or by employing *p* as a premise in an explicit line of theoretical or practical reasoning.[93] It is also important to note that accepting *p* is rationally based on one's reasons (evidential or pragmatic) within a given context. First, adopting a positive mental stance toward a proposition, *p*, usually requires having some reason for doing so. Moreover, to take on a mental resolution of fixing and implementing *p* in some context will typically require some rational support that warrants this mental course of action.

Given this account of acceptance, it is no surprise that these attitudes will generate dispositions and causal tendencies to behave in certain ways. By accepting *p*, an agent might be disposed to assert *p*, to act as if *p* is true, to reason with *p*, and so on. There may be a multitude of these dispositions associated with a token attitude of acceptance. These dispositions can vary depending on how one implements a proposition in one's thinking and reasoning. However, these dispositions need not

[92] This disposition can be partly explained by the possession of certain beliefs having contents that are similar to the content of what is accepted. If, for example, an agent accepts *p* by implementing *p* in her thinking, and this agent takes a positive mental stance toward *p* by considering *p* to be plausible, then it is likely that she might also have a corresponding belief *that p is plausible*. This belief can partly give rise to her positive mental stance toward *p* when she contemplates *p* and her reasons for its acceptability. See chapter 1 for further discussion.

[93] As I argued in the previous chapter, this notion of acceptance is broader than Cohen's view (see Cohen 1992). For Cohen, accepting *p* involves taking on an explicit mental "policy" of employing *p* as a premise in a line of reasoning (Cohen 1992, p. 4). Taking on a mental policy of premising a proposition, *p*, is sufficient to accept *p*. However, it is not necessary. My view is more inclusive: there are multiple ways one can accept a proposition. To accept *p* requires a resolution to implement and manipulate *p* in one's thinking.

always be manifested when one accepts *p*. An absence of some disposition need not always result in a cognitive failure. Moreover, criticism is not always warranted when an agent fails to manifest one or more of these behavioral tendencies. Hence, this dispositional account does not imply any normative element associated with acceptance. I contend, however, that there is good reason to think that acceptance generates normative cognitive commitments.

If acceptance generates cognitive commitments, then these commitments must differ from the commitments stemming from belief. The commitments are different precisely because a rationally held attitude of acceptance has different properties. First, one need not regard *p* as being true if one merely accepts *p*. Hence, acceptance would not generate Doxastic Veridicality. Second, acceptance is rationally motivated by both epistemic and pragmatic considerations. Indeed, one need not be irrational for accepting *p* for practical reasons. Consequently, acceptance would not generate Doxastic Rationality. Finally, accepting *p* can be done at will and without consideration for the available epistemic reasons one has for or against *p*. In these cases, acceptance can be restricted to certain contexts given the practical pressures one has for accepting *p*.[94] If this correct, then acceptance lacks the stability that belief enjoys. Hence, acceptance would not engender Doxastic Integrity.

Nevertheless, acceptance involves a *normative* element. For it seems that one cannot appropriately accept a proposition without taking on some mental course of fixing *p* in thought for one's purposes or reasons. More precisely, if an agent, *S*, accepts *p* at some time, *t*, then *S* ought to adopt some mental resolution to implement and manipulate *p* in her thinking, reasoning, or inquiry within a given context or range of contexts where *S* has sufficient reasons warranting the acceptability of *p*. Otherwise, *S* would not have a correct acceptance. One would be in some state of cognitive error if they claimed to accept *p* and yet refused or failed to implement this proposition in thought. My contention is that acceptance engenders the following cognitive commitment:

[94] I discuss this in detail in chapter 1. See also Bratman (1999).

Acceptance Commitment: To accept *p* involves a normative commitment to take a positive stance toward *p*, and to implement and manipulate *p* in thinking or reasoning.

The above Acceptance Commitment is a normative constraint on our attitudes of acceptance. It prescribes how an agent ought to treat a given proposition that she accepts. One satisfies this normative commitment by implementing the target proposition in thought. For example, one might satisfy this cognitive commitment by employing *p* as premise for making inference and determining what one ought to do in a given context. In this case, one would live up to the normative standard generated by one's attitude of acceptance. Acceptance Commitment, I contend, is an essential normative constraint on attitudes of acceptance.

This normative constraint helps to distinguish acceptance from other attitudes and behavioral states. If, for example, an agent manifests some disposition to assert *p* and put *p* forward in her reasoning without explicitly committing to a mental course of action to implement *p* in her thinking, then she is merely acting as if *p* is true. And if she puts *p* forward in thought while also failing to take any positive mental stance toward *p* when she contemplates *p*, then she is best understood as simply entertaining or supposing this proposition. However, neither acting as if *p* is true, nor supposing *p*, is sufficient for accepting *p*. Acceptance, then, is governed by the Acceptance Commitment. This commitment accurately describes the normative constraint inherent in attitudes of acceptance.

Acceptance also generates an additional cognitive commitment. When one accepts a proposition, one typically does so for reasons. These reasons can be evidential or pragmatic considerations in support of the target proposition. Hence, rationally held acceptance is sensitive to epistemic or pragmatic considerations. For when one genuinely accepts *p*, one will be disposed to take a positive mental stance toward *p*. Moreover, one will commit to using and implementing *p* in thinking and reasoning. This presupposes that one ought to have reasons to take a positive mental stance toward *p*. For example, if one accepts *p* and considers *p* to be a plausible explanation for some state-of-affairs, then one should base this on the available reasons one has. Presumably, one ought to have

some epistemic reasons for why one takes *p* to explain the occurrence of that state-of-affairs. To lack a rational basis for accepting *p* implies that the corresponding positive mental stance toward *p* is not appropriate or rational. In this case, one is displaying a sort of cognitive error.

Similar considerations apply to the resolution to cognitively implement a proposition, *p*. Acceptance is a voluntary affair. To voluntarily commit oneself to a mental course of implementing and positing *p* in thinking is typically appropriate when it is based on reasons. If one lacks a rational basis for taking on this mental resolution, then one would not seem to be in a well-founded and appropriate state of accepting *p*. Acceptance is normatively connected to reasons.

Moreover, accepting *p* seems to constrain its possessor to be sensitive to new information and reasons concerning the acceptability of *p*. Consider a case where an agent initially accepts *p* for good reasons that she has. These reasons can be epistemic or pragmatic.[95] Let us further suppose that at some future time she comes to lose those reasons. At this future time, she recognizes that she has sufficient reasons not to endorse and implement *p* in her thinking. If, however, she still takes a positive mental stance toward *p* while implementing *p* in her thinking, then she does not seem to hold an appropriate attitude of acceptance. Indeed, her reasons do not warrant an acceptance of *p*. In this case, she is displaying a sort of cognitive error. She ought to correct herself by revising or dropping her acceptance.[96] Hence, rational acceptance that *p* normatively requires a disposition or willingness to be responsive to rational considerations for *p*.

Acceptance therefore generates the following normative commitment:

---

[95] Recall that acceptance is sensitive to a wide range of reasons. An epistemic acceptance primarily depends on epistemic reasons and epistemic goals, whereas a pragmatic acceptance will primarily be based on practical reasons and pragmatic goals. On this view, one could pragmatically accept things that one does not believe.

[96] The agent here could rationally drop her acceptance that *p* and still occasionally act as if *p* were true. She might assert *p* and employ *p* in her practical reasoning. However, acting as if *p* were true is not synonymous with accepting *p*. Accepting *p* also requires taking a positive mental stance toward *p* based on one's reasons in a context.

> **Acceptance Rationality:** To rationally accept *p* involves a normative commitment to base one's acceptance on reasons (either evidential or pragmatic) that one has in a given context, and to be sensitive to new information and rational considerations for one's acceptance.[97]

Acceptance Rationality is a constitutive commitment of attitudes of acceptance. It normatively governs the way agents ought to form and revise attitudes of acceptance. To correctly accept *p* is to fix and implement *p* in one's thinking for one's given reasons (either evidential or pragmatic), and to be sensitive to new rational considerations for endorsing *p*. In normal situations, if an agent rationally accepts *p* and satisfies the corresponding commitments, then she is in a position to cite her reasons for her acceptance. Moreover, if she mistakenly bases her acceptance on the wrong kind of reasons, then she will typically be pressured to revise her acceptance. This pressure is a normative one. Acceptance Rationality explains the normative relationship associated with attitudes of acceptance.

Both Acceptance Commitment and Acceptance Rationality normatively govern the way one cognitively treats the content of a token acceptance in thought. However, these commitments also involve a normative permissiveness. If one accepts *p*, then one need not regard *p* as being true. This mental state is therefore compatible with an agnostic attitude toward the truth of the target proposition. Moreover, some theorists, like Cohen, have argued that, given the voluntary nature of acceptance, on could, in principle, accept inconsistent or contradictory propositions (Cohen 1992, p. 35). This seems right. One could not explicitly believe the conjunction *p* and not-*p*. However, one

[97] I have previously argued that we can draw a distinction between *epistemic* and *pragmatic* acceptances (see chapter 1). An agent's acceptance that *p* is an *epistemic acceptance* if and only if it is primarily based on epistemic reasons and formed in pursuit of some epistemic goal. On the other hand, an agent's acceptance that *p* is a *pragmatic acceptance* if and only if it is primarily based on practical considerations and formed in pursuit of some practical end. Hence, the content of the commitment of *Acceptance Rationality* will vary depending on the type of acceptance involved. If one epistemically accepts *p*, then one ought to be sensitive to new epistemic considerations for *p*. One would be in error if one epistemically accepted what one knew to be contradictory or epistemically unjustified. However, if one pragmatically accepts *p*, then one is normatively bound to be sensitive to pragmatic considerations for or against *p*. In these cases, it is possible for one to pragmatically accept things that contradict one's current beliefs and epistemic acceptances. See chapter 1 for a more details.

could accept this contradiction. Perhaps one would have some practical reason that makes it pragmatically rational to endorse this contradiction in thought or inquiry.[98]

It follows that the commitments of acceptance would normatively *permit* one to doubt the truth of the target proposition. Acceptance therefore engenders an element of *agnostic permissiveness*:

> **Agnostic Acceptance Permissiveness:** If an agent, *S*, accepts *p*, then *S*'s cognitive endorsement of *p* need not be consistent with *S*'s other beliefs. Nor is *S* committed to consider or regard *p* as being true.

Agnostic Acceptance Permissiveness, as I am construing it here, is not a cognitive commitment like Acceptance Commitment and Acceptance Rationality. For it does not normatively constrain one to treat a proposition in a specific way. However, it is still a normative element of acceptance, in that it *permits* one to accept things that one does not believe (or believes to be false).

Now, the commitments of acceptance are similar to the commitments engendered by decisions and intentions. First, they are commitments generated by mental states. Second, they are voluntarily brought about by agents in the same way that one consciously forms a decision or intention. However, like belief, the commitments associated with acceptance are not properly understood to be commitments of the will. For they are cognitive commitments: they normatively constrain the way one treats a proposition in thought.

This, however, does not mean that commitments of the will cannot accompany our attitudes of acceptance. Consider a case where an agent, *S*, accepts *p* by positing and implementing *p* in her

[98] To be clear, some cases of acceptance would preclude the possibility of accepting both *p* and not-*p*. If, for example, I accept *p* for epistemic reasons, and I do so to posit and implement *p* in my reasoning as a plausible explanation for some event, then my acceptance would prohibit me from rationally accepting not-*p*. In this case I have an *epistemic acceptance* that *p*, and thus my reasons would not permit a simultaneous epistemic acceptance that not-*p*. To knowingly accept inconsistent propositions would likely be done for purely pragmatic reasons. For example, consider a case where an agent is threatened with death if she does not take up a long-term mental endorsement and implementation of both *p* and not-*p* in her thinking. Perhaps she recognizes that she has sufficient *practical* reasons to use both *p* and not-*p* as premises in her reasoning. In this case, her acceptance of this contradiction is a *pragmatic* one. This view coheres with other accounts claiming that one can accept contradictions (see, for example, Cohen 1992). I say more about the different species of acceptance in chapter 1.

thinking. Suppose that *S* mentally implements *p* by premising *p* in a line of theoretical or practical reasoning. In this case, *S*'s employment of *p* as a premise in her reasoning would likely be the product of a preceding decision or intention to implement *p* in a specific way. In this case, we might plausibly understand *S* to be subject to a commitment of the will stemming from her decision or intention: *S* is normatively constrained to utilize *p* as a mental premise in her reasoning. Consequently, commitments of the will can accompany cognitive commitments of acceptance. One satisfies the cognitive commitments of an acceptance that *p* by implementing *p* in one's thought for good reasons that one has. One can then decide to mentally implement *p* in a specific way, which *ipso facto* generates a commitment of the will to carry out that specific mental implementation. Hence, these two kinds of commitments are compatible.

Acceptance, like belief, engenders normative commitments. The argument for this claim is similar to the argument provided in the previous section. First, it was shown that an attitude of accepting *p* will generate certain constraints on how one is to treat and implement *p* in thinking and reasoning. Failure to conform to these constraints implies some sort of cognitive error. Indeed, such failures would typically manifest a disposition to be self-critical and a disposition to correct how one endorses and treats the target proposition. If this is true, then our attitudes of acceptance must have some normative element. My contention is that the normative element is best explained by commitments. These commitments are constitutive elements of acceptance. For it is not correct to say that we accrue these normative constraints through promises, agreements, or moral obligations. Hence, acceptance generates unique cognitive commitments that normatively constrain our thinking and reasoning. These commitments accurately characterize our normative relationship with the propositions we accept. Acceptance is not simply a state that manifests certain causal dispositions to behave in certain ways. Rather, it is an inherently normative cognitive endorsement.

5. THE COGNITIVE COMMITMENTS OF SUPPOSITION

Like beliefs and acceptances, suppositions also generate cognitive commitments. The cognitive commitments of suppositions, however, are unique constraints with a weak and tentative normativity. In this section I briefly discuss these commitments.

Suppositions have distinct features that differentiate these attitudes from both beliefs and acceptances. To suppose a proposition, *p*, involves a weak and tentative cognitive endorsement of *p*. When supposing *p*, one momentarily puts *p* forward in thought as being tentatively "given". A supposition that *p* need not require its possessor to regard *p* as being true. Nor does one have to take a positive mental stance toward the content of what is supposed in thought. That is to say, one need not be disposed to think that *p* is either plausible, credible, useful, or beneficial. Indeed, an agent's suppositions can conflict with other beliefs she might possess. Moreover, suppositions can be based on either pragmatic or evidential reasons. Accordingly, one can bring about a supposition at will for pragmatic reasons while ignoring the evidential reasons for or against the target proposition. Suppositions based on reasons (evidential or pragmatic) are characterized as *rational* suppositions. On the other hand, *arational* suppositions have not rational basis. These suppositions are more akin to momentary imaginings that occur in thought. Typically, when one has an arational supposition that *p*, one simply finds oneself consciously contemplating *p* in thought.

Given the unique qualities of suppositions, it is no surprise that these cognitive endorsements would not engender the same cognitive commitments inherent in our beliefs. Indeed, suppositions lack Doxastic Veridicality, Doxastic Rationality, and Doxastic Integrity. For suppositions are not conceptually tied to the truth. Nor are they irrational if they lack any strong evidential basis. Finally, suppositions need not be stable and persistent cognitive endorsements. A supposition might be a momentary positing of a proposition in thought.

Moreover, suppositions do not share the central features associated with acceptance. To accept *p* involves a commitment to posit and implement *p* in thinking or reasoning. Supposing *p*, however, lacks this sort of commitment. Supposing a proposition can be a transitory state. One need not adopt a mental resolution to utilize and manipulate the target proposition of a supposition in one's thought. Consequently, suppositions do not generate the Acceptance Commitment.

Furthermore, suppositions can be arational and thus lack a rational basis. Consequently, it would not be correct to insist that suppositions generate some commitment of rationality. Moreover, suppositions can conflict with one's beliefs and acceptance. This sort of inconsistency is not a mark of error or irrationality. Hence, supposing *p* would not instantiate any normative commitment of consistency with other cognitive endorsements one might possess.

The foregoing suggests that, if there are any normative commitments associated with suppositions, then these commitments are fundamentally unique. My contention is that suppositions generate three *weak* cognitive commitments. The first commitment is what I refer to as "Suppositional Commitment". The second commitment involves a type of permissiveness shared by acceptance. One is normatively permitted to suppose a proposition that is inconsistent with one's attitudes of belief and acceptance. I refer to this commitment as "Agnostic Suppositional Permissiveness". The third commitment is called "Unjustified Permissiveness". This commitment normatively permits one to suppose propositions that one lacks justification for. I discuss each commitment in turn.

Suppositions involves a cognitive "fixing" of their propositional contents in thought. More precisely, to explicitly suppose *p* involves a transitory or temporary positing of *p* in thought, such that one consciously postulates and weakly affirms *p* in one's thinking. To *temporarily* and *weakly* affirm a proposition, *p*, is to put *p* forward in thought in such a way that the cognitive endorsement is highly susceptible to revision in the light of new evidence, practical considerations, contextual matters, and reasoning concerning the truth, plausibility, or usefulness of *p*. Nevertheless, this transitory positing

of *p* still involves a mental endorsement of *p*. It follows that an explicitly formed attitude of supposition will constrain its possessor to cognitively posit the content of what is supposed. Suppositions therefore engender the following constitutive commitment:

> **Suppositional Commitment:** An explicit attitude of supposition that *p* engenders a normative commitment to momentarily present *p* in thought in such a way that *p* is weakly affirmed, fixed, or put forward as a postulate in one's thinking or reasoning.

Suppositional Commitment is a cognitive commitment. For it concerns the treatment of propositions. Moreover, this commitment is normative: it prescribes how one ought to mentally treat a proposition that one supposes. An agent would not be in a genuine or appropriate state of supposing a proposition, *p*, if she did not explicitly endorse *p* in this way.[99]

The normativity of Suppositional Commitment differs from the normativity of Acceptance Commitment. When one accepts a proposition, one is committed to take up a mental resolution to continuously endorse and implement that proposition in one's thinking given one's rational considerations in a context. Hence, the commitment inherent in an attitude of acceptance is more stable than the commitment associated with a supposition. A useful comparison can be made with the commitments of decisions and intentions (see section 2). According to Gilbert, the commitments of personal intentions lack the "*trans-temporal reach*" that commitments of personal decisions enjoy (Gilbert 2013, p. 902). One can simply drop an intention without having to first consciously rescind

[99] Perhaps it is tempting to understand this commitment as not being purely normative in the same way that Doxastic Veridicality is not purely normative. That is to say, one cannot violate this commitment. One supposes *p* if and only if one weakly posits *p* in one's thought. However, we have some reason to doubt this claim. It seems possible for one to violate Suppositional Commitment. Consider, for example, a case where an agent, *S*, comes to suppose *p* at time, t_1. At this time *S* recognizes that *p* is unjustified. However, through some cognitive malfunction at a later time, t_2, *S* mistakenly thinks that *p* is true. Consequently, *S* irrationally comes to permanently fix *p* as being true in her thinking at t_2. Clearly there is something normatively wrong with *S*. I suspect that the normative error is partly due to *S*'s violation of Suppositional Commitment: *S* does not satisfy her commitment to tentatively and weakly affirm *p* in her thinking. Moreover, if *S* transitions into a state of believing *p* at t_2 when she also recognizes that she lacks justification for *p*, then she violates Doxastic Rationality. The mental transition here is indicative of two failures to satisfy two cognitive commitments.

that intention (Gilbert 2013, p. 902). The trans-temporal reach of Suppositional Commitments is similar to that of intentions. One may simply lose a supposition that *p* without having to consciously decide or intend to stop supposing *p*. Accepting *p*, however, involves a stable and persistent commitment to mentally utilize *p* based on one's available reasons within a context. Suppositional Commitment is a weak cognitive commitment.

The second commitment engendered by suppositions involves a normative permissiveness. One can have two contradictory suppositions. Moreover, suppositions can conflict with other cognitive endorsements. One can suppose *p* while simultaneously believing not-*p*. Following Rescher, suppositions can be "belief-contravening" (Rescher 1964, pp. 3-4). In these cases, I am not irrational or in error if my supposition directly conflicts with my beliefs. Moreover, suppositions need not be classified as irrational if they fail to be consistent with attitudes of acceptance. Consequently, suppositions generate the following commitment of permissibility:

> **Agnostic Suppositional Permissiveness:** If an agent, *S*, supposes *p*, then *S*'s supposition does not have to be consistent with *S*'s other beliefs and acceptances.

Finally, suppositions are sometimes classified as *arational*. These suppositions are not consciously formed on the basis of reasons. In many cases, one possesses an arational supposition that *p* by simply finding oneself positing and contemplating *p* in thought. In these cases, one does not consciously will oneself to endorse *p* for certain reasons that one has. It is important to note, however, that there is nothing inherently wrong with possessing an arational supposition. It therefore follows that suppositions engender the following normative permissiveness:

> **Unjustified Permissiveness:** If an agent, *S*, supposes *p*, then *S*'s supposition permits her to weakly and tentatively endorse *p* in thought without having any rational basis (evidential or pragmatic) for the truth, plausibility, or practical usefulness of *p*.

Agnostic Suppositional Permissiveness and Unjustified Permissiveness are normative commitments. Both commitments prescribe how one *ought to* cognitively treat the contents of one's suppositions. An agent need not be self-critical if her suppositions conflict with her attitudes of belief and acceptance.

Nor would she be self-critical if she momentarily supposed *p* when she recognizes that she has no reason to do so. One is permitted to suppose propositions that one thinks are false and unjustified.

One final point is in order. Our cognitive commitments of suppositions can be accompanied by commitments of the will. This is most conspicuous in cases of explicit suppositional reasoning. One might suppose *p* and come to accrue a commitment of the will when one forms a personal decision or intention to momentarily employ this supposition in a line of reasoning. More precisely, an agent, *S*, acquires a commitment of the will to undergo suppositional reasoning at some time, *t*, when *S* (i) fails to believe or accept *p* at *t*, (ii) supposes *p* momentarily, and (iii) forms a personal intention or decision at *t* to employ her supposition that *p* as a premise in her theoretical or practical reasoning. Accordingly, when one voluntarily decides to undergo suppositional reasoning, her reasoning process is normatively constrained by several kinds of commitments: cognitive commitments associated with the attitude of supposition, and commitments of the will stemming from the intention or decision. Both kinds of commitments are normatively compatible.[100]

I have articulated an account of cognitive commitments that are constitutive elements of the attitude of supposition. My contention is that we have good reason to think that this account is true. The argument for this claim is similar to the argument given for the cognitive commitments of belief and acceptance. First, I have argued that an attitude of supposition will impose certain constraints on its possessor. These constraints are cognitive, in that they constrain how one *ought* to treat the target proposition of a supposition. Some of these constraints are permissive. For example, one is permitted to suppose unjustified propositions. Moreover, one can suppose a proposition that conflicts with one's beliefs and acceptances. In these cases, one would not be in error. These remarks suggest that that

[100] It is possible to undergo suppositional reasoning that is not initiated by a prior decision or intention. One might find oneself reasoning with a supposition. In this case, one is not under any normative pressure to continue one's suppositional reasoning if they did not decide or intend to do so.

suppositions have an inherent normative element. The best explanation for this normativity is that our attitudes of supposition generate normative cognitive commitments.

6. The Commitments of Hypothesizing

To hypothesize about some state-of-affairs or event, *E*, typically involves accepting a proposition, *p*, as a possible and plausible explanation for *E*. Hypothesizing involves a mental resolution to implement a working hypothesis in one's thinking, reasoning, and inquiry. If this is correct, then hypothesizing essentially involves cognitive commitments associated with the attitude of acceptance. There is, however, reason to think that the normativity of hypothesizing is not exhausted by the *cognitive commitments* of acceptance. In this section I show that hypothesizing also generates *commitments of the will*.[101] These commitments are constitutive elements of the mental *processes* of hypothesizing.

For an agent to competently hypothesize about some hypothesis, *p*, she must be open to the possibility that *p* could be false. Indeed, one would not really be engaged in a process of hypothesizing if one was not open to this sort of possibility. Hence, if an agent sincerely hypothesizes with a working hypothesis, *p*, then she ought to keep an open mind as to the plausibility of *p*. Moreover, she should be receptive and sensitive to new information that may or may not support the truth of *p*.

This suggests that genuine hypothesizing involves some additional mental commitments. Specifically, hypothesizing involves certain decisions and intentions concerning how one utilizes a hypothesis in reasoning and inquiry. Let me explain.

First, when an agent hypothesizes, she typically decides or intends to employ a hypothesis in her inquiry, reasoning, and testing. For instance, she might explicitly put a hypothesis, *p*, forward as a premise in her reasoning to see what propositions follow from *p*, and to see how other propositions rationally support *p*. This mental activity would stem from a personal decision or intention to reason

[101] For a more detailed discussion of commitments of the will, see Gilbert (2013; 2014).

with *p* and rationally scrutinize *p*. If this is the case, then the corresponding decision or intention generates a commitment of the will. This commitment can be expressed as follows:

> **Commitment of Hypothesizing:** If an agent, *S*, hypothesizes with a proposition, *p*, such that *S* decides or intends to employ *p* in her reasoning, inquiry, or testing, then *S*'s decision or intention normatively commits her to utilize *p* as a premise in her reasoning and inquiry, or as a postulate for rational scrutiny and evidence gathering.

This commitment normatively constrains how one ought to treat a working hypothesis. For if an agent decides to treat *p* as a working hypothesis, then she would commit herself to a certain course of mental action of implementing *p* in her inquiry. Other things being equal, the agent has sufficient reason to satisfy this commitment. That is to say, she ought to employ *p* in her reasoning and inquiry. Failure to comply with this commitment while still holding an intention or decision to hypothesize with *p* results in a normative failure. Such failures license criticism and normative pressure to correct oneself.

Moreover, hypothesizing involves a decision or intention to be open-minded about the truth or plausibility of a working hypothesis. For one would not appropriately hypothesize if one was not inclined to be sensitive to new information and evidence concerning the truth or falsity of the target hypothesis. If this is correct, then hypothesizing generate the following commitment of the will:

> **Non-dogmatism:** If an agent, *S*, engages in hypothesizing about a proposition, *p*, such that *S* treats *p* as a working hypothesis, then *S* is normatively committed to being open-minded about *p*'s truth or falsity, and *S* commits herself to being receptive and responsive to new information that can establish *p*'s truth-status.

Non-dogmatism is a commitment of the will generated by an intention or decision to be open-minded in cases of hypothesizing. This commitment normatively constrains how one reasons, and how one considers and treats evidence for or against a proposition being treated as a working hypothesis.

Hypothesizing is a normative mental process engendering a plurality of commitments. For hypothesizing involves an attitude of acceptance generating cognitive commitments. Moreover, hypothesizing generates commitments of the will. These commitments constrain how one ought to

employ a working hypothesis in reasoning, and how one ought to treat new evidence for or against that hypothesis. To hypothesize correctly requires one to live up to these normative commitments.

7. Cognitive Commitments and Commitments of the Will

Cognitive commitments are similar to commitments of the will. Both kinds of commitments place normative constraints on their possessors. Moreover, both commitments are unilaterally brought about and revised by the agents themselves through changes in their own mental states. Nevertheless, these two types of commitment differ in important ways. This section provides a more detailed discussion of the fundamental differences between cognitive commitments and commitments of the will. What follows is a careful demarcation of these two types of normative constraints.

The first fundamental difference between cognitive commitments and commitments of the will concerns their so-called "direction of fit". Following Coliva, beliefs are properly understood to be representational states that can be correct if, and only if, they represent their contents correctly (Coliva 2019, p. 8). Similar considerations apply to acceptance and supposition.[102] Hence, we understand cognitive endorsements as having what Williamson and Searle have called a "mind-to-world direction of fit" (Williamson 2000, p. 2; Searle 2001, p. 37).[103] That is to say, cognitive endorsements fix their target propositions in thought as either being true, plausible, approximate, given, and so on. The corresponding cognitive *commitments* also have a mind-to-world *normative* fit: by cognitively endorsing a proposition, *p*, one is committed to mentally postulate and portray *p* as either being true, plausible, approximate, or given, depending on the type of attitude involved. One is normatively compelled to

[102] Of course, the direction of fit is not the same. One need not think that *p* is true if one accepts *p*. Suppositions are also compatible with agnosticism toward their target propositions.
[103] I also discuss *direction of fit* in chapter 1. It is also important to note that Williamson and Searle are primarily focused on the "mind-to-world direction of fit" associated with belief (Williamson 2000, p. 2; Searle 2001, p. 37). The direction of fit associated with acceptances and suppositions will differ slightly. Nevertheless, given that acceptances and suppositions are cognitive endorsements, they still have a mind-to-world direction of fit, rather than a world-to-mind direction of fit.

mentally fix, endorse, and implement a proposition in a specific way. Hence, one's cognitive commitment involves a mind-to-world normative fit.

Consider, for example, a case where I believe that aliens abducted me. Accordingly, my belief is a representation of a particular state-of-affairs. In this case, it is a misrepresentation of how things are. My belief is false. Nevertheless, my belief commits me to regard its content as being true. Acceptance and suppositions generate similar commitments. If I accept a proposition, *p*, then I am committed to fix and implement *p* in my thinking for my given reasons in a context. If, however, I merely suppose *p*, then I am committed to tentatively and weakly fix *p* as given in my thinking or reasoning. Cognitive endorsements settle how one cognitively treats a proposition in one's thinking and reasoning. The way in which I fix a proposition in thought will vary depending on the type of cognitive endorsement involved. However, all cognitive endorsements generate cognitive commitments with a mind-to-world direction of fit.

Personal intentions and decisions have a different direction of fit. In Searle's terms, we can say they have a "world-to-mind direction of fit" (Searle 2001, p. 38). Following Searle, decisions and intentions are not characterized as true or false (Searle 2001, p. 38). Unlike cognitive endorsements, a decision or intention need not aim at representing how things actually stand in the world. Rather, a decision or intention to do Φ simply involves settling matters concerning what one ought to do, what course of action one should take, or what kind of event one should bring about. If, for example, I decide to climb Mount Everest, then my decision does not represent a certain state-of-affairs as obtaining. As Searle notes, a decision or intention aims at bringing about a state-of-affairs to conform to the content of that decision or intention (Searle 2001, p. 38). Hence, intentions and decisions have a world-to-mind direction of fit. Moreover, the corresponding commitments of decisions and intentions will also have a unique *normative* fit. For personal commitments of the will normatively

constrain their possessors to perform some further action(s) or bring about some state-of-affairs. Unlike cognitive commitments, personal commitments of the will have a *world-to-mind normative fit*.

The second major difference between these commitments involves their normative scope and conditions of satisfaction. When an agent decides or intends to do Φ, then she has a commitment of the will to bring about Φ. Other things being equal, she ought to carry out her decision to do Φ (following Gilbert 2014, p. 64). It follows from this that our commitments of the will have a wide normative scope that ranges over a wide class of actions and states-of-affairs. One's commitment of the will can constrain one to perform physical, bodily actions. In some cases, one might be normatively bound to form new mental states and intentions. For example, if I decide to climb Mount Everest, then my decision commits me to perform a variety of further actions. I must purchase a ticket to Nepal. I should then travel to the base of Mount Everest. I must then climb Mount Everest. Most importantly, I ought to first learn how to climb mountains. Moreover, I am also committed to form new mental states of intentions to perform each of these further actions. Hence, satisfying our commitments of the will stemming from personal decisions and intentions will typically involve performing bodily actions, forming new mental states, and bringing about external states-of-affairs.

Cognitive commitments are different. These commitments have satisfaction conditions that are restricted to the internal, mental domain. Indeed, to have a cognitive commitment stemming from a cognitive endorsement is to be normatively bound to mentally treat a proposition in a certain way. One's cognitive commitments do not normatively constrain one to bring about external states-of-affairs or perform physical, bodily actions.

An interrelated point is that many commitments of the will require conscious effort and mental determination to successfully carry out the means needed to satisfy the contents of those commitments. If I decide to climb Mount Everest, then I can easily fail to satisfy the corresponding commitment of the will. This commitment has complex conditions of satisfaction that require

conscious effort, mental determination, and a sequence of successfully executed behavior. This sort of complexity need not be present with cognitive commitments. Some cognitive commitments are automatically satisfied. Recall that a token belief engenders Doxastic Veridicality, which is a constitutive commitment of this attitude of belief that cannot be violated. According to Coliva, if an agent, *S*, consciously believes *p* at time, *t*, then *S* is normatively constrained to treat or regard *p* as being true at *t*, such that *S* cannot simultaneously believe both *p* and not-*p* at *t* (Coliva 2019, p. 7). Doxastic Veridicality individuates our attitudes of belief. One cannot fail to satisfy this commitment. Similar considerations apply to cognitive commitments stemming from acceptance and supposition. Many of these commitments are satisfied without conscious effort, deliberation, or successful performance of actions. For our cognitive commitments are restricted to the internal, mental domain. These commitments constrain how we ought to treat propositions in our thinking.

Moreover, some cognitive commitments only generate normative standards of *permissiveness*. These commitments do not have conditions of satisfaction that require their possessors to treat a proposition in some specific way. Rather, they prescribe how one is *permitted* to treat a proposition in thought. For example, both acceptance and supposition generate some kind of normative commitment of *agnostic permissiveness*. These commitments normatively permit one to accept or suppose a proposition that conflicts with one's other mental states and beliefs. However, one is not normatively bound or obligated to accept or suppose inconsistent things. Commitments of the will stemming from personal decisions and intentions lack this kind of normative permissiveness. Consequently, cognitive commitments and commitments of the will have different satisfaction conditions.

A third source of normative difference has to do with the acquisition conditions of cognitive commitments and commitments of the will. Many personal decisions and intentions are consciously brought about for pragmatic reasons. In many cases, it can be rational for an agent to voluntarily decide or intend to do Φ even when she recognizes she has insufficient evidential support for doing

Φ. I might, for example, have good practical reasons to run a marathon. Given these reasons, I could form a decision to run a marathon while also recognizing that I have little evidence to suggest that I will actually run a marathon. My decision is rationally based on my available pragmatic considerations. Consequently, commitments of the will can be voluntarily acquired for pragmatic reasons without having to be based on sufficient epistemic reasons. Cognitive commitments associated with acceptance and supposition can be acquired in a similar way by simply deciding to accept or suppose a proposition for primarily pragmatic reasons. In such cases, one could accept or suppose *p* while ignoring the epistemic considerations for or against *p*. In doing so, one is not inherently irrational.

However, some cognitive commitments cannot be appropriately acquired voluntarily and without regard to the available epistemic reasons one has. Indeed, explicitly formed beliefs have unique acquisition conditions. As Coliva correctly points out, the commitments of an agent's rational belief that *p* are acquired involuntarily through a judgment that *p* obtains given her epistemic reasons for the truth of *p* (Coliva 2019, p. 6). A rational belief that *p* must be based on evidence for *p*. One cannot rationally believe *p* by voluntarily ignoring the evidence for or against *p* (see chapter 1). Hence, the cognitive commitments associated with a belief that *p* cannot be rationally acquired when one voluntarily ignores the relevant epistemic considerations concerning the truth or falsity of *p*.

A fourth way in which our cognitive commitments can differ from our commitments of the will concerns how these commitments are abandoned and revised. One can terminate a commitment of the will in one of two ways: (i) one can satisfy this commitment, or (ii) one can revise or abandon this commitment. How one ought to rationally revise a personal commitment of the will can depend on whether the commitment stems from a decision or an intention. According to Gilbert, if an agent intends to Φ, then she can cease to have this commitment by simply dropping her intention to Φ or replacing it with a new and different intention to do something other than Φ (Gilbert 2013, p. 902).

However, if her commitment to Φ stems from a personal decision, then she can rationally terminate or revise this commitment by consciously rescinding her initial decision to Φ (Gilbert 2013, p. 901).

However, cognitive commitments can be rationally altered without any conscious rescission or decision to do so. According to Coliva, we typically revise or change our commitments of belief in the same way we revise personal intentions: one can revise one's belief that *p* without having to first consciously rescind one's judgment that *p* is true (Coliva 2019, p. 6).[104] For example, an agent might simply come to believe that not-*p*. In this case, she could form a new belief that not-*p* while simultaneously dropping her initial belief that *p*. This revision does not require her to first consciously rescind her belief that *p* (Coliva 2019, p. 6). The same applies to cognitive commitments associated with acceptance and supposition. An agent does not have to revise her commitments of acceptance and supposition by having to first consciously undo these commitments. She could simply drop these cognitive endorsements or form new acceptances and suppositions with inconsistent contents.

A Fifth difference deals with the context-dependency associated with these commitments. Commitments of the will are motivated in part by one's pragmatic reasons. Consequently, a commitment stemming from a personal decision or intention is typically constrained to those contexts where one has pragmatic pressures to maintain one's decision or intention. If, for example, I decide to make pasta tonight based on my current desire or preference for pasta, then my decision is guided my available practical considerations in a given context. Specifically, my decision is practically motivated by my desire to eat pasta. Insofar as my practical considerations stay the same, I need not change my mind and thus revise my commitment of the will to make pasta.

Some cognitive commitments are also bound to a context for pragmatic considerations. Acceptances and suppositions brought about for pragmatic reasons will typically be restricted to those contexts in which their possessors have the same (or similar) pragmatic pressures. However, cognitive

[104] Coliva discusses Gilbert's (2013) notions of *intentions* and *decisions* (see Coliva 2019).

commitments engendered by beliefs are context-independent. Rational beliefs are sensitive to epistemic reasons. If an agent rationally believes *p*, then her belief is based on strong epistemic support for *p*. She will appropriately continue to believe *p* through different contexts regardless of whether she has practical pressures in some contexts to revise her belief. The cognitive commitments of belief persist insofar as the belief persists. Hence, cognitive commitments occupy a wider range of contexts.

The final difference I will discuss centers on inconsistency and normative conflicts. First, personal commitments of the will can conflict with each other. I might, for example, decide to Φ. I might also come to form an intention whose satisfaction would preclude my Φ-ing. If I form this intention without rescinding my decision, then I have put myself in an inappropriate normative situation to do inconsistent things. I ought to achieve consistency with my planning and goals by making the necessary changes to my system of intentions and decisions. However, conflicts between different cognitive commitments need not entail any normative tension. I can rationally believe *p* and (pragmatically) accept not-*p*. I can be agnostic toward *p* while also supposing both *p* and not-*p*. Our suppositions and acceptances can conflict with our beliefs without error. Given the different cognitive commitments involved, these states are normatively *permitted* to be inconsistent with one's beliefs and other cognitive endorsements. In these cases, one would not necessarily be irrational.[105] An agent's inconsistent cognitive commitments do not always engender a normative tension that calls for her to revise her mental states. Hence, cognitive commitments are different from individual commitments of the will stemming from decisions and intentions.

[105] Some cognitive commitments have a normative force that *requires* one to mentally treat a proposition in a certain way. For instance, to believe *p* engenders Doxastic Rationality: one should believe in accordance with one's epistemic reasons and be sensitive to new information regarding the truth or falsity of *p*. One *ought* to satisfy this commitment. However, some cognitive commitments *permit* certain ways of mentally treating a proposition in thought. Agnostic Acceptance Permissiveness *permits* one to suppose a proposition that is inconsistent with one's beliefs. This does not mean that one always *ought to* suppose propositions that are belief-contravening. Nevertheless, these permissive commitments are constitutive elements of their associated attitudes. Consequently, one *ought* to abide by them in order to properly possess those attitudes.

8. Conclusion: The Normative Structure of Thought

Cognitive endorsements generate a plurality of normative commitments that prescribe how one ought to mentally treat or implement the target propositions of one's attitudes of belief, acceptance, and supposition. Moreover, hypothesizing engenders its own unique normative commitments constraining how one ought to treat a working hypothesis in reasoning and inquiry. The commitments of hypothesizing involve cognitive commitments and commitments of the will. The following tables specify the fundamental commitments associated with consciously formed cognitive endorsements.

Belief	
Doxastic Veridicality	If an agent, *S*, at time, *t*, believes *p* by judging that *p* obtains, then *S* is committed to regard *p* as being true at *t* (or obtaining at *t*), such that *S* cannot also consciously believe not-*p* at *t* without revising her initial belief that *p*
Doxastic Rationality	If an agent, *S*, believes *p*, then *S*'s belief commits her to base that belief on sufficient epistemic grounds supporting the truth of *p*, and to be sensitive to new information and evidence concerning the truth or falsity of *p*.
Doxastic Integrity	If an agent, *S*, forms a belief that *p* by judging that *p* is true, then *S* accrues a normative commitment to maintain *p* as being true in some doxastic form, such that, *S*'s belief can only be rationally revised or terminated via another judgment concerning the truth of *p* (e.g., a judgment that not-*p* obtains), or that, given the available evidence, not-*p* is just as likely as *p* to be true.

Table 6: Cognitive Commitments of Belief

Acceptance	
Acceptance Commitment	To accept *p* involves a normative commitment to take a positive stance toward *p*, and to employ or use *p* in thinking or reasoning.
Acceptance Rationality	To rationally accept *p* involves a normative commitment to base one's acceptance on reasons (either evidential or pragmatic) that one has in a given context, and to be sensitive to new information and rational considerations for one's acceptance.
Agnostic Acceptance Permissiveness	If an agent, *S*, accepts *p*, then *S*'s cognitive endorsement of *p* need not be consistent with *S*'s other beliefs. Nor is *S* committed to consider or regard *p* as being true.

Table 7: Cognitive Commitments of Acceptance

SUPPOSITION	
Suppositional Commitment	An attitude of supposing *p* engenders a normative commitment to momentarily present *p* in thought, such that *p* is weakly affirmed or put forward as a postulate in thinking or reasoning.
Agnostic Suppositional Permissiveness	If an agent, *S*, supposes *p*, then *S*'s supposition does not have to be consistent with *S*'s other beliefs and acceptances.
Unjustified Permissiveness	If an agent, *S*, supposes *p*, then *S*'s supposition permits her to weakly and tentatively endorse *p* in thought without having any rational basis (evidential or pragmatic) for the truth, plausibility, or practical usefulness of *p*.

Table 8: Cognitive Commitments of Supposition

HYPOTHESIZING	
Cognitive Commitments	
Acceptance Commitment, Acceptance Rationality, and Agnostic Acceptance Permissiveness	Processes of hypothesizing are instantiated by attitudes of acceptance and therefore generate the cognitive commitments associated with that state. Specifically, hypothesizing engenders the cognitive commitments of *Acceptance Commitment*, *Acceptance Rationality*, and *Agnostic Acceptance Permissiveness* (see Table 7).
Commitments of the Will	
Commitment of Hypothesizing	If an agent, *S*, hypothesizes with a proposition, *p*, such that *S* decides or intends to employ *p* in her reasoning, inquiry, or testing, then *S*'s decision or intention normatively commits her to utilize *p* as a premise in her reasoning and inquiry, or as a postulate for rational scrutiny and evidence gathering.
Non-dogmatism	If an agent, *S*, engages in hypothesizing about a proposition, *p*, such that *S* treats *p* as a working hypothesis, then *S* is normatively committed to being open-minded about *p*'s truth or falsity, and *S* commits herself to being receptive and responsive to new information that can establish *p*'s truth-status.

Table 9: Commitments of Hypothesizing

Adopting these commitments leads to some important and surprising results. First, we have a straightforward answer to the question as to whether our cognitive endorsements are properly understood in purely non-normative, dispositional terms. Our cognitive endorsements generate normative commitments. Consequently, our beliefs, acceptances, suppositions, and process of hypothesizing are inherently normative. It would be a mistake to think that cognitive endorsements are non-normative states that simply dispose their possessors to think and behave in certain ways. We cannot simply take what Millar referred to as the, "contemplative stance" toward our consciously

formed cognitive endorsements. Hence, the purely dispositional picture fails to accurately capture the normative pressures agents accrue when they cognitively endorse propositions.

A second benefit of the cognitive commitment picture is that it adequately describes and elucidates the complex normative structure of our cognitive endorsements. First, it shows how different cognitive endorsements generate different normative constraints. Some of these normative constraints dictate how one ought to treat the target propositions of one's cognitive endorsements.[106] Other commitments dictate what kind of rational support is required for a cognitive endorsement to be rationally held.[107] However, there are some normative constraints that are merely permissive. That is to say, these commitments simply permit agents to mentally endorse or fix a proposition in some specified way. By clarifying this normative structure, we come to possess a better understanding of the rationality of our beliefs, acceptances, and suppositions. Moreover, demarcating these commitments advances our understanding of how it can be rational or normatively permitted to possess different cognitive endorsements that mentally fix inconsistent propositions in one's thought.

Finally, this account provides a satisfactory way to demarcate cognitive commitments from commitments of the will. Cognitive commitments have unique normative features. For example, commitments of the will have a different kind of normative direction of fit. Personal decisions and intentions normatively constrain their possessors to either carry out some further action, form some new mental states, or bring about some state-of-affairs. However, cognitive commitments normatively constrain how a proposition is treated and implemented in thought. Hence, cognitive commitments have a mind-to-world normative fit. Consequently, our notion of *commitment* is best understood to be a broad conceptual category that includes cognitive commitments. Cognitive commitments play a crucial role in constraining how agents ought to think and reason.

---

[106] Doxastic Veridicality is an example of this sort of commitment.
[107] See, for example, Doxastic Rationality and Acceptance Rationality.

III. SEMANTIC COMMITMENTS

Adam C. Sanders

0. INTRODUCTION

Many philosophers maintain that a given belief that *p* will generate certain commitments to believe the propositions implied by *p* (Levi 1997; Millar 2004; Bilgrami 2006; Shpall 2014; and Liberman and Schroeder 2016). This is a *diachronic doxastic commitment* that normatively governs how one ought to form new beliefs over time. If there are such commitments, then satisfying them would play an important role in rational belief-revision. This chapter focuses on the normativity of diachronic doxastic commitments. I aim to do several things. First, I briefly compare and contrast these commitments with cognitive commitments. I show that these commitments differ in their normative properties. Second, I articulate several challenges facing many standard accounts of diachronic doxastic commitments. I contend that any plausible view of diachronic doxastic commitments ought to address these worries. Finally, I advance the "semantic commitment" view. This account preserves the intuitive idea that we are committed to believe what follows from our currently held beliefs. Moreover, this alternative theory avoids many challenges facing standard views of diachronic doxastic commitments. I therefore argue that the semantic commitment view is a more plausible theory.

The upshot with appealing to my account is that it adequately captures a certain kind of normative pressure constraining how one ought to revise one's system of beliefs over time. Hence, countenancing semantic commitments facilitates a deeper understanding of epistemic rationality. Moreover, this theory helps to demarcate the different kinds of commitments that agents can acquire. It is shown that semantic commitments differ from both personal commitments of the will and cognitive commitments. Our behavior and thinking are constrained by a plurality of commitments.[108]

[108] It should be noted that our cognitive commitments will be consistent. However, semantic commitments can conflict with one another. More precisely, it is possible for an agent to have distinct semantic commitments to believe inconsistent things. I address this issue in the next chapter.

The structure of the chapter is as follows. In section 1 I briefly describe several standard views of diachronic doxastic commitments. In section 2 I contrast these diachronic doxastic commitments with cognitive commitments and personal commitments of the will. Section 3 articulates several challenges facing the standard accounts. In section 4 I advance the semantic commitment view. This view adequately handles the challenges raised in the previous section. I argue that the semantic commitment view is a more tenable theory. Section 5 considers the idea that our *non-belief* cognitive endorsements generate similar diachronic commitments. I argue, however, that acceptance and supposition do not generate similar semantic commitments. I conclude the chapter in section 6.

1. Standard Accounts of Diachronic Doxastic Commitments

It is commonly thought that agents are constrained to believe the things that follow from their current beliefs. For example, consider the following remarks by Frank Jackson:

> Someone who believes that P, and that if P then Q, *ought* to believe that Q. It is not simply that, by and large, they do believe that Q. It is that if they don't, there is something *wrong* (Jackson 2000, p. 101).

The basic idea here is that, all things being equal, an agent has sufficient *reason* or sufficient *normative pressure* to believe the things that follow from her current beliefs.[109] Consequently, she *ought* to believe the consequences of her beliefs. Failure to do so would be a mark of theoretical irrationality.[110]

[109] Many theorists claim that normative pressures, reasons, and commitments can have different degrees of normative strength that can determine how one should behave (see, for example, Raz (1975) p. 25; Liberman and Schroeder (2016), pp. 104 – 120; and Lord and Maguire (2016), p. 10). Accordingly, the normative pressure to believe what follows from one's beliefs can amount to what Lord and Maguire refer to as a "decisive" reason (Lord and Maguire 2016, p. 10). An agent's reasons to do Φ are "decisive" when they outweigh any competing reasons to not do Φ (Lord and Maguire 2016, p. 10).

[110] It is possible for one to have pragmatic reasons not to believe certain consequences of one's beliefs. That is to say, it might be pragmatically rational to refrain from forming the relevant beliefs. Hence, one's pragmatic reasons can conflict with one's epistemic reasons. Nevertheless, it seems that the failure to form the beliefs that one is committed to form would still involve some degree of theoretical irrationality. In what follows I focus on theoretical and epistemic reasons for belief-revision.

What is the source of this normative pressure? Different views posit different answers to this question.[111] However, my primary focus centers on one class of views claiming that our beliefs generate a kind of *commitment*.[112] According to these accounts, one is committed to believe the consequences of one's beliefs (Levi 1997; Millar 2004; Bilgrami 2006; Shpall 2014; Liberman and Schroeder 2016).

On Isaac Levi's view, a belief is a kind of doxastic *commitment*: to possess a belief is to be responsible for that belief, and to accrue certain obligations to fulfill the norms of rationality that govern doxastic states (Levi 1997, pp. 1, 6, 10-12). He maintains that to possess a belief that *p* is to acquire a commitment to believe everything that is logically entailed by *p* (Levi 1997, pp. 10-14).

Akeel Bilgrami also thinks that beliefs can be thought of as a kind of commitment.[113] He claims:

---

[111] There can be different views concerning the normative pressure we have to believe the consequences of our beliefs. Some philosophers think that there are *requirements* or *rules* of rationality that normatively guide our practices of belief-formation. Frege, for instance, held the view that rules of logic could be understood as descriptive "laws of truth" that do not lie outside of thought, and prescriptive laws or norms governing thinking and reasoning (Frege 1964, p. 325). According to this latter sense, Frege further claimed that "the laws of logic ought to be guiding principles for thought" (Frege 1964, p. 12). Other theorists contend that agents ought to abide by external rules of rationality (see, for example, Broome 1999). In this chapter I focus on those views appealing to commitments. I will say more about rational requirements in chapter 5. I will ultimately sketch a partial theory of theoretical rationality that makes room for both rational requirements and commitments governing cognitive endorsements and belief-revision.

[112] These commitments would seem to be generated by explicit full beliefs. Both Levi and Bilgrami take beliefs as commitments to be non-reducible to any associated dispositional beliefs an agent might possess (Levi 1997, pp. 6, 10-14; Bilgrami 2006, p. 213).

[113] For Bilgrami, to possess any kind of commitment, including a belief as a commitment, necessarily involves dispositions to take on certain "reactive attitudes" (Bilgrami 2006, p. 226). He claims that a "necessary and defining condition" for possessing a commitment is to "*be prepared to have certain reactive attitudes, minimally to be self-critical or to be accepting of criticism from another, if one fails to live up to the commitment or if one lacks the disposition to do what it takes to live up to it; and one must be prepared to do better by way of trying to live up to it, perhaps by cultivating the disposition to live up to it*" (Bilgrami 2006, p. 226, original italics). One might take issue with Bilgrami's view. Consider, for example, a situation in which an agent, *S*, at time, t_1, comes to form the belief that *p*. Let us also suppose that *S* is committed to believe the logical consequences of her beliefs. Moreover, *S* believes that conjunction introduction is a legitimate inference rule. If *S*'s belief-set also includes a belief that *q*, then, according to Bilgrami's account, *S* is

> …to desire something, to believe something, is to think that one *ought* to do or think various things, those things that are entailed by those desires and beliefs by the light of certain normative principles of inference (those codifying deductive rationality, decision-theoretic rationality, perhaps inductive rationality, and also perhaps to some broader forms of material inference having to do with the meanings of words as well) (Bilgrami 2006, p. 213).

To believe *p* is, according to Bilgrami, to generate a *commitment* or an "internal ought" that normatively binds one to believe those propositions that follow (logically, materially, and so on) from *p* (Bilgrami 2006, p. 213).[114]

According to Alan Millar, if one's current beliefs imply a further proposition, *p*, then one takes on an "implication commitment" to forming a belief that *p* "*if one gives any verdict on* [*p*] *at all*" (Millar 2004, pp. 73-64, original italics). On Millar's view, to give a *verdict* on a proposition, *p*, is to take up one of the following three cognitive stances toward *p*: (1) the stance of belief that *p*; (2) the stance of disbelief that *p* (i.e., believing not-*p*); and (3) the stance of suspending judgment or withholding a belief

---

committed to believe the conjunction, $p \wedge q$, at t_1. Indeed, *S* is committed to believe the conjunction of *all* the propositional contents of her belief-set at t_1. However, if *S*'s belief-set is very large, it is not clear that *S* would necessarily have the disposition to exhibit certain negative reactive attitudes of self-criticism if she could not believe a large, complex proposition with conjunctive content composed of the contents of her individual beliefs. Grasping this complex proposition is outside her cognitive abilities. Given that she already believes the individual conjuncts, it is plausible that she would lack any inclination to correct herself and take further steps to try to live up to this commitment. Perhaps the following disjunctive claim could be defended: Either the dispositions to exhibit reactive attitudes of self-criticism are not necessary components of these doxastic-formation commitments, or one does not have doxastic-formation commitments to believe all the propositions that logically follow form one's current beliefs. I will say more about reactive attitudes and the relation between doxastic commitments and one's cognitive abilities when I develop the semantic commitment theory.

[114] I discuss the implication relations that are relevant to these commitments in sections 3 and 4.

concerning *p* (Millar 2004, p. 74).[115] Hence, if *p* follows or is implied by one's current beliefs, then one is normatively committed to believing *p* if one mentally takes a position on *p*'s truth.[116]

To illustrate the basic idea underlying these views, consider the following case:

> ***Tomato:*** At time, *t*, Joe possesses the belief (B1) *If it rains today, then I will not water the tomato plants*, and he possesses the belief (B2) *that it rains today*.

In the *Tomato* case, Joe's beliefs at time, *t*, engender a normative commitment to form a new belief. Specifically, Joe is committed at *t* to form the belief, (B3) *that I will not water the tomato plants*. Insofar as Joe continues to hold (B1) and (B2), Joe *ought* to form (B3), if he takes a cognitive position on the truth of the target proposition. For (B3) is implied by *Modus Ponens* from (B1) and (B2) at *t*. His beliefs at *t* therefore commit him to form (B3). Consequently, Joe can satisfy this commitment by forming (B3). If, however, he continues to hold (B1) and (B2), then he violates this commitment if either of the following occur: (1) he suspends judgment on the proposition expressed by (B3), or (2) he believes the negation of the proposition expressed by (B3).

We can identify the core commitment behind these standard views as follows:

> **Doxastic-formation Commitment (DC):** If an agent, *S*, possesses beliefs at time, *t*, that entail a further proposition, *p*, then *S*'s beliefs at *t* normatively commit her to cognitively endorse *p*, such that *S* can only satisfy this commitment by believing *p*.[117]

---

[115] According to Millar, an agent's belief generates a "psychological commitment" to employ the target proposition of that belief "as an assumption in [her] thinking, should the need to do so arise" (Millar 2004, pp. 122-123). This psychological commitment, he claims, establishes what he calls the "*Implication Ideal*", which he defines in the following way: "For any Π, θ, if θ is implied by Π, then avoid believing Π while giving a verdict on θ other than belief" (Millar 2004, p. 76). He claims that the Implication Ideal grounds our normative commitments to believe the things implied by our current beliefs (Millar 2004, p. 77).

[116] I take it that Millar's point here is that one does not automatically violate a doxastic-formation commitment if one does not immediately come to form the relevant beliefs. Rather, these normative constraints diachronically persist and prescribe what sort of propositions one should come to believe given one's current beliefs. If one believes *p*, and if *q* is a proposition that follows from *p*, then one continues to be committed to believing *q* until one either (1) drops the belief that *p* or (2) believes *q*. And one can only violate this commitment if one takes a mental stance toward *q* other than belief (e.g., believing not-*q* or suspending judgment on *q*).

[117] Most theorists maintaining that an agent's beliefs commit her to believe the consequences of her beliefs seem to hold some version of this view. As I understand these theorists, one need not be aware of the propositions that follow from one's currently held beliefs. I say more about this later.

DC is a conditional statement specifying a normative commitment that an agent acquires when she forms beliefs. Accordingly, an agent is committed to believe *p* if *p* is entailed by the content of one of her beliefs (or her belief-set), *B*. Insofar as she does not revise *B*, she ought to come to believe *p*.[118]

Several clarifications are in order. First, advocates of DC share the view that these commitments are intimately related to rules of logic and inference. Typically, one will be normatively committed to believe *p* insofar as *p* follows from one's beliefs in accordance to some rule of logic or inference. It should be noted, however, that rules of reasoning and logic are distinct from the commitment expressed in DC. As Liberman and Schroeder correctly point out, normative commitments associated with DC are "*agent-dependent*," such that "they are the result of some state or activity of the agent" (Liberman and Schroeder 2016, p. 111).[119] Agents bring about commitments to believe certain propositions depending on their current mental states. On the other hand, rules of logic and reasoning are independent and external to an agent's mental states. Sam Shpall argues for the same point.[120] He contends that, unlike our commitments to believe the consequences of our beliefs, the normative rules and requirements of rationality are not "grounded" in the mental states of agents (Shpall 2014, pp. 154-155).

Second, the commitment expressed by DC can be terminated in several ways. If *S*'s beliefs commit her to believing *p*, then *S* can terminate this commitment by believing *p*. In this case, *S* satisfies

---

[118] Following Millar, I take it that these commitments provide a normative reason to believe a proposition that follows from one's current beliefs (Millar 2004, p. 74). However, Millar claims that the reason involved is not the kind of reason that would epistemically support or warrant the *truth* of a proposition entailed by one's beliefs (Millar 2004, p. 74). This is correct. In what follows, I argue for a specific view of the commitments governing belief-formation. However, I will not provide a complete analysis of this notion. At this point, it will be enough to show that our beliefs generate a normative pressure that constrains how we ought to form new beliefs given our current beliefs. Moreover, I argue in the next chapter that these commitments are not to be identified with reasons, considering that these commitments might not weigh against each other in the way reasons do.

[119] See also Shpall's discussion of this feature (Shpall 2014, pp. 151-152, 154-155).

[120] Shpall calls the commitment associated with DC a "rational commitment" (Shpall 2014). I prefer a different label considering that cognitive commitments are also important for theoretical rationality.

her commitment. However, *S* can also terminate the commitment by abandoning her grounding or basing belief(s) that generate the commitment (Liberman and Schroeder 2016, p. 111). Following Shpall, we can say that the commitment specified by DC is "*escapable*" (Shpall 2014, p. 153). Hence, one is not always required to believe the consequences of one's current beliefs.[121]

Now, on the face of it, DC seems plausible. We typically think that people are normatively compelled to believe what follows from their current beliefs.[122] Hence, appealing to the *Doxastic-formation Commitment* (i.e., DC) would seem to capture and explain this normative force. In what follows I contrast this commitment with cognitive commitments and personal commitments of the will.[123] I then proceed to articulate several challenges for views endorsing DC. Finally, I advance an alternative theory that posits *semantic commitments*. This theory avoids many of the challenges facing the DC view. Hence, it is a more tenable theory of the normative commitments of belief-revision.

2. Demarcating the Doxastic-formulation Commitment from Other Commitments

Let us momentarily suppose that beliefs engender the commitment expressed by DC. If this is the case, then these commitments differ from both cognitive commitments and commitments of the will.

First, DC differs from the cognitive commitments of belief. Consider the three cognitive commitments defended in chapter 2. These commitments can be expressed (roughly) as follows[124]:

> **Doxastic Veridicality:** To believes *p* at time, *t*, is to be normatively committed to regard *p* as being true at *t*, such that one cannot simultaneously consider not-*p* to be true at *t*.

---

[121] Consequently, the normativity associated with these doxastic commitments differ from the normativity of rational requirements. As Shpall notes, "requirements are not escapable" (Shpall 2014, p. 154). I will say more about rational requirements in chapter 5.
[122] Shpall (2014) and Liberman and Schroeder (2016) discuss this idea in detail.
[123] I adopt Margaret Gilbert's view of personal commitments of the will (see Gilbert 2013 and 2014). I will briefly discuss some of the features of cognitive commitments and commitments of the will. For a more detailed discussion see chapter 2.
[124] Recall that my notions of *Doxastic Veridicality* and *Doxastic Rationality* are borrowed from Annalisa Coliva's view (see Coliva 2016, 2019). I expanded on her account (see chapter 2).

> **Doxastic Rationality:** To believe *p* at time, *t*, is to normatively commit oneself to be disposed to take one's evidence as being sufficient to establish the truth of *p* at *t*, and to be sensitive to new evidence and information concerning the truth or falsity of *p*.
>
> **Doxastic Integrity:** To believe *p* at time, *t*, is to be normatively committed to continue to regard *p* as being true (either occurrently or dispositionally) in accordance with one's evidence, until one rationally revises one's judgment concerning the truth or falsity of *p* at some future time, t_n.

Each of these commitments are constitutive elements of explicit and consciously-formed *attitudes* of belief. They are *cognitive* commitments: If one explicitly believes *p*, then that token attitude is partly individuated by commitments that normatively govern how one regards *p* in thinking and reasoning.

DC is a different kind of normative constraint. This commitment partially depends on the *propositional content* of beliefs and the *logical relations* that hold between those contents and other propositions. To see why, consider the following case where *S* believes both p and $p \rightarrow q$. Call this set of beliefs, "Γ". *S*'s belief-set, Γ, generates a doxastic-formation commitment, DC, to believe *q*. But DC is not generated by any single member of Γ. The commitment, DC, to believe *q* does not result from a belief that *p*. Nor does DC result from a belief that $p \rightarrow q$. Hence, DC is dependent on multiple doxastic attitudes, their contents, and the logical relations that hold between those contents that other propositions. Consequently, DC is not an inherent or constitutive aspect of a token *attitude* of belief. Rather, it is an *extraneous* commitment associated with the *contents* of beliefs and belief-sets.

Now, DC resembles the commitments stemming from personal intentions and decisions. Following Margaret Gilbert, when an agent, *S*, intends or decides to do some act, Φ, then *S*'s intention or decision brings about an individual "commitment of the will" to perform Φ (Gilbert 2013, pp. 899-905). Other things being equal, *S* is under a normative and rational constraint to do Φ (Gilbert 2014, pp. 6, 64). DC has a similar diachronic normative force. To possess a commitment to believe some proposition, *p*, is to accrue a normative constraint to bring about some state-of-affairs: other things being equal, one ought to bring about a mental state of believing *p*.

However, DC differs from individual commitments of the will in several important ways. First, DC and commitments of the will are based on different kinds of mental events. Individual commitments of the will stem from intentions and decisions, whereas DC is generated by representational states of belief. Second, personal decisions and intentions can be rationally formed at will for purely pragmatic reasons. An agent can voluntarily decide to do Φ without considering the epistemic reasons in favor of doing Φ. On the other hand, rational belief-formation is not a voluntary affair that one can undertake by ignoring one's evidential considerations for the target proposition. Rational beliefs are based on good evidence and epistemic reasons. Hence, one cannot rationally bring into existence doxastic-formation commitments to believe certain propositions by simply deciding to do so while ignoring the relevant epistemic reasons.[125] Finally, a commitment to believe *p* based on one's current beliefs is a commitment to *regard* something as being the case. However, a commitment to do Φ based on one's decision or intention to do Φ, is a commitment to *make* something the case.

DC therefore has a unique set of normative properties that differ from the normativity associated with cognitive commitments and individual commitments of the will. It follows from this that DC is a distinct kind of commitment governing belief-formation.

3. Challenges Facing Standard Views of Doxastic-Formation Commitments

Standard views that adopt DC face several challenges. I discuss these issues in turn.[126] I contend that any plausible account appealing to DC must provide a satisfactory solution to these challenges.

[125] Many theorists hold that a rational belief that *p* is based on good epistemic reasons for *p*, such that one cannot voluntarily ignore the available evidence for or against *p* (see, for example, Cohen 1992, p. 21; Coliva 2016, pp. 33-34; Pritchard 2016, p. 90).

[126] Some of the issues I articulate in this section have been acknowledged by other theorists. These issues center on the kind of implication relations associated with our doxastic-formation commitments, and whether one can be committed to believe trivial things (see, for example, Bilgrami 2006). However, I also identify some additional challenges that require more attention.

3.1 The Implication Puzzle

It is not clear how we are to understand the "implication" relation invoked by DC. According to Levi, we are committed to believe the *logical consequences* of our current beliefs. On the other hand, Bilgrami argues that rational commitments concern more than the class of propositions that can be deductively deduced from our current beliefs. He claims that "obviously more than logical inference will be relevant; various more material sorts of inference will be relevant, depending on" the meanings of the contents of one's beliefs (Bilgrami 2006, pp. 213-214). Moreover, he thinks that our doxastic states can generate normative commitments guided by *inductive* rules of inference (Bilgrami 2006, p. 213).

This is problematic. It is important to provide a clear and cogent account of the content of our doxastic-formation commitments. Hence, if a theory restricts the content of a doxastic-formation commitment to include one kind of implication relation, then there needs to be a satisfactory explanation as to why other implication relations should be precluded. However, many standard views are not always clear as to which implication relations are essential to our doxastic-formation commitments. Moreover, it is not clear as to whether all agents are under the same normative constraints to form beliefs according to the same rules of logic and inference. If a standard view of DC is to be a tenable theory concerning the normativity of our belief-formation, then we need a satisfactory account of the implication relations that are essential to doxastic-formation commitments.

3.2 The Problems of Trivial Beliefs, Redundant Beliefs, and Logical Omniscience

The propositional contents of our beliefs logically entail an infinite number of propositions. If our beliefs engender commitments guided by logical entailment, then it follows that one's current beliefs will engender normative commitments to form an infinite number of new beliefs. Consequently, standard DC views that incorporate the relation of logical entailment are forced to claim that, all things being equal, agents *ought* to become logically omniscient.

On the face of it, however, a commitment to be logically omniscient seems dubious. For it is not obvious that an agent's finite set of beliefs generates a heavy normative burden to satisfy an infinite number of commitments to believe an infinite number of propositions that, realistically, none of us could live up to. As Harman correctly notes, agents simply lack the cognitive capacity to form an infinite number of beliefs (Harman 1986, p. 12). Hence, positing commitments that normatively bind us to be logically omniscient seems counterintuitive.

Levi is aware of this unrealistic normative burden. He claims that the totality of commitments generated by one's beliefs cannot be fulfilled given one's inability to cultivate the kind of cognitive capacities required for logical omniscience (Levi 1997, pp. 12-14). To assuage this worry, he suggests that we should understand these doxastic commitments as being akin to "religious vows [that] often incur obligations only an angel could fulfil" – for just as we are sympathetic and understanding of those who fall short of living up to religious vows, we should also be understanding of our own limitations in fulfilling doxastic commitments (Levi 1997, pp. 1, 14). Even though we fall short of fulfilling our commitment to believe the logical consequences of our beliefs, Levi thinks that our doxastic commitments will further commit us to strive to improve "the extent to which [an agent] meets the demands of logical omniscience" (Levi 1997, pp. 14-15).

Another possible response is articulated by Bilgrami. According to Bilgrami, the heavy normative burden generated by our beliefs does not pose a serious challenge to his view. For he thinks that a given commitment can have "excusing conditions," that is, conditions that, if satisfied, would excuse one from living up to that commitment (Bilgrami 2006, p. 372, fn. 8). He insists that we can be pardoned from satisfying our commitments to believe an infinite number of propositions if the costs of doing so would be incredibly high – for we simply lack the time or capabilities to form that many beliefs (Bilgrami 2006, p. 372, fn. 8).

One might, however, be dissatisfied with these responses. These views still insist that a given agent's belief-set, Γ, at some time, t, will normatively commit her to believe an infinite number of propositions that follow from the contents of Γ. Call this new entailed set of beliefs, "$\measuredangle$". An astronomically large number of the propositions that comprise $\measuredangle$ are trivial and uninteresting. Consequently, S is committed at t to believe an incomprehensibly large number of trivial and uninteresting propositions. However, to claim that we have this heavy normative burden would conflict with another plausible rational requirement governing belief-formation that Harman defends:

> *Clutter Avoidance* One should not clutter one's mind with trivialities (Harman 1986, p. 12).

If we accept *Clutter Avoidance*, then it is not clear how to reconcile this requirement with our normative commitments to believe an infinite number of *trivial* propositions.

Consider, for example, a case where I believe the proposition, *that it is raining*. Call this proposition, "*p*". The content of my belief entails an infinite number of other propositions. For example, disjunction introduction rules of classical logic insist that *p* entails any disjunction that has *p* as one of its disjuncts. Hence, according to standard views of DC, my belief in this situation commits me to forming the further belief *that p or q*, where "*q*" denotes *any* proposition. Thus, I would be committed to forming the belief *that it is raining or aliens will abduct me tonight*, the belief *that it is raining or aliens will abduct me in two nights*, the belief *that it is raining or aliens will abduct me in three nights*, and so on *ad infinitum*. Considering that we are supposing that these are *normative* commitments, I ought to satisfy them. The problem, however, is that this diagnosis of my normative situation does not seem correct. Moreover, it seems that I ought not to believe an infinite number of trivial things.

Furthermore, one might be suspicious of the idea that commitments generated from beliefs have the sort of excusing conditions that Bilgrami speaks of. It is not obvious that one can be excused

from satisfying a commitment of rationality simply because the costs of doing so would be too high.[127] If I am rationally required to maintain some appropriate set of beliefs that stand in appropriate logical relations, then I ought to satisfy that commitment and get my beliefs in order. It does not seem correct to say that I can maintain an inappropriate set of beliefs and yet still be considered rational simply because correcting my inconsistent or inappropriate belief-set would be too difficult. Furthermore, it is not clear what the threshold would be for when one's doxastic-formation commitments become too costly. If these commitments exist, then, on pains of being irrational, either one must satisfy them, or one must abandon them by rationally eliminating the grounding attitudes that engender them.

3.3 The Problems of Objectivity and Cognitive Incapability

Another potential problem for DC centers on the cognitive abilities and limitations that agents have. The ability to grasp propositions and the conceptual content of one's beliefs, along with the degree of understanding one has concerning the things that follow from one's beliefs, will vary significantly between individual agents. If this is the case, then some agents will be limited in the kind of inferences they can make, and the kind of propositions they can competently identify as being consequences of their current beliefs. This raises the question as to whether one can be normatively committed to believe the things that they do not comprehend or recognize as following from their beliefs.

Consider a case where an agent, *S*, believes a proposition, *p*. Suppose that *p* implies some complex proposition, *q*, that exceeds *S*'s cognitive grasp. That is to say, *S* lacks the cognitive abilities and concepts required to *fully* comprehend the content of *q*. Let us further suppose that *S* comes to hear about *q* through the testimony of some agent. However, given her cognitive position, *S* suspends judgment toward *q* and thus fails to believe *q*. How are we to make sense of *S*'s normative situation?

[127] It is plausible that some commitments have excusing conditions. However, I am focusing on commitments of epistemic rationality.

According to standard accounts of DC, *S* has failed to live up to her doxastic-formation commitment. All things being equal, *S* is in an irrational state. However, it seems implausible to charge *S* with irrationality. For it could be argued that, given her limited conceptual resources and her inability to grasp how *q* follows from *p*, *S* has done the appropriate thing by suspending judgment. For she lacks the cognitive ability to make a rational judgment concerning the truth of *q*. Consequently, it is not the case that *S* is in an irrational state.

The challenge for a proponent of DC, then, is to adequately address this issue and explain whether or not one can be under a normative constraint to believe things that one cannot understand or grasp. My contention here is that doxastic-formation commitments are subjective and dependent on one's cognitive abilities, understanding, and conceptual resources. Agents lack any normative constraint to believe the things that fall outside their understanding. Indeed, any plausible account of doxastic-formation commitments should accommodate the subjective differences in cognitive abilities and conceptual resources.

3.4 The Satisfaction Challenge

Proponents of DC typically maintain that the *only* way to satisfy a doxastic-formation commitment to cognitively endorse a proposition, *p*, is to form the relevant belief that *p*. To take any other stance (e.g., believing not-*p* or suspending judgment on *p*) is a violation of one's commitment.

I contend that we have good reason to reject this view. First, belief is not the only way to cognitively endorse a proposition. Sometimes we *accept* propositions that we do not believe. On my view, an attitude of *accepting p* involves two major components: (1) taking a positive mental stance toward *p*, and (2) adopting a mental resoluteness to posit and implement *p* in one's thinking, given one's available reasons for *p*.[128] One's rationally held acceptance is typically voluntarily brought about.

[128] See chapter 1 for a detailed discussion of acceptance and how this state differs from belief.

Acceptance is sensitive to both epistemic and pragmatic considerations (see, for instance, Cohen 1992, pp. 4-5, 12; Bratman 1999, pp. 20-21). Moreover, one can accept *p* by mentally implementing *p* in different ways. For example, one can accept *p* by simply putting *p* forward in thought as either given, plausible, approximate, or beneficial depending on one's reasons and aims. In doing so, one mentally endorses *p* and takes up a disposition to be sensitive to new evidence and information concerning the truth, plausibility, or usefulness of *p*. Another form of acceptance involves committing oneself to implementing *p* as a premise in reasoning for epistemic or pragmatic reasons.[129] Finally, *acceptance* is a robust conceptual category that includes different kinds or species of acceptance depending on one's goals. *Epistemic* acceptance is primarily based on evidential considerations and epistemic goals. On the other hand, *pragmatic* acceptance is primarily based on pragmatic considerations and practical goals. These different forms of acceptance need not involve a commitment to think that the target proposition is actually true.[130] Consequently, acceptance is a distinct kind of cognitive endorsement.

Given that acceptance is a type of cognitive endorsement, it is plausible to say that we can sometimes *partially* satisfy a commitment to believe *p* by merely accepting *p*. Consider the following situation. At time, t_1, Linda believes *p* and she believes $p \rightarrow q$. Suppose that Linda has not yet taken a mental position on *q*. Linda's beliefs at t_1 would engender a doxastic-formation commitment to believe *q*. If Linda comes to believe *q* at some future time, t_2, then she satisfies her commitment.

But consider a different case. Suppose that Linda's counterpart, Linda*, believes *p* and believes that $p \rightarrow q$ at t_1. Suppose that at t_2 Linda* fails to believe *q*. Perhaps she is not entirely convinced by

[129] Cohen claims that "…to accept that *p* is to have or adopt a policy of deeming, positing, or postulating that *p* – i.e., of including that proposition or rule among one's premisses for deciding what to do or think in a particular context, whether or not one feels it to be true that *p*" (Cohen 1992, p. 4). My view of acceptance is broader. Mental policies of putting *p* forward as a premise in reasoning can be sufficient for acceptance when one is also disposed to take a positive mental stance toward *p*. However, a mental policy of premising *p* is not necessary for accepting *p*.

[130] Of course, accepting *p* still involves an endorsement of *p* and a mental commitment to posit *p* as *given* in one's thinking. But this does not mean that one is committed to think that *p* is actually true.

q at t_2.[131] Nevertheless, Linda* recognizes that *q seems* to follow from her beliefs. In light of this, she *accepts q*, and she mentally posits *q* in her thinking.

Unlike Linda, Linda* does not come to believe *q*. However, it is mistaken to say that Linda* has categorically violated her normative commitment to endorse *q*. For she cognitively endorses *q* and takes on a mental resolution to implement *q* in her thinking. Moreover, she correctly recognizes that *q* seems to follow from her current beliefs. In fact, Linda* would be normatively better off than any other agent in her situation that merely suspended judgment on *q*, or came to believe not-*q*.

This suggests that our doxastic-formation commitments can be *fully* satisfied or *partially* satisfied. To fully satisfy a doxastic-formation commitment to endorse *p* requires believing *p*. To partially satisfy this commitment requires accepting *p*.[132] One might reformulate (DC) as follows:

---

[131] Depending on one's view of belief-formation, we might also point to cases in which one is not able to form beliefs via correct inferences due to stress or fatigue. In these situations, one might not believe what follows from their current beliefs. However, they could still accept these deducible propositions.

[132] The satisfaction worry is exacerbated in cases involving so-called "hinges". Roughly, Wittgenstein claims that we endorse certain fundamental propositions that function as the "grounds" for our believing other propositions (Wittgenstein 1969, remarks 203 – 209). These grounding propositions lack proper justification warranting rational belief, and thus they operate as "hinges" that we simply take for granted (Wittgenstein 1969, remarks 341 – 343). For instance, the proposition, *that there is an external world* is a hinge that we endorse in order to justify believing other propositions. According to some contemporary hinge theorists, the attitude that one takes toward this particular hinge (and other hinges) is distinct from rational belief (see, for example, Coliva (2015) and Pritchard (2016)). On Coliva's view, one assumes or accepts hinge propositions (Coliva 2015, p. 128 – 129). On Pritchard's view, one's attitude towards a hinge proposition is a non-belief "hinge commitment" (Pritchard 2016, p. 92, 102). Hence, if we accept the view that we lack sufficient justification to rationally believe hinge propositions, then we will have cases where one is normatively committed to believe a hinge proposition that they cannot rationally believe. To borrow an example from Moore, an agent might reason as follows: (1) I have a hand; (2) If there is a hand then there is an external world; so, (3) There is an external world (for a discussion of Moore's proof see Coliva 2015, pp. 86 – 96). Now, in this situation, if the agent believes the contents of (1) and (2), then her beliefs would, according to standard theories employing DC, commit her to form the belief *that there is an external world*. However, as Coliva points out, the content expressed in (3) is not justified considering that one must first assume the truth of (3) to justify (1) (Coliva 2015, pp. 93 – 95). For Coliva, one can only accept or assume the hinge *that there is an external world* (Coliva 2015, pp. 128 – 129). Consequently, one would always be in error for failing to form an unwarranted belief toward this hinge. But this would be the wrong diagnosis of this scenario. If there are doxastic-formation commitments, then they should not normatively require one to believe things that lack sufficient justification. On my alternative view, we can begin to alleviate this worry by countenancing partial satisfaction conditions. Accordingly, taking a non-belief

(DC*) If an agent, *S*, believes a proposition, *p*, and *p* implies some proposition, *q*, then *S*'s belief engenders a normative commitment, *C*, to cognitively endorse *q*, such that *S* can either fully satisfy *C* by believing *q*, or partially satisfy *C* by accepting *q*.

(DC*) stipulates a more plausible formulation of the satisfaction conditions for doxastic-formation commitments.[133] I conclude that any tenable theory of doxastic-formation commitments ought to make room for partial satisfaction conditions.

4. Semantic Commitments

Many of the challenges discussed in the previous section arise because the standard views of doxastic-formation commitments do not focus on the *subjective* cognitive positions that individual agents possess. In this section I advance an alternative theory that posits what I refer to as "semantic commitments". A *semantic commitment* is a type of doxastic-formation commitment that is grounded in an agent's beliefs, cognitive abilities, understanding, and dispositions to form new beliefs. One's grasp of the *meanings* of things and the logical relations that hold between propositions and one's beliefs is essential to the type of doxastic-formation commitments agents can accrue. It is for this reason that these normative constraints are referred to as "semantic commitments". After articulating the

endorsement of acceptance toward a hinge proposition implied by one's current beliefs does not result in a normative failure. Moreover, my view has the resources to eliminate the problem entirely. For on my view, one is not committed to believe things one already believes (I discuss this in the next section). I could also insist that one is not committed to believe things that one already believes *or* accepts.

[133] An alternative view might dispense with different satisfaction conditions and hold that the content of a diachronic doxastic commitment changes in cases where one is not in a cognitive position to easily form rational beliefs whose contents are the propositions deducible from one's current beliefs. In these cases, one might only accrue a commitment to *accept* the propositions that follow from one's beliefs. Accordingly, it is not always the case that there are partial satisfaction conditions. Perhaps this is a possible view of the normativity of diachronic reasoning. However, beliefs indicate how one takes the world to be. When an agent reasons from her current beliefs to form new cognitive endorsements, she is normatively pressured to believe the things that she takes to follow from her beliefs. If *p* follows from her current beliefs, and her current evidence rationally warrants a belief that *p*, then she has a commitment to believe *p*. If she fails to believe *p*, then she can still partially satisfy her commitment by accepting *p*. But it might not be correct to say that her commitment will change to one of accepting *p* simply when she is not in a position to easily satisfy the original doxastic commitment.

semantic commitment view, I argue that the semantic commitment account is a more tenable theory of doxastic-formation commitments.

To explicate the semantic commitment account, let us first recall some salient features of belief. To believe a proposition, *p*, is to regard *p* as being true.[134] As Alan Millar puts the point, a "belief constitutively aims at truth" (Millar 2004, p. 123). Hence, a belief is understood to be a cognitive state of settling matters concerning the truth and epistemic acceptability of a proposition.

As a first approximation, the semantic commitment account holds that an agent's belief-set, *Γ*, can commit her to forming a new belief that *p*. In this case, the contents of *Γ* will imply *p*. Consequently, *Γ* partly commits the agent to settle matters concerning the truth of *p*.[135] Semantic commitments are therefore commitments of content: they partly depend on an agent's beliefs, the contents of those beliefs, and the logical relations between those contents and further propositions. Like the DC commitment of standard views, semantic commitments are *extraneous* commitments that are not inherent in token attitudes of belief.[136] This is the first major feature of semantic commitments.

A second feature concerns the content of the commitment. If *S*'s belief-set, *Γ*, implies a further proposition, *p*, and *Γ* partly commits *S* to believe *p*, then *S* is normatively bound to settle matters regarding the truth of *p*. The content of the commitment engendered by *Γ* includes a constraint to bring about a new state-of-affairs, namely, an additional mental representation of treating *p* as being true. This is a general feature of the content of any given semantic commitment. A semantic commitment normatively requires its possessor to revise her current belief-set by cognitively settling

---

[134] Many theorists hold this view (or some similar view). See, for example, Wayne Davis (1988) p. 170. I discuss the conceptual connection between beliefs and the truth in chapter 1.

[135] It is important to keep in mind that an agent's belief-set will only *partly* commit her to form further beliefs. The content of this commitment is also dependent on the agent's understanding, conceptual stock, and cognitive abilities. I will discuss this in more detail.

[136] Recall that a commitment associated with a cognitive endorsements is *extraneous* if and only if it is not an inherent or constitutive part of that attitude. If it is not an inherent or constitutive part of the attitude, then it is partly dependent on the content of the attitude and the logical relations holding between that content and other propositions.

matters on the truth and plausibility of a proposition that follows from how she currently takes the world to be.

Now, one is not committed to believe every proposition that follows from one's currently held beliefs. This is an important point that demarcates the semantic commitment view from rival theories. On the semantic commitment view, one is committed to cognitively endorse those propositions that one is cognitively *able* to endorse. Consider the following cases.

> ***Case 1:*** Al believes that mammals can nurse their offspring with milk. He also believes that all mammals have vertebrae. Moreover, Al possesses the concept of *disjunction* and comprehends classical rules of logic. Hence, he comprehends and accepts disjunction introduction rules. At some future time, Al is hiking in the woods and perceives what he takes to be a mammal nursing her young. From his perceptual states he forms the belief, (B1) *that the creature before me is a mammal*. He then contemplates the possibility that he is mistaken. While maintaining (B1) he comes to form the further belief (B2) *that either the creature before me is a mammal or the creature before me is not a vertebrate*.

Now, contrast this case with a different kind of case.

> ***Case 2:*** Bea believes that mammals can nurse their offspring with milk. Bea, however, is not familiar with many other salient features of mammals. For instance, she lacks the belief that mammals are vertebrates. Indeed, she does not fully grasp the meaning of the terms, "vertebrae" and "vertebrate". Moreover, Bea partially grasps the concept of *disjunction*. Furthermore, she is not able to competently employ certain inference rules like disjunction introduction. At some future time, Bea is hiking in the woods and sees what she takes to be a mammal nursing her young. From this perception she forms the belief, (B1) *that the creature before me is a mammal*. At that time, she meets Al. Al tells her that he possesses the belief (B2) *that either the creature before me is a mammal or the creature before me is not a vertebrate*. After considering Al's testimony, Bea does not form the same belief (B2). She is not confident that the proposition expressed by (B2) is true. Rather, she withholds belief regarding the proposition expressed by (B2).

In the above cases, the agents come to possess different doxastic states based on their initial beliefs. Al comes to believe (B2) *that the creature before me is a mammal or the creature before me is not a vertebrate*. The proposition expressed in this belief follows via disjunction introduction from the content of his belief (B1) *that the creature before me is a mammal*. Bea in *Case 2* also believes (B1) *that the creature before me is a mammal*. However, Bea fails to form the further belief (B2) *that the creature before me is a mammal or the creature before me is not a vertebrate*. Rather, she withholds belief regarding the proposition

expressed in (B2).[137] This discrepancy in doxastic states is due to differences concerning the agents' cognitive positions. For Al has different beliefs than Bea. He also has a deeper understanding of mammalian physiology. Moreover, his conceptual repertoire differs from Bea's, and he comprehends and accepts different rules of inference. Consequently, Bea is not in the cognitive position to rationally form the same disjunctive belief (B2) that Al possesses.

If we suppose that a standard view of DC is correct, and thus maintain that we have doxastic-formation commitments to cognitively endorse or believe all the propositions that are entailed by our current beliefs, then we would have to say that Bea in *Case 2* has done something wrong when she fails to believe the disjunctive belief (B2) *that the creature before me is a mammal or the creature before me is not a vertebrate*. For the proposition expressed by (B2) logically follows via disjunction introduction from her initial belief (B1) *that the creature before me is a mammal*. Hence, Bea would be committed to form the belief (B2). By failing to believe the proposition expressed by (B2) when she contemplates it, Bea would not live up to her commitment. By withholding belief, she has not done what she was normatively required to do. According to standard views of DC, she is irrational. But this diagnosis of Bea's situation is inappropriate. It is incorrect to say that Bea is in error for withholding belief regarding the proposition expressed by (B2). To see why, consider the following line of reasoning.

First, the existence of a commitment will license criticism when there is a failure to appropriately satisfy the content of that commitment.[138] For to be committed to doing Φ is to accrue a type of normative constraint: other things being equal, one ought to bring about Φ. To do otherwise

[137] It might be tempting to credit Bea with an attitude of suspended judgment toward the proposition expressed by (B2). Following Jane Friedman, suspending judgment seems to be an explicit state that differs from mere non-belief (Friedman 2013, pp. 177 – 180). However, given that Bea does not have the conceptual resources to fully grasp this proposition, she might lack the ability to suspend judgment on the proposition expressed by (B2). Instead, it is more likely that she simply withholds belief. Accordingly, she consciously refrains from taking any doxastic stance toward the content of (B2). At any rate, my argument will succeed if Bea either suspends judgment or consciously withholds belief.
[138] Many theorists would agree on this point (See Bilgrami (2006), Gilbert (2013), and Coliva (2016)).

is to act in error. It follows that commitments license criticism when one fails to live up to one's commitments. This includes a willingness to be self-critical of oneself when one fails to honor and satisfy one's commitments (see Bilgrami 2006, p. 226). Second, if one possesses a commitment to believe *p*, such that one ought to form this belief instead of taking any other mental stance toward *p* (e.g., believing not-*p*, suspending judgment on *p*, or consciously withholding belief regarding *p*), then one is subject to criticism (including self-criticism) when one fails to live up to this commitment. Hence, if Bea has a genuine commitment to form the belief (B2) in *Case 2*, then she is subject to criticism (including self-criticism) for failing to form (B2) and for taking some other mental stance toward the target proposition. Third, it seems inappropriate to subject Bea to genuine criticism for failing to form the belief (B2) in *Case 2*. Indeed, she lacks the cognitive abilities and understanding needed to fully grasp the content of (B2). Bea therefore lacks a genuine commitment to form the belief (B2) in *Case 2*.

The claim that it is inappropriate to subject Bea to criticism for failing to form the belief (B2) is defensible for several reasons. Given the cognitive abilities, conceptual resources, and belief-system that Bea possesses in *Case 2*, she is not in the cognitive position to *rationally* form the belief (B2). For she does not fully grasp the content of (B2). Nor is she able to competently see how (B2) follows from her initial beliefs. The intuition here is that withholding belief on the proposition expressed by (B2) is the rational thing for her to do. Consequently, it would be inappropriate to criticize her for failing to form this belief. Moreover, it is permissible for Bea to defend her mental stance given her cognitive situation. If it is not obvious that others would be licensed to criticize Bea in *Case 2* for failing to from (B2), and Bea can justifiably defend her position to withhold belief in this situation, then Bea is not under any normative constraint to form the belief (B2). It follows that Bea has not violated any normative commitment to believe (B2). This consequence, however, militates against the DC view.

My contention is that the proper diagnoses of *Case 1* and *Case 2* is that Al and Bea have different doxastic-formation commitments. Specifically, they have different *semantic commitments*. A *semantic commitment* is a normative constraint to cognitively endorse the propositions that follow from one's current beliefs. However, unlike standard views appealing to DC, this particular commitment depends on the unique cognitive positions that individual agents possess. We can begin to clarify this idea by drawing on Robert Audi's notion of a "disposition to believe" (Audi 1994). Audi claims that "a disposition to believe is a readiness to form a belief" (Audi 1994, p. 424). However, for an agent, *S*, to be ready to form a belief that *p* requires that *S* possesses certain cognitive capabilities and dispositions. Similarly, to have a semantic commitment to believe *p* requires being in the appropriate *cognitive position* with the appropriate doxastic dispositions to rationally form a believe that *p*.

An agent's *cognitive position* at a given time, *t*, is comprised of her mental states, understanding, cognitive abilities, and cognitive dispositions that guide her thinking and reasoning. For my purposes, I focus on three main elements that contribute to one's cognitive position. On my view, these elements ground the semantic commitments one has to believe or cognitively endorse certain propositions.

First, an agent's cognitive position at time, *t*, includes her set of beliefs at *t*. These beliefs indicate how she takes the world to be. Moreover, the contents of her beliefs provide a basis for reasoning and inference. An agent can posit her beliefs as premises in her thinking and reasoning.

Second, an agent's cognitive position at *t* includes the concepts she has at *t*. An agent's stock of concepts helps to determine which propositions she can grasp and understand. This facilitates the formation of new cognitive endorsements and puts her in the position to incorporate new information into her system of beliefs. Moreover, one's conceptual repertoire partly determines which inferences and propositional operations one can perform. If, for example, an agent has the concept of *conjunction*, then she can mentally integrate her beliefs with other propositions and beliefs that she possesses.

Finally, an agent's cognitive position at time, *t*, will include her understanding and cognitive abilities to conduct certain lines of reasoning. Reasoning can involve different kinds of inference. One can, for example, employ a deductive inference in reasoning. Reasoning by *Modus Ponens* is a good example of a process of deductive reasoning. From both (1) p and (2) $p \rightarrow q$, one can infer (3) q. The proposition, q, logically follows from (1) and (2). Indeed, if (1) and (2) are true, then (3) must be true. But one can also reason inductively. One pattern of inductive reasoning goes as follows. An agent posits in her thought that every observed *F* has been a *G*. She then concludes from this that the next *F* she observes will be a *G*. Unlike deductive reasoning, the truth of the premises in inductive reasoning do not guarantee the truth of the conclusion. But true premises employed in inductive reasoning can make the conclusion likely. Finally, we can employ inference to the best explanation or abductive reasoning. In cases of inference to the best explanation, one might posit some proposition, p, as obtaining in her thought. From this she might infer that q is the best explanation of p.

Agents can differ in how they understand the relations between propositions. Moreover, they can differ in their abilities to employ rules of logic and inference in reasoning.[139] Indeed, the inference rules that one accepts, utilizes, or competently employs in reasoning will determine what sort of inferences one can make and which propositions one can cognitively endorse. In cases of doxastic reasoning, one employs beliefs as mental premises.[140] If an agent, *S*, believes p, and *S* posits this belief

[139] I do not provide an account of what reasoning consist of. I take reasoning to be a kind of mental process of positing certain propositions as premises and drawing conclusions. Some philosophers have discussed the relationship reasoning has with rules of inferences (see Boghossian 2014 and Broome 2014). For my purposes, we can supplement my rough remarks on reasoning with John Broome's account of reasoning via rule-following. Broome claims that if you engage in reasoning, then you will "follow a rule", such that the "rule *guides* you and you actively follow it" (Broome 2014, p. 624). For Broome, following an inference rule, *R*, in reasoning is to be "intentionally guided" by *R*, such that "you intend to comply with" *R* (Broome 2014, p. 631). Here the intention to comply with *R* can be "deliberate" or one of habituation (Broom 2014, p. 631). I think this is a plausible account of reasoning. However, my argument here does not depend on the truth of this account. The semantic commitment view is compatible *mutatis mutandis* with other views of doxastic inference and reasoning.

[140] Doxastic reasoning is a mental process of belief-revision that involves positing certain beliefs as mental premises. Not all reasoning is doxastic. Sometimes we suppose a proposition, p, to see what

as a premise in her reasoning, then *S* takes *p* to be a true premise. She can deduce certain propositions from this premise. If her doxastic reasoning is aimed at truth, she can, in optimal circumstances, come to correctly believe what is implied by her premise belief. In some situations, her doxastic reasoning can lead her to revise or abandon her initial premise beliefs. What is important, however, is that the specific norms and rules of inference that one can competently grasp and employ in thought will affect which propositions one can come to justifiably believe in doxastic reasoning. These abilities and dispositions will therefore shape one's cognitive position at a given time.

To illustrate how this shapes one's cognitive position, consider a possible situation involving a creature that can employ inductive reasoning, but denies that one can rationally believe the consequences of inductive inferences. Indeed, this individual believes that one should only believe the obvious deductive consequences of one's premise beliefs in conscious reasoning. This creature would not share the same rational standards that normal human beings typically abide by. Hence, it seems incorrect to say that this creature would be committed to believe the inductive consequences of its current beliefs. Unlike typical human agents, this individual occupies a distinct cognitive position that is subjectively stricter on the type of inferences one can justifiably employ to form new beliefs. Consequently, this agent would have different semantic commitments.

This applies to other inference rules and norms of reasoning. The crucial point is that the inference rules that an agent is willing and able to competently employ in doxastic reasoning will contribute to her cognitive position. For it is these rules and norms of reasoning that dictate which propositions she can justifiably deduce and believe from her premise beliefs.[141]

---

follows from *p*. This is suppositional reasoning. Suppositional reasoning does not involve thinking that the proposition supposed is true. Once can be agnostic toward the truth of what is supposed. Or one could believe in the negation of what is supposed (see Rescher 1964, pp. 3-4).

[141] Several clarificatory points are in order. First, this does not mean that for an agent to reason at time, *t*, she must have a conscious belief at *t* regarding the type of rule that she is using when she performs her reasoning. According to John Broome, reasoning according to some inference rule can proceed via a tacitly held or habituated intention that is not consciously and deliberately formed (Broome 2014,

The final aspect of semantic commitments I will discuss here centers on their termination and satisfaction conditions. First, a semantic commitment can be abandoned by simply giving up the belief (or set of beliefs) that give rise to that commitment. For example, an agent, *S* might possess a belief and cognitive position that semantically commits her to believe *p*. Consequently, *S* is normatively committed to believe *p* when *S* cognitively settles matters concerning the truth of *p*. However, let us suppose that *S* recognizes that *p* is false or doubtful. Perhaps *S* has good epistemic reasons supporting not-*p*. Following standard accounts of doxastic-formation commitments, *S* can rationally abandon her commitment to believe *p* by rationally revising her grounding beliefs. Semantic commitments can be rationally terminated when the beliefs that engender them are rationally terminated. In these cases, one revises the cognitive position that is a precondition for the semantic commitment.

The other way in which one can rationally eliminate a semantic commitment is by satisfying that commitment. If the content of *S*'s belief-set, *Γ*, implies some further proposition, *p*, and *S* is committed to cognitively endorse *p* given her cognitive position, then *S* can fully or partially satisfy this commitment. *S* fully satisfies this commitment if she forms the appropriate belief that *p*. She only partially satisfies this commitment by accepting *p*. As noted earlier, acceptance is a kind of cognitive endorsement that is distinguished from belief. To accept a proposition, *p*, involves a commitment to utilize *p* in thought, and to be disposed to take a positive mental stance toward *p* when one contemplates *p*.[142]

Consider a case where an agent, *S*, has a semantic commitment to believe *p*. In this situation, *S* might lack sufficient evidence to rationally believe *p*. Moreover, she could lack sufficient evidence to revise her grounding beliefs that imply *p*. In this situation, *S* could still accept *p* in light of the fact

---

pp. 630-631). Second, I am not suggesting that agents must always employ these inference rules correctly. As Broome points out, we sometimes reason incorrectly by employing and following incorrect rules of inference (Broome 2014, p. 624). Rather, the idea here is that one has the ability and willingness to grasp and reliably utilize (to some appropriate degree) an inference rule in reasoning.

[142] See also chapter 1.

that she recognizes that *p* is implied by her beliefs. By accepting *p*, she only partially satisfies her normative commitment. Intuitively, agents in these cases would be doing something normatively right. They do not irrationally believe the implied proposition, nor do they revise their grounding beliefs without good reason to do so. Moreover, they recognize or comply with correct inference rules and come to cognitively endorse what follows from their beliefs. If this is right, then semantic commitments have partial satisfaction conditions.

As I argued in the previous section, it is reasonable to think that a doxastic-formation commitment to believe *p* is not categorically violated if one merely accepts *p*. Of course, an agent does not fully live up to her normative standards by partially satisfying her semantic commitment. In these cases, she is not quite right in how she ought to think about the target proposition, *p*. However, partially satisfying a semantic commitment is normatively better than an outright violation of that commitment (e.g., believing not-*p* when one has a semantic commitment to believe *p*). Satisfying our semantic commitments comes in degrees.

We are now in the position to fully articulate the semantic commitment view.

> **Semantic Commitment:** If an agent, *S*, has a belief (or non-empty set of beliefs) that implies a proposition, *p*, via some correct norm or rule of inference, *R*, and *S* possesses the relevant concepts, understanding, and cognitive abilities to settle matters concerning *p*'s truth, such that she can competently form a new belief that *p* via *R*, then *S* is normatively committed to cognitively endorse *p* when she mentally determines the truth or epistemic acceptability of *p*. *S* can fully satisfy this commitment by believing *p*, or she can partially satisfy this commitment by accepting *p*.

The *semantic commitment* view, I contend, is a plausible theory of doxastic-formation commitments. First, an agent, *S*, will typically accrue a genuine *constraint* to cognitively endorse a proposition, *p*, when *p* follows from her beliefs, and *S* is in a cognitive position to competently grasp and believe *p*. Any failure to satisfy this constraint in the appropriate ways can license criticism. This suggests that the constraint is inherently *normative*: an agent's beliefs and cognitive position will *normatively* prescribe how she *ought* to think about certain propositions. Moreover, the source of this

normativity is best explained by the existence of a *commitment*. For this normative element is brought about by the subjective mental states and cognitive abilities that agents possess.

Second, appealing to semantic commitments provides an intuitive picture of the normative pressure and rational motivation agents experience when undergoing doxastic-reasoning. Other things being equal, an agent will be disposed to either (1) believe what she competently recognizes as being a consequence of her current beliefs, or (2) revise her current belief-system and cognitive position. Failure to satisfy her commitment results in error. In these cases of failure, rational agents will be disposed to correct themselves. The semantic commitment theory can accommodate this normative pressure. Moreover, the semantic commitment view can correctly diagnose and handle a wide range of cases, including *Case1* and *Case 2* above. This account has explanatory and predictive power. Consequently, the semantic commitment theory is a plausible view of rational doxastic-formation.

I further contend that the semantic commitment account is preferable to standard views of doxastic-formation commitments. This claim is established by the following Modus Ponens argument.

(P1) The semantic commitment account avoids or ameliorates several challenges and puzzles facing the standard accounts.

(P2) If the semantic commitment account avoids or ameliorates several challenges and puzzles facing the standard accounts, then the semantic commitment account is a more tenable and preferable theory of doxastic-formation commitments.

(C) The semantic commitment account is a more tenable and preferable theory of doxastic-formation commitments.

The second premise (P2) is fairly obvious. Hence, I will proceed to defend (P1). The semantic commitment account can surmount and alleviate the challenges and puzzles discussed in section 3.

First, the semantic commitment view provides a simple and straightforward answer to the so-called "Implication Puzzle" (section 3.1). Let us recall that the semantic commitment view does not maintain that all agents are normatively bound to revise their beliefs according to the same rules of inference. Nor does this view take a specific stance on which implication relations are relevant to our doxastic-formation commitments. Rather, the implication relations relevant to an agent's semantic

commitment are determined by her subjective cognitive position. If *S*'s belief implies *p* by some norm or inference rule, *R*, and *S* is in a position to rationally accept and competently employ *R* in her doxastic reasoning to form a new belief that *p*, then *S* can accrue a commitment to believe *p*.[143]

Second, the semantic commitment view ameliorates the worries concerning logical omniscience and the formation of trivial beliefs (section 3.2). No doubt an agent's semantic commitments will normatively compel her to form a large number of beliefs if she maintains her current set of beliefs. However, the normative force of semantic commitments is restricted to accommodate the subjective abilities that agents possess. Consider a case where an agent, *S*, believes *p* and *S* possesses the relevant concepts to grasp and utilize disjunction introduction. Indeed, *S* believes that she can deductively deduce any disjunction from her belief that *p*. On the semantic commitment account, it is not correct to say that *S* is committed to believe an infinite number of disjunctive propositions that have *p* as a part of their content. For *S* is not in the cognitive position to form an infinite number of beliefs. Similar considerations apply to commitments to believe a single, complex proposition composed of an infinite number of atomic propositions. Typical human beings lack the cognitive capacities to comprehend and believe overly complex things. Hence, one's cognitive position will restrict the kinds of doxastic-formation commitments one can accrue. This helps alleviate some of the issues concerning infinite belief-formation and logical omniscience.

Moreover, the semantic commitment view alleviates some worries associated with trivial commitments. Specifically, the view can handle commitments to believe *redundant* things.[144] On

[143] The semantic commitment theory is similar to Bilgrami's view, in that it allows for a wider class of implication relations (see section 1).

[144] To be clear, the semantic commitment view might not fully solve the worry regarding infinite trivial commitments. However, commitments to believe redundant things is a problem for standard views of DC. The semantic commitment view alleviates this worry. Moreover, the semantic commitment view could deal with other cases of trivial commitments. For instance, one's belief that *p* might entail a trivial belief that *q*. However, one might not be in the cognitive position to cognitively endorse *q*. Hence, one would not be normatively constrained to believe *q*.

Bilgrami's account, an agent's belief that *p* logically entails *p*, and thus commits her to believing *p* (Bilgrami 2006, p. 373 fn. 8). This, however, seems counterintuitive. Are diachronic doxastic commitments redundant in this way? On my view, one does not accrue a semantic commitment to believe something that one already believes.[145] A semantic commitment is understood to be a commitment to bring about a *new* state-of-affairs, namely, a new or revised system of beliefs. To have a semantic commitment is to be normatively constrained to bring oneself to settle matters concerning the truth and plausibility of a further proposition not currently believed. One cannot bring about a new state-of-affairs in which one settles matters on the truth of a proposition, *p*, when one has already determined whether *p* is true. The semantic commitment theory rules out these redundant commitments.

Third, the semantic commitment theory avoids the related challenges associated with objectivity, cognitive inabilities, and variations in understanding (section 3.3). The ability to form beliefs is dependent on having certain cognitive abilities, reasoning capacities, and dispositions. Moreover, an agent's cognitive position is partly constituted by her stock of concepts. Indeed, an agent's conceptual repertoire will partly dictate which propositions she can comprehend and come to believe. Hence, one can be committed to believing *p* only when one has the cognitive resources to do so. The semantic commitment view accommodates these cognitive variations and maintains that one is only committed to believe those things that one is in a cognitive position to endorse. This accords

[145] To be clear, Bilgrami has a different conception of doxastic commitments. On his view, a "commitment…is something that is possessed within a point of view", such that it makes sense for one to say or think that one should believe *p* when they already possess this belief (Bilgrami 2006, p. 373 fn. 8). The semantic commitment view is different. This view posits commitments to mentally settle matters on new propositions, form new beliefs, and thus revise one's current system of beliefs. Moreover, I take it that the type of normative force that Bilgrami discusses in the cited remarks above is captured (at least partly) by what I refer to as "cognitive commitments" associated with belief. Specifically, to believe *p* would engender the cognitive commitments of *Doxastic Veridicality* and *Doxastic Integrity*. Consequently, one ought to regard *p* as being true and continue to believe *p* in some doxastic form until one revises one's judgment concerning the truth of *p*.

with the intuition that a doxastic-formation commitment to believe *p* is not some objective normative constraint that applies to any agent that possesses any set of beliefs that implies *p*. Rather, our doxastic-formation commitments are subjective and dependent on individual cognitive abilities.

Finally, the semantic commitment view incorporates the plausible idea that our doxastic-formation commitments have multiple satisfaction conditions (section 3.4). One can fully satisfy a semantic commitment by forming the appropriate belief, or one can partially satisfy this commitment by forming an attitude of acceptance. An agent is not fully irrational or categorically in error if she merely accepts a proposition, *p*, that she has a commitment to believe or endorse. For if she correctly recognizes that *p* follows from her beliefs and thus comes to accept *p*, then she is in a better normative position than those states of suspending judgment on *p* or believing not-*p*. Satisfying our doxastic-formation commitments can come in degrees.

I have argued that the semantic commitment view avoids or ameliorates many of the challenges facing standard accounts of doxastic-formation commitments. If the semantic commitment view can avoid or ameliorate these worries, then it is a more tenable and plausible theory of doxastic-formation commitments. I conclude that the semantic commitment view is a more tenable theory.

5. Semantic Commitments and Acceptance and Supposition

One question we might ask is whether our attitudes of acceptance and supposition generate similar semantic commitments. Before concluding this chapter, I briefly consider this question and argue that semantic commitments are restricted to beliefs.

Let us first consider acceptance. Accepting *p* need not involve any commitment to accept those propositions that follows from *p*. One might accept *p* and take on a commitment to employ *p* as a premise in reasoning. In this case, one might be disposed to accept certain propositions that

follow from *p*.[146] However, accepting *p* need not *always* involve a disposition to use *p* as a premise in an explicit line of reasoning. Acceptance is constituted by *some* mental resolution to utilize a proposition in thought or reasoning, and to take a positive mental stance toward that proposition. Moreover, accepting *p* is compatible with being agnostic toward *p*'s truth. One might accept *p* for pragmatic reasons. Consequently, one has a *pragmatic acceptance* (see chapter 1). In this case, one might appropriately resist accepting some of the propositions that follow from *p*. For example, I might pragmatically accept *that my neighbors are not stealing my mail*. I might recognize that I lack sufficient epistemic reasons supporting the truth of this proposition. However, I accept it primarily because of certain practical reasons I have in a context. For instance, accepting that proposition allows me to satisfy my goal of having pleasant interactions with my neighbors. But even if I accept this proposition, I need not be committed to accept or believe the propositions that logically follow from it.

Furthermore, it is possible for an agent, *S*, to pragmatically accept a contradiction, (*p* & not-*p*).[147] If *S* also believes or accepts *ex contradictione quodlibet*, then *S* will recognize and accept that (*p* & not-*p*) entails any proposition.[148] However, it does not seem correct to say that *S* would have an infinite number of semantic commitments to accept an infinite number of propositions. Acceptance is rationally based on reasons (evidential and pragmatic) that one has in a given context. Hence, *S* ought to only accept those propositions that are rational for her to accept. We therefore lack any good reason to accept the claim that our attitudes of acceptance will always generate semantic commitments.[149]

---

[146] Note that Cohen claims that acceptance involves a "policy" of employing a proposition in reasoning (see Cohen 1992, p. 4). My view of acceptance is different. Policies of premising propositions are a type of acceptance. But a policy of premising *p* is not necessary for accepting *p* (see chapter 1).

[147] It is important to keep in mind that one would be unable to rationally accept contradictions for epistemic reasons. An agent's rational acceptance of (*p* & not-*p*) would be a *pragmatic* acceptance.

[148] See Pietz and Rivieccido (2013) for a brief discussion of this principle in classical logic.

[149] There are many cases we can construct to demonstrate this point. For instance, an agent might lack sufficient *epistemic* reason to believe a religious proposition. Nevertheless, she could accept that proposition for pragmatic reasons. Perhaps accepting it facilitates some practical goal that she has.

Similar considerations apply to suppositions. Suppositions are a distinct class of cognitive endorsements (see chapter 1). One feature they share with acceptances is that to suppose *p* need not require one to consider *p* as being true. Our suppositions do not engender any semantic commitments for the same reasons that acceptances do not generate these commitments. Moreover, suppositions do not involve a commitment to utilize a proposition in thought. Nor is one disposed to take any positive mental stance toward the target proposition of a supposition. Indeed, our suppositions can be brief imaginings and weak mental postulations of propositions. Given the tentative nature of these states, it is not plausible to insist that our suppositions generate commitments to suppose the propositions that follow from our suppositions. Finally, it should be noted that one can suppose a proposition, *p*, while recognizing that one lacks good reason to think that *p* is plausible or useful. In these situations, it is not appropriate to claim that one's supposition that *p* normatively compels one to suppose or cognitively endorse the propositions that follow from *p*.

I contend that acceptances and suppositions do not engender semantic commitments. Semantic commitments are therefore restricted to an agent's beliefs and her subjective cognitive position.

6. Conclusion

I have advanced an account of semantic commitments. Semantic commitments are normative constraints governing belief-formation. Unlike standard accounts, this view focuses on the subjective cognitive positions that agents possess. If an agent's set of beliefs imply some proposition, *p*, via some correct inference rule or norm, *R*, and *S* is in a cognitive position to grasp *p* and competently use *R* to believe and settle matters concerning *p*'s truth, then *S* is committed to believe *p* when *S* is in a position to do so. *S* can fully satisfy this commitment by believing *p*, or she can partially satisfy this commitment by accepting *p*. This account, I contend, provides a plausible theory of our normative commitments

to cognitively endorse certain propositions that follow from our beliefs. I further argued that this view avoids or ameliorates several challenges and worries facing standard views of doxastic-formation commitments. Hence, the semantic commitment theory is a more tenable and preferable account.

Several important consequences follow from this account. First, accommodating semantic commitments helps to explicate the normative structure of our thinking. Semantic commitments have a distinct kind of normative force. Unlike cognitive commitments that normatively govern how one ought to treat a proposition in thought, our semantic commitments normatively constrain us in how we should form new cognitive endorsements over time. Second, semantic commitments have multiple satisfaction commitments. One can fully or partially satisfy the content of these commitments. Consequently, these normative constraints differ from other kinds of requirements, obligations, and commitments. For example, it does not seem right to say that one could partially satisfy a moral obligation or requirement to perform some act. Either one satisfies this obligation, or they do not. Semantic commitments therefore present themselves as a unique kind of normative constraint that governs our thinking.

IV. Doxastic Normative Tension and the Primacy of Cognitive Commitments

Adam C. Sanders

0. Introduction

Semantic commitments contribute to diachronic epistemic rationality. If an agent's beliefs and cognitive position commit her to believing a further proposition, *p*, then, other things being equal, she ought to satisfy her commitment and form the belief that *p*. Semantic commitments therefore play a role in determining how we ought to think and reason. One problem with this account is that it allows for a certain type of normative tension. Specifically, it is possible for there to be cases involving conflicting commitments to believe inconsistent things. This issue also arises for most (if not all) standard accounts of doxastic-formation commitments.[150] This is problematic. For if satisfying semantic commitments is an important aspect of epistemic rationality, then we need a way to reconcile or eliminate conflicting and problematic semantic commitments.

In this chapter I aim to address and resolve the issue of conflicting semantic commitments. I further demonstrate how our semantic commitments normatively interact with our cognitive commitments to promote rational belief-revision. I establish several claims. First, I provide a skeptical argument against a view I refer to as the "Different Weights" solution. Let us call this view "DW". According to DW, the following conjunction is true: It is the case that both semantic commitments are like reasons in the sense that they have different known weights or degrees of strength, and we can appeal to only the weight of these commitments to determine how one ought to resolve inconsistent semantic commitments. I argue, however, that we lack sufficient reason to endorse this conjunction. Consequently, we lack sufficient reason to endorse DW. Second, I advance a normative procedure maintaining that we should appeal to our cognitive commitments to reconcile inconsistent

[150] A rough picture of the standard accounts is presented in the previous chapter. For a more detailed discussion of these views see Levi (1997), Millar (2004), Bilgrami (2006), Shpall (2014), and Liberman and Schroeder (2016).

and problematic semantic commitments. Finally, I argue for what I refer to as the "Primacy Thesis", which states that our cognitive commitments are normatively fundamental or prior to semantic commitments, and thus they play a more important role in governing our epistemic lives.

The arguments advanced here establish several important claims. Taken together, these claims provide a plausible account of the normativity of doxastic reasoning and belief-formation. The consequences of this account are both important and interesting. First, it provides us with a deeper understanding of the unique normative features of our semantic commitments. Semantic commitments can be distinguished from other normative elements. Moreover, this view explicates how semantic commitments can be guided by our more fundamental cognitive commitments. This sheds further light on how semantic commitments function in relation to other normative elements. The picture that emerges from this discussion is that rational thinking and belief-formation is normatively governed by a complex hierarchy of commitments.

The structure of this chapter is as follows. In section 1 I discuss cases of doxastic normative tension. The specific cases I focus on involve agents with one or more conflicting semantic commitments to believe inconsistent things. I further discuss DW and its proposed method for resolving conflicting semantic commitments. In section 2 I provide an inductive argument for the claim that we lack sufficient reason to accept the view that semantic commitments have different degrees of known normative weight. I further argue that we lack sufficient reason to accept the claim that we can appeal to only the normative weights of semantic commitments to determine how we should revise inconsistent semantic commitments. It follows that we lack sufficient reason to endorse DW. Section 3 advances a normative procedure for dealing with inconsistent semantic commitments. On this view, alleviating the normative tension stemming from conflicting or inconsistent semantic commitments is guided by appeals to other properties of beliefs, namely, the associated rational support and cognitive commitments. In section 4 I argue for the Priority Thesis stating that cognitive

commitments are more fundamental than semantic commitments. I conclude the chapter in section 5 by discussing the hierarchy of commitments governing rational thinking and belief-formation.

1. DW AND THE PROBLEM OF DOXASTIC-FORMATION NORMATIVE TENSION

If an agent, *S*, has a set of beliefs that imply a proposition, *p*, via some rule or norm of inference, *R*, and *S* has the conceptual resources, understanding, and cognitive standing to grasp *p* and believe *p* by competently using *R*, then *S* has a semantic commitment to believe *p*. Semantic commitments engender normative constraints to revise and update one's belief-set through doxastic reasoning and belief-formation. Other things being equal, one ought to satisfy the content of one's semantic commitment to believe and cognitively endorse *p* when one is in the position to settle matters concerning *p*'s truth.[151]

Now, it is typically agreed that rational agents ought to have consistent or non-contradictory beliefs (Quine and Ullian 1970, pp. 16-17; Broome 2013, pp. 154-157). Consider a case where an agent, *S*, at time, *t*, believes several propositions, $p_1, p_2, p_3, \ldots, p_n$. Call this set of beliefs, "*Γ*". *S*'s belief-set, *Γ*, is inconsistent at *t* if it is not the case that $p_1, p_2, p_3, \ldots, p_n$ can simultaneously be true at *t*. In this case, *S* is not doxastically right. *S* ought to revise her beliefs and restore consistency. Consequently, if rationality is partly governed by our semantic commitments, then this normative constraint should steer agents away from forming inconsistent belief-sets.

The problem, however, is that the semantic commitment account allows for conflicting normative constraints to believe inconsistent things. Moreover, this problem extends to standard

[151] The semantic commitment account was defended in chapter 3. A semantic commitment to believe *p* puts a normative constraint on its possessor to bring about a new state-of-affairs, namely, a belief that *p* (one could *partially* satisfy the commitment by accepting *p*). Other things being equal, one ought to satisfy this commitment. This normative force is a general feature of all commitments. See, for example, Margaret Gilbert's work on personal commitments of the will stemming from intentions and decisions (Gilbert 2013, 2014).

accounts of doxastic-formation commitments (see Liberman and Schroeder 2016, p. 116). Consider a case where an agent, *S*, believes (B1) *p* and she believes (B2) *q*. Suppose that *S* competently employs a process of reasoning and comes to recognize and believe that *p* entails *r*. Moreover, during this time she also comes to recognize and believe that *q* entails not-*r*. If *S* continues to hold (B1) and (B2), then she acquires two semantic commitments to believe both (B3) *r* and (B4) not-*r*. This is unacceptable. For *S* ought not to believe contradictory things. She ought to reconcile this normative conflict. The problem, however, is that it is not clear how an agent should effectively eliminate this kind of normative tension.

If our semantic commitments contribute to diachronic epistemic rationality, then there needs to be some acceptable procedure to deal with these cases of normative conflict. One possible proposal has been discussed by Liberman and Schroeder. They note that some theorists might hold that our doxastic-formation commitments function like reasons in the sense that they can weigh against one another to determine what one ought to do (Liberman and Schroeder 2016, p. 116). Following Lord and Maguire, we can say that one's reasons for performing Φ are "decisive" when they outweigh any competing reasons to not do Φ (Lord and Maguire 2016, p. 10).[152] Having a decisive reason to Φ will normatively require one to bring about Φ (Lord and Maguire 2016, p. 9). Perhaps doxastic-formation commitments (including semantic commitments) function in the same way. When one's beliefs and cognitive position engenders multiple commitments to believe inconsistent things, then the stronger or weightier commitment is decisive, and thus normatively determines which propositions one ought to believe or revise.

The strategy I want to first consider insists that semantic commitments have different degrees of weight and that we need to *only* look to the more weighty commitments to dictate how one ought

[152] Many theorists hold that reasons can have different degrees of weight or strength (see, for example, Raz 1975, p. 25; and Liberman and Schroeder 2016, pp. 104-120).

to form or revise their beliefs. This is the DW view. The following figure illustrates the strategy employed by DW theorists.

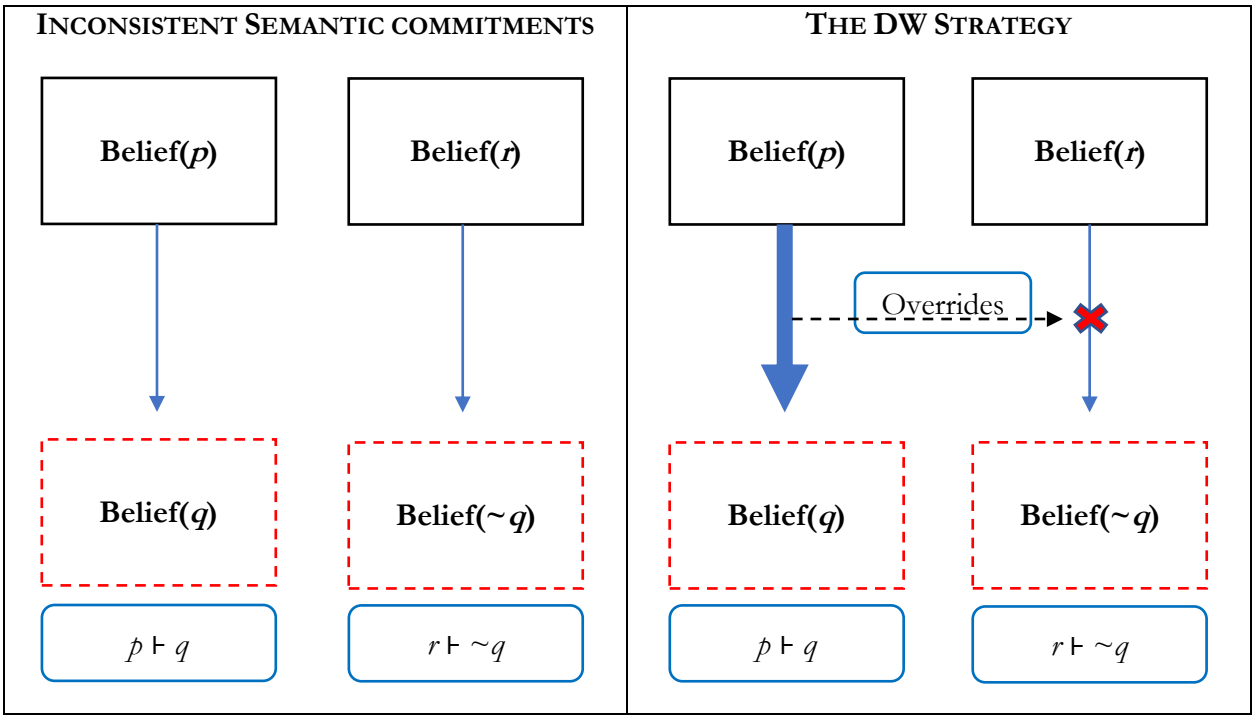

Figure 2: The Different Weights Solution

Some clarifications are in order. Let "⊢" specify a logical relation (mind-independent) of deducibility, such that, "$p \vdash q$" means "q is deducible from p". Let boxes with solid outlines represent beliefs that an individual agent holds at a given time. Let boxes with dashed (broken) outlines represent possible beliefs the individual agent could form at some future time. Finally, let the arrows indicate semantic commitments generated by currently held beliefs. The thicker the arrow, the stronger the normative force attached to the semantic commitment.

The left box of *Figure 2* represents a case where one has multiple beliefs that engender multiple semantic commitments to believe inconsistent things. The right box of *Figure 2* represents the DW strategy. According to the DW strategy, our semantic commitments have different normative weights. In cases where there are multiple semantic commitments to believe inconsistent things, some of those semantic commitments will have more normative strength that outweigh the other semantic

commitments. Stronger semantic commitments override weaker ones. Accordingly, we need to *only look to the weight of the semantic commitment*. That is to say, we do not need to appeal to the other properties of our beliefs (e.g., rational support, true contents, etc.) to determine how one ought to handle conflicting semantic commitments.

The DW strategy depends on the truth of the following conjunction:

> **(N)** It is the case that both (i) semantic commitments have varying degrees of normative weight that can be known, and (ii) we can appeal to only the normative weightiness of the known corresponding semantic commitments to determine how one should resolve or revise inconsistent and conflicting semantic commitments.

My contention, however, is that the DW view (and any view that endorses the conjunctive thesis, N) is problematic. In the next section I show that we have sufficient reason to doubt or reject N. Hence, we have sufficient reason to doubt DW. I further argue that the process of determining which semantic commitments to satisfy or revise is guided by some other normative standard or procedure.

2. A Skeptical Argument Against the Different Weights Solution

My strategy in this section is to cast doubt on the plausibility of DW (i.e., the Different Weights solution). This is not to say that I will argue that DW is false. Rather, I am merely arguing for the skeptical position stating that we lack sufficient reason to endorse N. The cogency of DW depends on the cogency of the conjunctive thesis, N. Consequently, we lack sufficient reason to endorse DW. My argument proceeds by demonstrating that we have sufficient reason to doubt one or more of the conjuncts of N. I discuss each conjunct in turn.

2.1 Doubting The First Conjunct of N

If semantic commitments have different degrees of normative strength, then there needs to be some fact that explains how these commitments acquire these different weights. In this section I discuss

several possible views that purport to explain how our semantic commitments can acquire different normative weights. None of these views are conclusively established. Moreover, I argue that we cannot always know whether or not a given belief (or set of beliefs) engenders a stronger semantic commitment than the semantic commitment generated by another belief (or set of beliefs). I therefore conclude that we have good reason to be skeptical of the claim that semantic commitments have different known normative weights.

Consider first the fact that some beliefs are more justified than others. One might think, then, that strongly justified beliefs will generate stronger semantic commitments than weakly justified beliefs. In situations involving inconsistent semantic commitments, one need only look to the weightier semantic commitments that stem from beliefs with strong rational support. These semantic commitments will outweigh (and sometimes override) the semantic commitments stemming from weakly justified beliefs.

I am sympathetic to the idea that one's available reasons and epistemic support will play a role in determining how one ought to revise inconsistent semantic commitments. Indeed, I will appeal to a similar strategy in the next section. However, it is not at all clear that the degree of rational support for a belief would causally produce a corresponding semantic commitment with a higher degree of normative force. Having a high degree of evidence for one's belief that *p* might dispose one to feel very confident that one's belief is true. A high degree of rational support can also enable one to provide more reasons for holding the target belief. This can also dispose one to maintain that belief. But it does not follow from this that the normative force of a semantic commitment associated with a given belief would be much weightier than the normative force generated by a belief that is justified on slightly weaker evidence. For it is equally plausible that semantic commitments have the same normative strength, and any perceived pressure to satisfy one semantic over another is merely due to (i) a recognition that one belief (or set of beliefs) has more rational support, and (ii) having certain

dispositions to maintain and form beliefs that are sufficiently justified. There needs to be a further reason suggesting that rational support causally affects a semantic commitments normative strength. Second, an agent may not always be in a situation to know which one of her beliefs generate stronger or weaker semantic commitments. Consider a case where an agent, *S*, has two beliefs that generate two distinct semantic commitments to believe inconsistent propositions. Suppose that *S*'s grounding beliefs both have a high degree of rational support. In this situation, *S*, may not know which one of her semantic commitments is weightier. At this point, it is equally plausible that the rational support of our beliefs, and not the semantic commitments, provide the best explanation for the normative force and guidance needed for determining which semantic commitments to satisfy or revise. More work is needed to conclusively show how agents can know or justifiably determine the normative weight of a semantic commitment that causally depends on the corresponding degree of justification.

Another possible view might hold that correct beliefs (i.e., beliefs with true contents) generate semantic commitments that are weightier than the semantic commitments stemming from incorrect beliefs (i.e., beliefs with false contents).[153] According to this strategy, when one has inconsistent semantic commitments, one need to only look to the semantic commitments generated by correct beliefs. The semantic commitments of correct beliefs will override the weaker semantic commitments stemming from incorrect beliefs.

This strategy is not convincing. Beliefs are correct only derivatively when their target propositions are true. And it is not obvious how the truth of a proposition would causally affect the normative weight of a belief's semantic commitment. Without any supporting reasons, we are warranted to take a skeptical stance toward the claim that the truth of a belief's propositional content will causally affect the normative strength of the corresponding semantic commitments. Moreover,

[153] It is typically thought that our beliefs are evaluated according to a so-called "Truth norm" claiming that beliefs are correct when they have true contents (see Wedgwood (2002) p. 272, Lynch (2009), p. 79, and Whiting (2013) p. 122). I discuss this norm later in the chapter.

when we believe propositions, we consider the target propositions to be true. Considering that we are fallible creatures, we can (and typically do) mistakenly believe false propositions. But when we contemplate or reason with beliefs that have false propositional contents, then we will posit those contents as being true in the same way that we would posit beliefs that actually have true contents. In these cases, it is not clear (at least from a first-personal view) that correct beliefs would engender semantic commitments that are more weighty than semantic commitments engendered by incorrect beliefs. Consequently, this strategy does not successfully establish the claim that our semantic commitments always have different degrees of normative strength. Nor does this strategy conclusively show that an agent could reliably or justifiably determine which one of her beliefs have false contents. It is therefore doubtful that one could always know which one of her beliefs have stronger or weaker semantic commitments.

A third view might appeal to some version of Quine's idea that our beliefs are interconnected into a "system" of beliefs, such that some beliefs at the center of this system (e.g., beliefs concerning logic and physics) are more firmly rooted in the belief-system than those beliefs at the system's edges (e.g., perceptual beliefs) (Quine 2004, pp. 50-51). One might advance a rough version of this Quinean picture and argue that those beliefs firmly rooted in the center of a belief-system engender weightier semantic commitments than those semantic commitments generated by beliefs at the system's edge.[154]

This, however, is not a promising strategy. First, it is not clear that the position of a belief in this system would causally determine a varying degree of strength of the corresponding semantic commitment. At most, a more firmly rooted belief might dispose one to be less inclined to abandon that belief. But this is not the same thing as having a much stronger semantic commitment to maintain

[154] This is a very rough picture of Quine's view. I am not providing a complete exegesis of Quine's view. I am only interested in views that might appeal to the general idea that entrenched beliefs (where *entrenchment* can be understood in different ways) will generate different normative commitments with different known normative weights.

that belief and cognitively endorse the propositions that are implied by it. More work is needed to show that entrenched beliefs causally affect the weights of semantic commitments. Second, it is not clear from the first-personal view that one would have sufficient reason to think that his or her entrenched beliefs will always generate stronger semantic commitments than non-entrenched beliefs. Nor is it clear how one could know how much normative force a semantic commitment from one entrenched belief has compared to the normative force of a different semantic commitment generated by a similarly entrenched belief. Hence, adopting this Quinean strategy does not conclusively show that one can *know* that one's firmly entrenched beliefs will *always* engender corresponding semantic commitments with stronger normative weights. We can plausibly doubt the first conjunct of N.

A final view I will consider here centers on the way beliefs are formed and how they justify other beliefs. According to Lehrer, we can draw a distinction between "basic beliefs" and "non-basic beliefs" that operate in foundationalist theories of justification (Lehrer 1974, p. 76). He claims that a belief is *basic* if (1) it is "self-justified" and thus does not derive its rational support from "any non-basic belief"; (2) it can "only be refutable by other basic beliefs"; and (3) it functions as the primary way to refute or justify all other beliefs (Lehrer 1974, p. 76). Basic beliefs, then, concern "basic facts" and observational (or perceptual) propositions (Lehrer 1974, pp. 76-77). On the other hand, a *non-basic* belief is simply a belief that is justified by some other belief or basic belief (Lehrer 1974, p. 76). This seems like a plausible distinction. One might argue that basic beliefs are foundational and therefore generate weightier semantic commitments than those commitments stemming from non-basic beliefs. In cases of inconsistent semantic commitments, we need to only look to the weight of the semantic commitments stemming from basic and non-basic beliefs.

This strategy, however, faces several challenges. Even if we endorse the distinction between basic and non-basic beliefs, it is not clear how the property of being a basic belief plays any causal role in determining the strength of its corresponding semantic commitment. How is it the case that a belief

that *p* has a weaker semantic commitment simply because it is justified by an antecedent basic belief? Without any argument to suggest otherwise, it is implausible to hold that a distinction between basic and non-basic beliefs will correspond to a distinction between weighty and non-weighty semantic commitments. Moreover, an agent may not always know or recall whether one or more of her beliefs are basic or not. In these cases, it is not clear that she will be able to justifiably determine which one of her beliefs generate the strongest (or weakest) semantic commitment. This strategy, then, does not establish the first conjunct of N.

I have considered several strategies that purport to show how there can be semantic commitments with different degrees of normative strength or weight. I argued that none of these strategies are sufficiently warranted. I further argued that these strategies do not successfully demonstrate that we can always know the normative weight of our semantic commitments. This is not to say that semantic commitments *cannot* have different normative weights. Indeed, there might be other strategies that I have not considered that can plausibly show that (i) semantic commitments have different normative weights, and (ii) we can know these different weights. Nevertheless, I have considered what I take to be the best strategies in defense of the first conjunct of N. The line of reasoning thus far provides a cogent inductive argument for the claim that we can rationally doubt the first conjunct of the thesis, N. If this is true, then we need not endorse N.

2.2 Doubting the Second Conjunct of N

If the arguments in 2.1 were not convincing, then I contend that we also have reason to doubt the second conjunct of N. Recall the second conjunct of N:

> … we can appeal to only the normative weightiness of the corresponding semantic commitments [of one's beliefs] to determine how one should resolve or revise inconsistent and conflicting semantic commitments.

I submit that we have good reason to think that this conjunct can be doubted. Let me explain.

First, it has already been shown in section 2.1 that it is unclear how the semantic commitments associated with beliefs (or sets of beliefs) could be attributed with different degrees of normative strength. But even if our beliefs do generate semantic commitments with different normative strengths, then it is not clear that this would *always* be a good guide for epistemic rationality.

Consider, for instance, entrenched beliefs. Suppose that our firmly entrenched beliefs engender strong semantic commitments that normatively outweigh the commitments stemming from non-entrenched or tentative beliefs (e.g., tentative and ordinary pre-reflective observational beliefs). However, it is possible for one to have irrational entrenched beliefs that routinely cause one to form bad, improper, or irrational beliefs. In these cases, it seems wrong to insist that the semantic commitments from the irrational entrenched beliefs are a good guide for subsequent rational belief-formation.

Similar considerations apply to the view that basic beliefs generate semantic commitments that are normatively stronger than the semantic commitments stemming from non-basic beliefs. Suppose that this is true. However, one might form an irrational basic belief. As a consequence of holding this irrational basic belief, one might be committed to form further irrational non-basic beliefs. It is incorrect to insist that the semantic commitment of this irrationally held basic belief would be a reliable guide for rationality and appropriate belief-formation. This point can be generalized to all basic beliefs: abiding by the semantic commitments generated by basic beliefs might not *always* be a good guide for epistemic rationality.

Perhaps semantic commitments engendered by true beliefs offer the best guide for rationally alleviating conflicting commitments to believe inconsistent things. If we are primarily guided by truth in these cases, and hence normatively compelled to believe only

those things that logically follow from our true beliefs, then we will be normatively guided to only hold consistent sets of beliefs. I am sympathetic to this view. However, we fall back into the problems discussed in subsection 2.1. One might have some false beliefs that, in conjunction with one's other beliefs, engender inconsistent semantic commitments. However, given the nature of belief (see chapter 1), one will regard the propositional contents of these beliefs (correct or incorrect) as being *true*. From the first-personal view, one will not always be capable of determining which belief is incorrect. Hence, abiding by the semantic commitments engendered by correct beliefs (i.e., beliefs with true contents) is not always the best guide for epistemic rationality.

Finally, we might look to the rational basis of our beliefs. Suppose that beliefs with more epistemic support generate stronger semantic commitments than beliefs with weaker rational support. Rationally held beliefs are those that are sufficiently warranted by epistemic reasons. Having sufficient epistemic reasons for a belief that *p* is a good guide for the truth of *p* (see chapter 1). Accordingly, it might be the case that semantic commitments engendered by rationally held beliefs are the best guide for dealing with inconsistent semantic commitments. That is to say, in cases where an agent has beliefs that semantically commit her to believe inconsistent or contradictory things, she ought to only believe the things that follow from her beliefs that are rationally supported to some significant degree. The semantic commitments engendered by beliefs with a high degree of rational support will outweigh or override the semantic commitments stemming from beliefs with a lower degree of rational support.

I think that rational evaluation is a good guide to belief-formation. However, the above proposal faces some challenges. First, one might have two beliefs that entail contradictory propositions. These two beliefs might have the same degree of rational support. Consequently, merely appealing to the *weight* of semantic commitments will not be a good guide for dealing

with the inconsistent commitments. Moreover, this view is faced with the challenge of redundancy. Why should we look at *only* the semantic commitments based on rational beliefs with a high degree of epistemic support when we can simply look to the epistemic support itself? I am thus skeptical that we should only look to the weight of semantic commitments to determine how we ought to resolve conflicting commitments to believe inconsistent things.

At this point we have reason to doubt one or more conjuncts of the above thesis, N. Furthermore, we can justifiably doubt that we need to only appeal to semantic commitments (and perhaps their weights) to appropriately resolve doxastic normative tension to believe inconsistent things. There may be no need to posit semantic commitments with different degrees of normative weight when we already have established norms and cognitive commitments at our disposal. Perhaps we can appeal to these further normative constraints for rational guidance. If this is the case, then we have sufficient reason to be skeptical of the cogency of N. Since DW depends on the truth of N, we lack sufficient reason to endorse DW.

3. A Normative Procedure For Alleviating Problematic Semantic Commitments

I do not doubt that we can plausibly distinguish basic beliefs from non-basic beliefs. I also accept that some of our beliefs are more entrenched than others. Moreover, it is obvious that we sometimes have false beliefs. Furthermore, our beliefs can have different degrees of rational support. However, my arguments in section 2 purport to show that we lack sufficient reason to rationally believe or endorse the following conjunction, N: It is the case that both (i) semantic commitments have varying degrees of normative weight that can be known, and (ii) we can appeal to only the normative weightiness of the known corresponding semantic commitments to determine how one should resolve or revise inconsistent and conflicting semantic commitments.

If my arguments in section 2 are convincing, then we are still left with the puzzle of how to deal with conflicting and problematic semantic commitments. Liberman and Schroeder propose that, when dealing with conflicting commitments that do not weigh against each other, we might construct a procedure that appeals to other requirements and obligations that tells us which commitment to give up or revise (Liberman and Schroeder 2016, p. 116). Liberman and Schroeder do not argue for a specific procedure. In this section I advance and defend a plausible normative procedure for dealing with incompatible and conflicting semantic commitments. According to this operation, we do not simply look to the semantic commitments *themselves*. Rather, we must exercise rational agency and undergo an *activity* or *process* of rational evaluation. On this view, we ought to abide by other principles and norms governing the fundamental features of beliefs to resolve conflicting semantic commitments. Specifically, we should appeal to principles and norms governing the consistency and rational support of beliefs. But this normative procedure, I contend, is guided by *cognitive commitments*.

There are many different norms and rules of rationality governing some specific feature of beliefs. I want to begin by first arguing that appealing to a *single* rule of rationality or norm of correctness will not be sufficient to deliver a complete and adequate procedure for dealing with all cases of conflicting semantic commitments. To illustrate why, consider a simple case involving an agent with two inconsistent semantic commitments to believe *p* and believe not-*p*. We could try to reconcile this normative tension in several ways. First, we could simply look to some external rule or standard of rationality that agents ought to comply with.[155] One of these rules claims that we are rationally required to refrain from having contradictory beliefs at a given time. Hence, we might think

[155] These rules are external in the sense that they do not arise from our subjective mental states (see Shpall 2014, pp. 154-155). Rather, our mental states should comply with these rules. This is a defining feature of rational requirements. Contrary to this, cognitive commitments and semantic commitments depend on our own mental states. I will say more about rational requirements in the next chapter.

that by complying with this rule, agents will be led away from satisfying their inconsistent semantic commitments to believe contradictory things.[156]

I agree that we should refrain from having contradictory or inconsistent beliefs. The problem, however, is that appealing to a rational rule of non-contradiction will not help us determine how one ought to rationally handle conflicting semantic commitments. At most, this rational standard tells us that agents are rationally forbidden to believe both *p* and not-*p* at a given time. This is irrelevant in determining which inconsistent semantic commitments an agent ought to revise, and how she ought to change her cognitive positions in cases of doxastic normative tension.

A second possible strategy utilizes some standard or requirement of rational support. Many philosophers correctly think that a belief that *p* is rational when it is based on good epistemic reasons supporting the truth of *p* (see Pritchard 2016, p. 90). Accordingly, some theorists of epistemic rationality proceed to specify norms governing the rational basis of beliefs.[157] For example, Richard Feldman articulates the following principle:

> **Rationality Norm:** For any person S, time t, and proposition p, if S has any doxastic attitude at all toward p at t and S's evidence at t supports p, then S epistemically ought to have the attitude toward p supported by S's evidence at t (Feldman 2000, p. 679).[158]

It might be argued that compliance with this principle will help agents reconcile inconsistent semantic commitments. Having inconsistent semantic commitments indicates that one lacks sufficient evidence to believe the inconsistent propositions entailed by one's beliefs. Moreover, one has evidence suggesting that the grounding beliefs that give rise to the inconsistent semantic commitments are

[156] Several theorists claim that our mental states are governed by some rule prohibiting contradictory beliefs. See, for example, Way (2010) p. 1059 and Broome (2013) pp. 154-155. I will say more about this requirement in the next chapter. There I will argue that rationality can be constituted by both requirements and commitments.
[157] For some discussions about the requirements and norms governing the rational basis of beliefs see Feldman (2000) and Brunero (2010).
[158] I should note that Feldman does not refer to this as the "Rationality Norm".

problematic and ought to be revised. According to the Rationality Norm, one ought to drop the beliefs that lack sufficient support. Revising one's system of beliefs would eliminate the normative tension.

I am sympathetic to this approach. But simply appealing to the Rationality Norm might fail to provide a *complete* normative procedure that effectively handles all cases involving inconsistent semantic commitments. As noted in the previous section, it is possible for an agent, *S*, to have good evidence and rational support for a set of beliefs that imply inconsistent or contradictory propositions. Appealing to the Rationality Norm prescribes that an agent should only believe propositions that have sufficient evidential support. Hence, *S* should not believe the implied inconsistency. Moreover, she should abandon or revise her grounding beliefs that give rise to the normative tension.[159] However, considering that her grounding beliefs are well-supported, she may not be in a position to correctly determine which one of her grounding beliefs should be revised. Consequently, she will not be in a position to determine which semantic commitments to give up. Moreover, she may not want to suspend judgment on all of the target propositions of her belief-set. Indeed, we want to reject the problematic beliefs while retaining the correct and rationally supported beliefs. What we need is a complete normative operation that guides agents in processes of rational evaluation and belief-revision. This includes a straightforward method of determining which beliefs to retain when one is rationally evaluating one's system of beliefs. My ultimate aim in this section is to advance a complete and effective normative operation that agents can follow in cases of doxastic normative tension.

A final strategy I will consider here involves appealing to the so-called "Truth norm". Many versions of this norm claim (roughly) that a given belief at time, *t*, is correct or appropriate if and only if the propositional content of that belief is true at *t*.[160] This is a norm that we can use to evaluate our

[159] Abandoning or revising the grounding beliefs could mean that one merely suspends judgment on the target propositions.

[160] Many theorists adopt a version of the Truth norm. See, for example, Wedgwood (2002) p. 272, Lynch (2009), p. 79, and Whiting (2013) p. 122). To be clear, there are different variations of this norm. Daniel Whiting, for instance, expresses this norm as follows: "On may believe that *p* if and only

beliefs at a given time. Accordingly, we might say that we ought to give up beliefs that are false. In doing so, we will eliminate inconsistent semantic commitments to believe or cognitively endorse inconsistent things.

This approach faces several problems. From a first-personal perspective, an agent with inconsistent semantic commitments may not be in a position to correctly discern which one of her premise or grounding beliefs are false. Otherwise, she would not be in the predicament that she is currently in. Presumably she takes herself to have sufficient evidence to support her beliefs at the time when her thinking commits her to endorse inconsistent propositions. Appealing to the Truth norm provides little help given that it only dictates that a given belief is correct if and only if it is true. The norm is therefore silent as to *how* an agent should go about revising her set of inconsistent beliefs to eliminate inconsistent semantic commitments. What we need is a clear and complete normative procedure that tells us how agents ought to deal with conflicting semantic commitments.[161]

A promising strategy that I will explore makes use of a *plurality* of internal norms: our cognitive commitments. I have already argued for the view that belief engenders cognitive commitments that individuate and normatively constrain the attitude of belief itself. On the face of it, then, it seems plausible that appealing to these commitments will provide a satisfactory procedure for alleviating normative tension brought about by conflicting semantic commitments. Recall that our beliefs generate the following commitments[162]:

> **Doxastic Veridicality:** To believes *p* at time, *t*, is to be normatively committed to regard *p* as being true at *t*, such that one cannot simultaneously consider not-*p* to be true at *t*.

---

if it is true that *p*" (Whiting 2013, p. 122). My argument here does not depend on any specific formulation of this norm.

[161] One might insist that the Truth-norm generates a corresponding norm to revise false beliefs. This may be true. However, the point here is appealing to the Truth norm is not sufficient to adequately guide us in all cases of belief-revision. For this norm cannot determine which beliefs one should revise in cases where one has an inconsistent belief-set but is not sure which beliefs are problematic or false.

[162] The commitments expressed by my notions of *Doxastic Veridicality* and *Doxastic Rationality* are taken from Coliva's view of doxastic commitments (Coliva 2016, 2019). See also chapter 2.

> **Doxastic Rationality:** To believe *p* at time, *t*, is to normatively commit oneself to be disposed to take one's evidence as being sufficient to establish the truth of *p* at *t*, and to be sensitive to new evidence and information concerning the truth or falsity of *p*.
>
> **Doxastic Integrity:** To believe *p* at time, *t*, is to be normatively committed to continue to regard *p* as being true (either occurrently or dispositionally) in accordance with one's evidence, until one rationally revises one's judgment concerning the truth or falsity of *p* at some future time, t_n.

These cognitive commitments are constitutive aspects of a token attitude of belief. They normatively govern how one ought to mentally treat the target proposition of a belief. Contrary to this, a semantic commitment is a kind of diachronic normative constraint to form new beliefs and thus revise one's belief-set. Moreover, semantic commitments are generated by one's beliefs, concepts, and cognitive abilities. Both kinds of commitments can interact in such a way that promotes epistemic rationality. My contention here is that cognitive commitments provide a basis for developing a normative procedure to maintain rational semantic commitments while revising problematic ones.[163] In cases of doxastic normative tension, our cognitive commitments collectively determine how one ought to modify one's cognitive position. Appealing to a plurality of cognitive commitments provides a *comprehensive* and *rational* way to avoid inconsistent normative pressures to believe inconsistent things.

The proposed normative procedure proceeds in steps. The first involves a mental act of suspending judgment or taking some non-belief stance. If, for example, one's set of beliefs at time, *t*, implies contradictory propositions, *p* and not-*p*, then one ought not to believe both *p* and not-*p* at *t*. Our beliefs aim at truth. This follows from the cognitive commitment of Doxastic Veridicality. If an agent believes *p*, then she is normatively constrained to regard *p* as being true, such that she cannot simultaneously regard both *p* and not-*p* as being true. Hence, one cannot consciously satisfy a semantic commitment to believe contradictory things at a given time. Doxastic Veridicality therefore guides

[163] The procedure I develop here might fit into what Liberman and Schroeder refer to as a "*Priority Picture*" (Liberman and Schroeder 2016, p. 113).

one to take some non-belief stance toward the contradictory propositions entailed by one's set of beliefs. At minimum, one is required to refrain from believing these contradictory propositions if one is in this conflicting cognitive position.[164] Similar considerations apply to cases involving semantic commitments to believe inconsistent things that are not explicit contradictions.

Of course, it could be argued that a non-belief stance is a violation of an agent's semantic commitment if that agent still maintains her initial premise or grounding beliefs that gave rise to her conflicted normative situation. In response to this, it is important to note that confronting any normative tension to believe inconsistent things indicates that one's system of beliefs and cognitive position is not quite right. For this cognitive position implies a new set of beliefs with inconsistent propositional contents. Considering that the target propositions of the members of an inconsistent set of beliefs cannot be true together, it follows that one should modify one's set of beliefs. The fact that one is in this cognitive position provides sufficient warrant to revise or suspend judgment on some of the grounding beliefs that generate the inconsistent semantic commitments. In doing so, one evades any inconsistent semantic commitments by abandoning the beliefs that partly engender them.

The normative force that facilitates the elimination or revision of one's grounding beliefs generating inconsistent semantic commitments is the cognitive commitment of Doxastic Rationality. This commitment binds an agent to believe in accordance with her available evidence and epistemic reasons that are sufficient to warrant belief. In cases involving semantic commitments to believe inconsistent things, one has objective evidence suggesting that the basing beliefs are inconsistent and cannot be true together. Hence, Doxastic Rationality guides one to make the appropriate doxastic revisions. Accordingly, one ought to abandon irrational grounding beliefs or suspend judgment on the propositional contents of those beliefs that generate inconsistent semantic commitments.

---

[164] There could be several ways to satisfy this requirement. Typically, one satisfies the requirement by suspending judgment. But the requirement could be satisfied without performing any mental act of suspending judgment. In these cases, one simply fails to take any mental stance.

This raises the question as to how we ought to determine which premise or grounding beliefs to revise. Hence, the final step of our normative procedure involves a Quinean approach to evidential or rational scrutiny. This move is also guided by Doxastic Rationality. Recall that rational beliefs are sensitive to good evidence and epistemic reasons. Once one drops or suspends judgment on her beliefs that play a part in generating inconsistent semantic commitments, one can proceed to evaluate the evidence one originally had for holding those beliefs. As Quine and Ullian put the point:

> …when a set of beliefs is inconsistent, at least one of the beliefs must be rejected as false; but a question may remain open as to which to reject. Evidence must then be assessed, with a view to rejecting the least firmly supported of the conflicting beliefs. But even that belief will have had some supporting evidence…so in rejecting it we may have to reject also some tenuous belief that had helped to support it. Revision may thus progress downward as the evidence thins out (Quine and Ullian 1970, pp. 16-17).

To determine which beliefs to keep, modify, or reject requires an evaluation of the rational basis of those beliefs. More precisely, our evaluation of beliefs is facilitated by a process of *rational scrutiny*. Processes of rational scrutiny are constituted by reflection, inquiry, and evidence gathering concerning the rational basis of propositions. I take rational scrutiny to be an intuitive way to rationally maintain one's belief-system. If an agent discovers that one or more of her grounding beliefs lacks adequate or sufficient epistemic support, then she ought to drop those beliefs. What is important to note here is that dropping beliefs that lack sufficient rational support is a constitutive part of complying with one's commitment of Doxastic Rationality. Failure to abandon these unsupported beliefs constitutes a failure to live up to this cognitive commitment. Consequently, one would be in error. Of course, an agent may find herself in a situation in which she still cannot determine which grounding beliefs should be abandoned or revised. In these cases, the agent should continue to suspend judgment on the appropriate propositions, and she should continue to undergo her process of rational scrutiny.[165] Once

[165] According to my view, one need not suspend judgment on every belief in one's belief-set. At most, one can keep suspending judgment on only the beliefs that one judges to be directly culpable in generating a conflict of semantic commitments.

the agent accumulates enough evidence to determine which beliefs should be retained or rejected, then she can proceed to make the necessary rational revisions to her belief-set.

Performing rational scrutiny is an important step in our normative procedure to ameliorate inconsistent semantic commitments. Moreover, it guides agents in effectively determining which grounding beliefs and corresponding semantic commitments one should retain. And this operational step is underpinned and rationally motivated by the cognitive commitment of Doxastic Rationality.

This normative procedure can be employed to resolve inconsistent semantic commitments. The procedure does not simply appeal to the weight of semantic commitments. Rather, it appeals to our cognitive commitments to determine which propositions an agent should suspend judgment on, and which grounding attitudes she should revise or reject. Doxastic Veridicality prohibits one to consciously believe contradictory things at the same time. If one is semantically committed to believe contradictory things, then one ought to take some non-belief stance toward these contradictory propositions. In addition to this, Doxastic Rationality guides one to suspend judgment or revise one's grounding beliefs that give rise to the conflicting or problematic semantic commitments. Furthermore, Doxastic Rationality facilitates rational scrutiny. One ought to believe for good reasons. Hence, any grounding belief ought to be revised or rejected if it lacks sufficient rational support. Our cognitive commitments therefore provide the normative basis for constructing a plausible procedure for dealing with inconsistent semantic commitments.

4. The Primacy of Cognitive Commitments

Cognitive commitments normatively guide us in problematic situations involving conflicting semantic commitments to believe inconsistent things. In this section I demonstrate that our cognitive commitments have a broader regulative role that is not simply restricted to those problematic cases

involving inconsistent semantic commitments. I further argue for what I refer to as the "Primacy Thesis":

> **Primacy:** Cognitive commitments are normatively fundamental or prior to semantic commitments, and thus they play a more important role in governing our epistemic lives.

Recall the account of *semantic commitments* defended in chapter 3.

> **Semantic Commitment:** If an agent, *S*, has a belief (or non-empty set of beliefs) that implies a proposition, *p*, via some correct norm or rule of inference, *R*, and *S* possesses the relevant concepts, understanding, and cognitive abilities to settle matters concerning *p*'s truth, such that she can competently form a new belief that *p* via *R*, then *S* is normatively committed to cognitively endorse *p* when she mentally determines the truth or epistemic acceptability of *p*. *S* can fully satisfy this commitment by believing *p*, or she can partially satisfy this commitment by accepting *p*.

This argument for the Primacy Thesis can be expressed as follows:

> **(P1)** Appealing to cognitive commitments will alleviate or eliminate inconsistent and problematic semantic commitments.
> **(P2)** Satisfying a given semantic commitment is rationally permitted or appropriate when doing so does not conflict with one's cognitive commitments.
> **(P3)** If premises (P1) and (P2) are true, then cognitive commitments are normatively fundamental or prior to semantic commitments, and thus they play a more important role in governing our epistemic lives.
> **(C)** Therefore, cognitive commitments are normatively fundamental or prior to semantic commitments, and thus they play a more important role in governing our epistemic lives.

The third premise (P3) seems true. Hence, establishing (P1) and (P2) will show that the Primacy Thesis is true. The propositions expressed in these two premises are intimately related. Taken together, they suggest that our semantic commitments are fundamentally dependent on our cognitive commitments. I defend each of these premises in turn.

First, (P1) is motivated by the arguments in the previous section. In cases of conflicting semantic commitments to believe inconsistent things, our cognitive commitments can determine which semantic commitments to revise. First, Doxastic Veridicality prohibits belief in contradictory things. Second, Doxastic Rationality binds agents to suspend judgment on inconsistent propositions and perform rational scrutiny in order to determine which beliefs ought to be revised. Complying with

one's cognitive commitments will eliminate inconsistent semantic commitments. Hence, our cognitive commitments guide and override our inconsistent semantic commitments.

Moreover, there can be problematic cases of semantic commitments that do not involve two or more conflicting commitments to believe inconsistent or contradicting things. The cases I have in mind concern "epistemically suspicious" semantic commitments. These cases involve agents with a commitment to believe a proposition that seems highly unlikely or absurd. Consider a case where *S* comes to competently recognize that the proposition, *that Bigfoot exists* follows from her current beliefs. In this situation, *S* has never seriously considered this proposition before. Nor is she aware of having any beliefs that directly entail that Bigfoot does not exist. However, when *S* comes to consider this proposition, she takes it to be counterintuitive or unlikely. Accordingly, she is reluctant to think that it is the case that Bigfoot exists. In this situation, *S* would experience some normative tension. In order to alleviate this tension, *S* can rationally scrutinize and evaluate her rational basis for the grounding beliefs that give rise to her semantic commitment to believe the proposition, *that Bigfoot exists*. This move is facilitated by her commitments of Doxastic Rationality. This cognitive commitment guides her in a process of rational scrutiny. It prescribes that she ought to revise her grounding beliefs that lack sufficient justification. Revising one's grounding beliefs will eliminate any problematic or epistemically suspicious semantic commitment. Doxastic Rationality therefore guides and overrides problematic or epistemically suspicious semantic commitments. Consequently, (P1) is true.

Let us now turn to a defense of (P2). Satisfying a given semantic commitment will always be rationally permitted when it does not conflict with one's cognitive commitments. Indeed, satisfying a semantic commitment might not be the *overall* rational thing to do when one takes into consideration the evidential basis for the belief(s) that engender that commitment. Following Sam Shpall, these doxastic-formation commitments are "pro tanto" and merely contribute to what one ought to believe (Shpall 2014, p. 156). Hence, these semantic commitments can be outweighed by further obligations

and rational requirements an agent might have (Shpall 2014, p. 156). Specifically, they can be overturned or permitted depending on one's cognitive commitments. For instance, an agent, *S*, might decide to evaluate the evidential basis for any number of her currently held beliefs. If one of those beliefs does not pass the test of rational scrutiny and is thus judged to lack sufficient justification, then *S* is rationally compelled to revise and drop that belief. Consequently, the semantic commitments stemming from her irrational or epistemically unsupported belief should not be satisfied. By rationally revising one or more grounding beliefs *S* can rationally eliminate the corresponding semantic commitments.[166] Doxastic Rationality prescribes this revision. What this demonstrates is that satisfying *S*'s cognitive commitments is crucial for determining which semantic commitments are rational.

This point can be generalized to all agents and to all semantic commitments. A semantic commitment that stems from one's belief will provide normative pressure to satisfy the content of that commitment. One satisfies that commitment (partially or fully) by forming the appropriate cognitive endorsement. However, satisfying a semantic commitment might not be the overall rational thing to do. In cases where a semantic commitment stems from an irrational or epistemically unsupported belief, then one ought not to satisfy that commitment. Rather, one ought to revise one's belief-set and cognitive position. This is prescribed by the relevant cognitive commitments. Rationally satisfying one's semantic commitments is therefore dependent on the appropriate satisfaction of one's cognitive commitments. Consequently, (P2) is true.

I have shown that premises (P1) and (P2) are true. If (P1) – (P3) are true, then the Primacy Thesis is true. I conclude that our cognitive commitments are normatively fundamental or prior to our semantic commitments, and thus they play a more important role in governing our epistemic lives.

[166] To be clear, one can eliminate a semantic commitment to believe *p* by abandoning the beliefs that give rise to this commitment. However, one could still come to believe *p* on independent grounds if one comes to possess independent evidence that is sufficient to warrant a rational belief that *p*.

5. Conclusion: A Hierarchy of Commitments Facilitating Epistemic Rationality

To recapitulate, I provided a skeptical argument that showed that we lack sufficient reason to rationally endorse DW. First, it is not clear that semantic commitments will always have different degrees of known normative weight. Second, it is not clear that we should always appeal to the normative weight of semantic commitments to determine how one ought to rationally hold and revise one's beliefs. Hence, in cases involving conflicting semantic commitments to believe inconsistent things, we need to appeal to some alternative standard or procedure to eliminate the normative tension of conflicting semantic commitments. I demonstrated that this normative tension can be alleviated by appealing to our cognitive commitments. I further argued that cognitive commitments are normatively prior to our semantic commitments, and thus they play a more important role in governing rational thinking.

One interesting consequence of this view is that semantic commitments differ fundamentally from other normative elements that guide human behavior and thinking. For instance, reasons are normative. Pragmatic reasons to do Φ can weigh against reasons not to do Φ. Other things being equal, the more decisive reasons will determine whether one should do Φ. Epistemic reasons function in a similar way. Epistemic reasons to believe *p* can weigh against the epistemic reasons one has for not-*p*. Semantic commitments, however, differ in the sense that they do not have degrees of weight. These commitments are a unique kind of normative constraint governing belief-revision.

Moreover, this account deepens our understanding of the normative structure that constrains rational thinking. Cognitive commitments are normatively fundamental and prior to semantic commitments. However, satisfying both of these commitments is a constitutive part of epistemic rationality. These two kinds of commitments are not in normative competition with each other. Rather, they function together to determine which propositions an agent can rationally believe, and how agents ought to revise their current belief-sets or cognitive positions. Rationally satisfying one's semantic commitments will facilitate rational diachronic belief-formation. However, this process is

governed by our fundamental cognitive commitments. In situations involving conflicting semantic commitments, our cognitive commitments can determine which semantic commitments one should revise or abandon. Moreover, cognitive commitments guide and permit the rational satisfaction of individual semantic commitments. Satisfying one's cognitive commitments requires that one revises or abandons any irrational grounding beliefs. Revising these beliefs will eliminate problematic semantic commitments to believe epistemically dubious things. Our semantic commitments are therefore dependent on the appropriate satisfaction of our cognitive commitments. Consequently, this account illuminates several important features of the normative structure governing human thinking. Epistemic rationality is partly constituted by a normative *hierarchy* of cognitive commitments and semantic commitments.

V. Robust Epistemic Rationality Through Commitments and Requirements

Adam C. Sanders

0. Introduction

Theories of rationality typically focus on how agents should form, hold, and maintain mental states. According to many theorists, rationality specifically governs the coherence and consistency of our mental states: agents ought to have attitudes that correctly cohere together.[167] For example, rationality might prohibit an agent from having some mental state, M_1, given that she already holds another mental state, M_2. Simultaneously possessing these attitudes is prohibited because M_1 and M_2 do not fit together in any appropriate way. It is commonly thought that rationality is constituted by a system of external, mind-independent *rules* or *requirements* that agents should comply with. If rationality contains a rule or requirement, *R*, prohibiting one from having both M_1 and M_2, at a given time, *t*, then one ought to comply with *R* and refrain from having both M_1 and M_2 at *t*. In addition to this, there are cognitive and semantic commitments that depend on an agent's mental states. These internal normative commitments also play an important role in governing how we ought to form and hold certain cognitive states. Consequently, these commitments must factor into a theory of epistemic rationality. This raises several important issues. First, how do these two kinds of normative constraints compare and relate to one another? Should we understand epistemic rationality as being constituted by both requirements and commitments? Are these two types of normative constraints compatible? Any plausible theory of epistemic rationality should address these issues.

In this chapter I focus on the nature and composition of epistemic rationality. I argue for what I call the "Robust View". On this account, epistemic rationality is construed as a system of external (mind-independent) requirements and internal (mind-dependent) doxastic commitments.[168] These two

[167] See, for example, Donald Davidson (1985) and Jonathan Way (2011).

[168] This idea has also been shared by other theorists. According to Coliva, our commitments stemming from our beliefs are "*constitutive* of epistemic rationality" (Coliva 2019, pp. 6-7). Sam Shpall also notes

distinct normative constraints, I contend, are compatible. Both requirements and commitments prescribe how an agent ought to form, revise, and maintain cognitive endorsements. Countenancing both external requirements and internal commitments provides a more plausible and complete picture of epistemic rationality. Moreover, incorporating cognitive and semantic commitments into our theory of epistemic rationality helps to resolve certain problems facing standard accounts of rationality.

Several caveats are in order. First, I focus on epistemic rationality. Hence, I am not concerned with practical rationality. However, I should note that if rationality as a whole includes both practical and epistemic rationality, and epistemic rationality includes both requirements and commitments, then rationality as a whole includes both requirements and commitments. Second, I will focus on epistemic rationality concerning belief. I demonstrate that epistemic rationality concerning belief is constituted by both commitments and requirements. But it is important to keep in mind that attitudes of acceptance and supposition are also rationally governed by normative commitments (see chapter 2). Moreover, I am sympathetic to the view that these other cognitive endorsements (especially acceptance) can be rationally governed by both commitments and requirements (see the conclusion of this dissertation). However, I will set this further issue aside and focus on belief.

The structure of this chapter is as follows. In the first section I briefly articulate a standard theory maintaining that epistemic rationality is composed of external, mind-independent rules. In the second section I contrast this picture with the view that belief is rationally governed by internal, mind-dependent commitments. In the third section I advance the Robust View. I defend the thesis that the Robust View is a viable and consistent theory. In section four I discuss the normative relation between doxastic commitments and epistemic requirements. I show that doxastic commitments are normatively prior. In section five I argue that the Robust View is a more tenable and preferable theory

---

that rationality should incorporate commitments (Shpall 2013). However, Shpall's view is fundamentally different given that the commitments he defends are different from the commitments that I defend. I say more about his view later in the chapter.

of epistemic rationality. This account provides a more complete picture of epistemic rationality. Moreover, it resolves some important issues facing standard theories of rational requirements. I conclude the chapter in section six.

1. REQUIREMENTS OF EPISTEMIC RATIONALITY

Rationality concerns the relationships and consistency of sets of mental states.[169] Our notion of *rationality* can have several meanings. According to John Broome, we might understand *rationality* in two ways: (1) as a property that agents can possess, or (2) as a collection of requirements that agents should satisfy (Broome 2007, pp. 361 – 362). Many contemporary theorists seem to focus on the latter sense of the term.[170] That is to say, rationality is taken to be a collection of requirements that govern the relationships and consistency of mental states. A fully rational agent is one who satisfies each requirement of rationality. In what follows, I focus on the epistemic rationality of belief. I will briefly discuss some of the requirements that purportedly constitute standard views of epistemic rationality.

1. 1 STANDARD EPISTEMIC REQUIREMENTS

The primary question to be addressed is the following: Which requirements factor into epistemic rationality? I will not provide a complete list of epistemic requirements. Rather, I briefly consider what I take to be some of the standard requirements that many theorists appeal to when discussing epistemic rationality. I will then discuss the logical formulation of these requirements in section 1.2.

The first requirement deals with the consistency of our beliefs. Many theorists adopt some version of the following epistemic rule:

[169] For more discussion on the idea that rationality governs relationships between mental states see, for example, Donald Davidson (1985), Broome (2005), and Jonathan Way (2011).
[170] See, for example, Broome (2007) and Way (2011).

> **Non-contradiction:** Rationality requires of *S* that (if *S* believes *p*, then *S* does not believe not-*p*).[171]

The *Non-contradiction* requirement seems fairly plausible. To believe *p* is to be in a mental state that posits *p* as being true. Consequently, possessing a belief that *p* would not appropriately fit with a belief that not-*p*. Rationality prohibits having this set of mental states.

A second standard requirement governs the kind of epistemic support that grounds our beliefs. It is typically thought that beliefs ought to be based on good epistemic reasons that support the truth of their target propositions.[172] A belief that *p* is irrational when it lacks sufficient epistemic support. This has led some theorists to posit an epistemic requirement governing the rational support of beliefs (see John Brunero 2010, p. 29). Consequently, we might countenance the following rule:

> **Justification:** Rationality requires of *S* that (if *S* believes that she has sufficient epistemic reasons and evidence establishing the truth of *p*, then *S* believes *p*).[173]

The *Justification* requirement governs certain sets of mental states, namely, beliefs about propositions and beliefs about one's evidence supporting the truth of those propositions. On the face of it, this

---

[171] This is a slightly modified version of Way's formulation of "Belief Consistency", which maintains that "Rationality requires of you that [if you believe that p, you do not believe that not-p]" (Way 2010, p. 1059). I have substituted "*S*" for "you". Other theorists adopt a similar requirement. For example, Broome notes that "[r]ationality requires of you that you do not both believe *p* and believe not-*p*" (Broome 2005, p. 322).

[172] This idea is endorsed by many philosophers. For example, Hume insisted that one should "proportion [one's] belief to the evidence" (Hume 1975, p. 110). Richard Foley has maintained that a given belief ought to be based on epistemic reasons that support the truth of the target proposition (Foley 1993, p. 16). Moreover, Foley has defended the following principle: For any person S, time t, and proposition p, if S has any doxastic attitude at all toward p at t and S's evidence at t supports p, then S epistemically ought to have the attitude toward p supported by S's evidence at t (Foley 2000, p. 679). Hence, if *S* takes her evidence to establish the truth of *p*, then *S* ought to believe *p*. For a belief should be based on good epistemic reasons supporting the truth of its target proposition. Finally, Duncan Pritchard has also claimed that a belief should be appropriately responsive to "rational considerations" (Pritchard 2016, p. 90).

[173] This requirement is based on a formulation that Brunero discusses. He expresses the rule as follows: "Rationality requires one to believe that *p*, if one believes that there is conclusive evidence that *p*" (Brunero 2010, p. 29). I refer to this rule as "Justification". I have also slightly modified the formulation of his rule in order to provide a consistent list of epistemic requirements that utilize the same symbols, terminology, and logical structure.

requirement seems plausible. *Justification* captures the intuition that one ought to believe *p* when one believes that one's available evidence establishes the truth of *p*. Failure to believe *p* while simultaneously holding the belief that *p* is epistemically supported to a sufficient degree results in a system of mental states that do not fit together in any appropriate way. Indeed, one ought to avoid this inappropriate set of mental states.

Finally, it is sometimes thought that an agent is under some kind of obligation or normative constraint to believe the deductive consequences of her current beliefs.[174] Specifically, epistemic rationality is taken to include some version of the following *Closure* requirement:

> **Closure:** Rationality requires of *S* that (if *S* believe *p* and believes *p* entails *q*, then *S* believes *q*).[175]

Closure seems *prima facie* plausible. An agent, *S*, would not be in a *fully* rational state if *S* had a set of beliefs, *Γ*, that entailed a proposition, *p*, and *S* formed some other non-belief state, *M*, toward *p* (e.g., believing not-*p*). The set of states containing *M* and the members of *Γ* would not cohere together in

---

[174] For example, Frank Jackson asserts that "[s]omeone who believes that P, and that if P then Q, *ought* to believe that Q. It is not simply that, by and large, they do believe that Q. It is that if they don't, there is something *wrong*" (Jackson 2000, p. 101). Now, many theorists contend that the normative force that compels agents to believe the consequences of their beliefs stems from doxastic *commitments* (See, for example, Levi 1997; Millar 2004; Bilgrami 2006; Shpall 2014; and Liberman and Schroeder 2016). I discuss several accounts that appeal to these kinds of commitments in chapter 3. However, I argue for a modified view that posits *semantic commitments*. Contrary to my semantic commitment view, we might hold that the normative force to believe the deductive consequences of one's beliefs stems from an external *requirement* or *rule*. At this moment I am focusing on epistemic requirements. But I will say more about how these rational requirements differ from cognitive and semantic commitments in section 2 of this chapter.

[175] Many theorists talk about rationality as a set of rules that consists of some *Closure* requirement. John Broome discusses a "*Modus ponens*" rule, which states that "[r]ationality requires of *N* that, if *N* believes *p* and *N* believes that if *p* then *q*, and if it matters to *N* whether *q*, then *N* believes *q*" (Broome 2008, p. 153). Jonathan Way articulates a similar requirement he refers to as "Closure-W", which states that "[r]ationality requires that [if you believe that p and believe that if p then q, then you believe that q]" (Way 2011, p. 228). One interesting question worth asking is whether an agent is only required to believe the deductive consequences of her beliefs. I have argued that one can be normatively committed to believe things that inductively follow from one's beliefs (see chapter 3). However, I will not pursue these issues here. For the time being I will take it for granted that *Closure* (as I have formulated it) is a standard requirement of epistemic rationality.

any appropriate way. Hence, rationality dictates that one should avoid these types of cognitive situations.

I take it that *Non-contradiction*, *Justification*, and *Closure* are plausible and uncontroversial requirements that partly constitute standard accounts of epistemic rationality. This does not suggest that these requirements exhaust the rules of epistemic rationality. For instance, we might follow Gilbert Harman here and insist that we ought to comply with the following principle:

> *Clutter Avoidance* One should not clutter one's mind with trivialities (Harman 1986, p. 12).

I am sympathetic to the idea that rationality (pragmatic or epistemic) includes some version of the Clutter Avoidance principle.[176] Unlike the other requirements previously mentioned, this principle is not *necessarily* concerned with the way an individual belief will cohere with one's currently held belief-set. Moreover, I admit that there can be other, similar requirements governing beliefs and belief-formation.

To be clear, it is an important task for theorists of rationality to advance a satisfactory and complete list of epistemic requirements. However, I will set this further issue aside. My arguments in this chapter are not held hostage to the truth of any particular view that purports to list all the requirements that exhaust epistemic rationality. As I will clarify later on, the primary aim of this chapter is to show that epistemic rationality is composed of both epistemic rules and doxastic commitments. In doing so, I must demonstrate that these two kinds of normative constraints are consistent, and that countenancing epistemic rules will not be normatively superfluous. The epistemic rules of Non-contradiction, Justification, and Closure provide the biggest challenge. Indeed, principles like Clutter Avoidance are straightforwardly consistent with our cognitive commitments and semantic

[176] Perhaps *Clutter Avoidance* is not a genuine *epistemic* principle. It might be associated with pragmatic rationality. Nevertheless, my point here is that we want a system of epistemic rationality that does not conflict with other plausible principles governing thinking or behavior.

commitments.[177] Moreover, I suspect that any additional principles or rules like Clutter Avoidance will not be normatively redundant or superfluous given that they demand fairly different things from us. For my purpose, then, what is important to keep in mind is that many standard accounts of epistemic rationality will incorporate some versions of Non-contradiction, Justification, and Closure. I take these to be the crucial rules of epistemic rationality.

1.2 The Normativity, Scope and Logical Formulation of Rational Requirements

Rational requirements govern how certain mental states should cohere together. Accordingly, if *R* is an *epistemic* requirement of rationality, then *R* will dictate that a given agent should have some set of beliefs that stand in a specific kind of relationship to each other. For example, *Non-contradiction* prescribes that agents should have non-contradictory beliefs: it is a requirement of rationality that if an agent, *S*, believes *p* at time, *t*, then *S* does not believe not-*p* at *t*. Rational requirements therefore have a normative force. According to many theorists, if a rule of rationality, *R*, requires that an agent brings about Φ, then *S* ought to bring about Φ (see Way 2010, p. 1058; Shpall 2014, pp. 156-157). Hence, rational requirements provide agents with reason to comply with them.[178]

It follows from this that the epistemic rational requirements stipulated above will normatively govern how an agent ought to form and retain certain relationships that hold between her individual

[177] My main concern is whether external rules of rationality conflict with internal commitments generated by beliefs. It is a separate question whether or not these rules of rationality will conflict with each other. Indeed, it seems that the rule of Closure can be in tension with the rule of Clutter Avoidance. One might, for example, believe *p* and believe that *p* entails *q*, and yet *q* is a trivial belief. According to Closure, one should believe *q*. But according to Clutter Avoidance, one should not believe *q*. It is possible that some of these rules can override others. It is also possible that cognitive commitments play a role in determining which rules of rationality one should follow. However, given my purposes I will set these further issues aside.

[178] The normative force of rational requirements has been discussed by many philosophers. According to Way, some theorists take for granted the following: "If rationality requires you to A, you should A" (Way 2010, p. 1058). Shpall also notes that rational requirements are "decisive", in that they determine what one ought to do (Shpall 2014, p. 157). I will say more about this later in the chapter.

beliefs. This, however, leads to certain issues concerning the scope and logical formulation of these requirements. According to the formulations above, epistemic requirements have what Broome calls "wide scope" (Broome 2007). More precisely, "rationality requires" ranges over a conditional proposition of the form, *if p then q*, where "*p*" and "*q*" refer to certain mental states that an agent possesses (see Broome 2007, p. 360; Brunero 2010, p. 29; Way 2010, p. 1059; and Shpall 2013, p. 718). Consequently, one is required to comply with this entire conditional proposition. For example, consider the *Justification* requirement. In its wide scope formulation, "rationality requires" ranges over the conditional proposition, *if S believes that she has sufficient epistemic reasons and evidence establishing the truth of p, then S believes p*. Hence, an agent is required to satisfy this entire conditional proposition. She can thus form a belief that *p* when she believes that *p* is sufficiently supported by her evidence, or she can drop the antecedent belief.[179] This leads to appropriate relationships that hold between beliefs.

We can contrast this logical formulation with what Broome and others refer to as "narrow scope" requirements (see, for example, Broome 2007). On narrow scope views, "rationality requires" will range over the consequent of a requirement in conditional form (see Broome 2007, pp. 359 – 360; Brunero 2010, p. 29; Way 2010, pp. 1057 – 1058; Shpall 2013, p. 718). For example, if we construe *Justification* to be a narrow scope requirement, then we would formulate it in the following way:

> **Justification*:** If *S* believes that she has sufficient epistemic reasons and evidence establishing the truth of *p*, then *S* is rationally required to believe *p*.

In *Justification**, the operator, "rationality requires", ranges over the consequent specifying that *S is rationally required to believe p*. Accordingly, *S* accrues a normative constraint to form a belief that *p* once she first believes that *p* is epistemically well-supported.

Many theorists prefer wide scope formulations. One reason for this is that wide scope requirements avoid a certain kind of normative "boot strapping" or "detachment" problem facing

[179] This disjunctive satisfaction condition is what many theorists refer to as "symmetry" (see, for example, Way (2010)). I will say more about this when I discuss the application of the Robust View.

narrow scope views (for a discussion of this issue, see Broome 2007, pp. 359 – 370; Way 2010, pp. 1057 – 1059; and Shpall 2013, p. 718). To illustrate this problem, consider *Non-contradiction*. Let us reformulate this rule so that "rationality requires" has a narrow scope ranging over the consequent of the relevant conditional proposition. Hence, we get the following narrow scope formulation:

> **Non-contradiction*:** If S believes p, then rationality requires that S does not believe not-p.

According to Way, there can be situations where an agent, *S*, believes *p* for irrational reasons – i.e., *S* believes *p* but recognizes that her evidence supports the truth of not-*p* (Way 2010, p. 1058). According to *Non-contradiction**, *S*'s irrational belief makes it so that *S* is rationally required *not* to believe not-*p*. But Way correctly points out that this is implausible: *S* ought to believe not-*p* given that her evidence conclusively supports the negation of *p* (Way 2010, p. 1058). This worry seems to generalize to many (if not all) narrow scope requirements. Wide scope formulations, however, seem to avoid this worry. Way notes that wide scope rules only insist that one should satisfy an entire conditional proposition and thus avoid any inappropriate sets of attitudes (Way 2010, pp. 1059 – 1060). If, for example, one believes *p*, then one can satisfy the wide scope version of *Non-contradiction* by either doing one of two things: (1) refraining from also believing not-*p*, or (2) abandoning the initial belief that *p* (see Way 2010, p. 1060). One is only rationally required to *either* do (1) or (2) in order to comply with *Non-contradiction*. Hence, if an agent irrationally believes *p*, then she is not required to refrain from believing not-*p*. Rather, she is licensed to drop her irrational belief.

The logical structure of rational requirements is a contentious issue. I will not rehearse the different arguments put forward by proponents of wide scope requirements. Nor will I address the different strategies employed to defend the narrow scope view.[180] I take it for granted that the *epistemic*

[180] Some theorists defend the view maintaining that some requirements are narrow scope requirements. See, for example, Niko Kolodny (2005), pp. 509 – 563 and Mark Schroeder (2004), pp. 337 – 364.

rules of *Non-contradiction*, *Justification*, and *Closure* are construed as wide scope rules.[181] These requirements are such that they apply to all agents. On my view, it is more intuitive to characterize these particular requirements as having a wide scope logical structure. This characterization helps to demarcate epistemic rules from narrow scope doxastic commitments (I discuss this in the next section). Indeed, I will demonstrate that we are left with a more plausible picture of epistemic rationality when we countenance both wide scope rational requirements and narrow scope doxastic commitments.

2. Contrasting Epistemic Requirements with Commitments of Belief

Epistemic requirements prescribe how one ought to rationally form and maintain beliefs. However, I have previously argued that doxastic commitments also contribute to epistemic rationality (see chapter 2). These two normative constraints differ in several important ways. The task at hand is to briefly specify the fundamental features distinguishing requirements from doxastic commitments.

First, recall that consciously formed beliefs engender the following normative commitments.

> **Doxastic Veridicality:** If an agent, *S*, at time, *t*, believes *p* by judging that *p* obtains, then *S* is committed to regard *p* as being true at *t* (or obtaining at *t*), such that *S* cannot also consciously believe not-*p* at *t* without revising her initial belief that *p*

> **Doxastic Rationality:** If an agent, *S*, believes *p*, then *S*'s belief commits her to base that belief on sufficient epistemic grounds supporting the truth of *p*, and to be sensitive to new information and evidence concerning the truth or falsity of *p*.

> **Doxastic Integrity:** If an agent, *S*, forms a belief that *p* by judging that *p* is true, then *S* accrues a normative commitment to maintain *p* as being true in some doxastic form, such that, *S*'s belief can only be rationally revised or terminated via another judgment concerning the truth of *p* (e.g., a judgment that not-*p* obtains), or that, given the available evidence, not-*p* is just as likely as *p* to be true.

Additionally, I have argued that our beliefs partly engender the following semantic commitment:

[181] To be clear, I am sympathetic to the idea that some requirements of rationality might have a narrow scope formulation. It does not follow from my view that *all* requirements of rationality (theoretical and practical) are wide scope rules.

> **Semantic Commitment:** If an agent, *S*, has a belief (or non-empty set of beliefs) that implies a proposition, *p*, via some correct norm or rule of inference, *R*, and *S* possesses the relevant concepts, understanding, and cognitive abilities to settle matters concerning *p*'s truth, such that she can competently form a new belief that *p* via *R*, then *S* is normatively committed to cognitively endorse *p* when she mentally determines the truth or epistemic acceptability of *p*. *S* can fully satisfy this commitment by believing *p*, or she can partially satisfy this commitment by accepting *p*.

Cognitive commitments are constitutive elements of a token belief that *p*. These commitments govern how one ought to mentally treat *p* and how one should respond to epistemic considerations concerning the truth or falsity of *p* (see chapter 2). On the other hand, *semantic commitments* are extraneous doxastic commitments that prescribe how one ought to revise one's current belief-set based on one's current beliefs, understanding, conceptual repertoire, and cognitive abilities (see chapter 3).

Doxastic commitments differ from epistemic requirements in several ways. First, doxastic commitments have a different kind of logical structure. Epistemic requirements are taken to have a wide scope formulation. Contrary to this, doxastic commitments have a narrow scope structure. Consider again each of the cognitive commitments articulated above. Each specification of those commitments takes the following form: If an agent, *S*, believes *p*, then *S* is committed to Φ. This is a conditional statement. Note that "Φ" denotes some way that an agent, *S*, will mentally treat *p*. Moreover, the way *S* should mentally treat *p* is dependent on the kind of cognitive commitment involved. But what is important to keep in mind here is that "committed" takes narrow scope over the consequent of this conditional statement. Similar considerations apply to semantic commitments. Hence, our doxastic commitments are formulated as narrow scope normative constraints.

The difference in logical structures illuminates and reflects several further ways that doxastic commitments differ from epistemic requirements. First, note that epistemic requirements of rationality are external rules that do not depend on the mental states of individuals. Hence, these rules apply to all agents at all times. As Shpall puts the point, these requirements are not "agent-dependent" (Shpall 2014, pp. 154-155). Doxastic commitments, however, depend on how one takes the world to

be. They are "agent-dependent" (Shpall 2014, pp. 154- 155).[182] If an agent believes *p*, then she (1) comes to acquire certain normative commitments that govern how she should treat *p*, and (2) comes to acquire certain semantic commitments governing what she should cognitively endorse based on her belief that *p*, her cognitive abilities, her understanding, and her conceptual repertoire. Doxastic commitments are therefore dependent on our mental lives and individual cognitive abilities. They are subjective constraints (see chapters 2 and 3).[183]

Moreover, doxastic commitments are distinct in that one can either satisfy one's commitments or rationally abandon them by rationally giving up the attitudes or cognitive position that engenders them. Following Shpall, we can say that commitments are "escapable" when their grounds are removed or revised (Shpall 2014, p. 153). Epistemic requirements are different. These requirements cannot be rationally abandoned, ignored, or escaped simply by revising one's beliefs (Shpall 2014, p. 154). They exist regardless of how we think or take the world to be. Indeed, we either comply with these external rules or we fail to comply with them.

Finally, epistemic requirements differ in their normative force. Many theorists take requirements or rules of rationality to be normatively decisive or conclusive: one simply ought to comply with the rules of rationality at all times (see, for example, Shpall 2014, pp. 156 -1 57). The requirement of *Non-contradiction*, for example, gives all agents at all times conclusive or decisive reason to not have contradictory beliefs at a given time. Similar considerations apply to other rules of rationality.

---

[182] To be clear, Shpall is working with a different conception of doxastic commitments. He focuses on "rational commitments" (see Shpall 2014). I have argued that our doxastic commitments come in two types: cognitive commitments and semantic commitments (see chapters 2 and 3). I will say more about how doxastic commitments factor into epistemic rationality in section 3. I will also note how this view differs from Shpall's account.

[183] This idea is shared by many theorists. See, for example, Annalisa Coliva on beliefs as commitments (Coliva 2016 and 2019). Moreover, Margaret Gilbert argues that individual intentions and decisions can generate "commitments of the will" (see Gilbert 2013 and 2014).

Contrary to this, doxastic commitments have different degrees of normative strength. First, some doxastic commitments are normatively conclusive. Specifically, the *cognitive commitments* of a belief that *p* decisively prescribe how one ought to mentally treat *p*. Moreover, the cognitive commitment of Doxastic Veridicality cannot be violated. This commitment is a constitutive component of a consciously held belief that partly individuates this particular kind of propositional attitude (see Coliva 2019 and chapter 2). On the other end of the spectrum, our semantic commitments are not always decisive in determining how one ought to revise one's belief-set. In some cases, one's semantic commitment to believe *p* can be outweighed or overruled by one's cognitive commitments. Borrowing some terminology from other theorists, we can say that these commitments are "pro tanto" and thus merely factor into how one should form or revise one's beliefs (see Broome 2000, p. 91; and Shpall 2014, pp. 156 - 157).[184] But these semantic commitments are not typically decisive on their own. If, for example, an agent, *S*, is semantically committed to form an irrational or epistemically absurd belief that *p*, and this semantic commitment is based on a set, Γ, of irrationally formed beliefs, then *S* should refrain from believing *p*. What *S* ultimately *ought* to do is undergo a process of rational scrutiny and come to revise her initial belief-set, Γ. This doxastic-revision is prescribed by her cognitive commitments (especially Doxastic Rationality). By appropriately revising Γ, *S* will eliminate and escape any irrational or epistemically dubious semantic commitments to believe irrational or epistemically

[184] According to many theorists, reasons have different degrees of weight that can be appealed to in order to determine how one should act and think (see, for example, Raz (1975), p. 25; Broome (2004), p. 91; Shpall (2014), p. 156; Liberman and Schroeder (2016), pp. 104-120; and Lord and Maguire (2016), p. 10). In some cases, one can have reasons to do Φ while also having conflicting reasons to not do Φ (see Dancy 2000, p. 4). If my reasons to do Φ are stronger than my reasons against doing Φ, then I should do Φ. Following Lord and Maguire, we can say that the reasons to do Φ in this case are "decisive" (Lord and Maguire 2016, p. 10). According to Broome, this suggests that reasons are "pro tanto" and can be "weighed against each other…to determine" what one should do (Broome 2004, p. 91). Contrary to this, Broome contends that requirements are "strict" (Broome 2004, p. 91). As Shpall puts the point, "all things considered concepts like 'requirement'…do not merely weigh in favor of something" (Shpall 2014, p. 156). Some doxastic commitments, namely, semantic commitments, are pro tanto in this sense. These commitments are not normatively strict. Rather, a semantic commitment can be outweighed by the normative force of one's cognitive commitments.

absurd propositions. Consequently, not all doxastic commitments have an inherently conclusive normative force. Indeed, semantic commitments are defeasible and can therefore be annulled or undercut by cognitive commitments. This further feature demarcates our doxastic commitments (cognitive and semantic) from epistemic rational requirements.[185]

3. Robust Epistemic Rationality Through Commitments and Requirements

At this point we have two distinct kinds of normative constraints that govern our beliefs and diachronic processes of belief-formation. Epistemic requirements are wide scope external rules that apply to all agents. Doxastic commitments (both cognitive and semantic) are narrow scope normative commitments that depend on the subjective mental lives of individual agents. This suggests that epistemic rationality might be best construed as a system of rules and commitments. However, this raises some serious concerns. First, it is not clear how these two kinds of normative constraints interact and fit together to promote epistemic rationality. Second, if we countenance both kinds of normative constraints, then we can ask whether this view leads to certain normative redundancies. If our beliefs engender normative commitments that contribute to epistemic rationality, then does this lead to the view that epistemic requirements are superfluous and unnecessary? In this section I address these issues. My contention is that it is a mistake to characterize epistemic rationality as a system comprised solely of requirements. On the contrary, I show that we ought to endorse the *Robust View* of epistemic rationality. Accordingly, doxastic commitments and epistemic requirements both play a role in governing how we ought to form and maintain rational beliefs.

[185] It is important to note that my account of doxastic commitments fundamentally differs from Shpall's account. Shpall seems to think that all commitments are "pro tanto" and "play a *non-decisive* weighing role" (Shpall 2014, p. 157). I do not share this view. I have argued for a more robust theory of commitments. On this view, some commitments have a decisive role in determining how one should think. Specifically, cognitive commitments decisively prescribe how one should mentally treat a proposition.

My primary aim in this chapter is to show that the Robust View is a plausible and preferable theory of epistemic rationality. In this specific section of the chapter I defend the thesis that the Robust View is a viable and consistent theory. In defending this thesis, I take the following claim for granted:

> **(P1)** The epistemic requirements of *Non-contradiction*, *Justification*, and *Closure* are the primary or standard rational rules governing beliefs.

These wide scope rules were articulated in section 1.1 above. Second, I defend the following claim:

> **(P2)** The epistemic requirements of *Non-contradiction*, *Justification*, and *Closure* are normatively compatible with the established doxastic commitments associated with consciously formed beliefs.

The relevant doxastic commitments are *Doxastic Veridicality*, *Doxastic Rationality*, *Doxastic Integrity*, and the *semantic commitments* partly engendered by our beliefs. For an epistemic requirement, *R*, to be *normatively compatible* with one's doxastic commitments requires that (1) *R* does not contradict or militate against one's doxastic commitments, and (2) *R* is not normatively superfluous and therefore unnecessary. Third, I posit the following material conditional:

> **(P3)** If (P1) and (P2) are true, then the Robust View is a viable and consistent theory of epistemic rationality.

I take (P3) to be uncontroversial. Hence, if (P1) – (P3) hold, then we get the following conclusion:

> **(C)** The Robust View is a viable and consistent theory of epistemic rationality.

Before proceeding, several clarifications are in order. The first clarificatory point concerns (P1). I admit that there can be alternative formulations of these requirements. I will not focus on these other formulations. I take it that my argument in this section can be modified to accommodate most alternative, wide scope formulations of these rules. Second, I am sympathetic to the idea that there could be additional requirements of epistemic rationality that I have not discussed. However, if there are any additional epistemic requirements, then it is not clear what these requirements might be. Moreover, it is not clear whether these additional epistemic requirements would pose any serious threat to the Robust View. Hence, I will set this issue aside. Finally, I should note that the idea that

rationality is constituted by both requirements and commitments is not a novel one. For example, Shpall has claimed that "requirements are not sufficient for an adequate theory of rationality" (Shpall 2013, p. 719). Shpall argues that rationality includes both requirements and "rational commitments" (Shpall 2013, pp. 719, 729 – 732). Nevertheless, the Robust View differs fundamentally from Shpall's account. First, Shpall endorses his own view of rational commitments. On my view, however, doxastic commitments are appropriately construed as *cognitive commitments* and *semantic commitments*. Second, our views diverge concerning the normativity and nature of commitments. For example, Shpall thinks that commitments are "pro tanto" and need not always be normatively "decisive" (Shpall 2013, p. 731). Contrary to this, I have argued that some commitments are not pro tanto. Cognitive commitments are decisive. Semantic commitments, however, are pro tanto and can be overruled (see chapter 4). Our two views are therefore distinct. With these clarifications in mind, let us turn to a defense of (P2).

First, it is fairly straightforward that epistemic requirements do not contradict (or militate against) our doxastic commitments. Consider, for example, Non-contradiction. This requirement is compatible with each doxastic commitment. For accepting Non-contradiction does not force us to abandon one or more of our doxastic commitments. This seems obvious once we recall the content of the cognitive commitment of Doxastic Veridicality. When one consciously forms a belief that *p*, this doxastic attitude engenders a normative commitment to regard *p* as being true, such that one cannot simultaneously believe not-*p* without first revising her doxastic states (see Coliva 2019).[186] Non-contradiction similarly insists that one should refrain from having contradictory beliefs.

Similar considerations apply to the requirement of Justification. This requirement dictates that one ought to believe *p* when one has a further belief that *p* is well-supported by one's evidence. A cognitive commitment of Doxastic Rationality normatively constrains one to believe in accordance

[186] I also discuss this in detail in chapter 2.

with one's evidence and to be sensitive and receptive to epistemic considerations. There is nothing inherently problematic with being under both of these normative constraints.

In a similar vein, the Closure rule insists that agents ought to believe propositions that deductively follow via Modus Ponens from their current beliefs. This is entirely compatible with semantic commitments that normatively constrain agents to believe the consequences of their beliefs when they are in the cognitive position to do so. Consider a situation in which an agent, *S*, believes *p* and believes *p* implies *q*. According to Closure, *S* ought to believe *q*. Insofar as *S* is in the cognitive position to believe *q*, she also takes on a semantic commitment to believe *q*. Even in unordinary or rare circumstances where *S* lacks the cognitive position to form a belief that *q*, she can still be rationally required to believe *q*. In this situation she merely lacks a semantic commitment to believe *q*. But this does not mean that she has a normative commitment to refrain from believing *q*. In this special situation, her lack of commitment is compatible with being under a rational requirement to form the relevant belief. Consequently, epistemic requirements are consistent with doxastic commitments.

Now, the most challenging part of a defense of (P2) is showing that epistemic requirements are not normatively redundant. Consider again the requirement of Non-contradiction. On the face of it, this requirement might seem redundant given Doxastic Veridicality. But I contend that this redundancy is merely apparent. Doxastic Veridicality governs consciously formed and occurrent attitudes of belief. When an agent, *S*, consciously forms the belief that *p*, then this occurrent doxastic state commits her to treat *p* as being true such that she cannot simultaneously treat not-*p* as being true (without revising her belief that *p*). On the other hand, Non-contradiction can be understood as a broad requirement that governs many kinds of doxastic attitudes. Specifically, it governs both occurrent beliefs and dispositional beliefs. A *dispositional* belief is taken to be an unconscious, background belief an agent possesses that can guide thinking and behavior (see Harman 1986, p. 6; Coliva 2019, p. 4; and chapter 1). Following Coliva, one does not merely acquire a dispositional belief

that *p* via a mental act of consciously judging that *p* is true (Coliva 2019, p. 4). Accordingly, we might say that one is *required* to make sure that one's total set of beliefs (occurrent and dispositional) are consistent and do not contradict with each other. The content of Non-contradiction differs from the content of Doxastic Veridicality. Consequently, Non-contradiction is not superfluous and redundant.

Let us now consider the requirement of *Justification*. One might insist that this requirement is redundant if beliefs engender *Doxastic Rationality*. This view, however, would not be correct. First, the requirement of Justification involves sets of beliefs. Specifically, it concerns a belief that *p* and its relationship to a second belief about the epistemic support for *p*. This requirement demands that an agent should form an additional state of believing *p* when she comes to possess a belief that *p* is sufficiently supported by her evidence. The commitment of Doxastic Rationality is different. If an agent, *S*, believes *p*, then this belief engenders the commitment of Doxastic Rationality to believe in accordance with one's evidence, and to be sensitive to new information concerning the truth or falsity of *p*. This commitment is not merely restricted to the coherence between sets of beliefs concerning propositions and their rational support. Rather, the content of Doxastic Rationality goes further by prescribing how one ought to be sensitive and responsive to evidence, and how one should modify one's beliefs in light of new information bearing on the propositional objects of one's beliefs. It is possible for an agent to believe *p* while satisfying Doxastic Rationality without having some further belief about the evidence supporting *p*. Consequently, Justification is not straightforwardly redundant.

Finally, the requirement of Closure differs from semantic commitments engendered by an agent's beliefs. For Closure concerns a specific type of inference rule, namely, logical deduction via Modus Ponens. An agent's semantic commitments, however, can concern a wider range of inference rules depending on her understanding, conceptual resources, and cognitive abilities. Semantic commitments are subjectively derived from one's cognitive position, whereas Closure is an external

rule that can apply to all agents in all situations. Hence, Closure is not redundant. It is a distinct normative rule that agents ought to comply with.

I have shown that the standard epistemic rules are consistent with our doxastic commitments. I have also shown that these rules are not normatively redundant. We can formulate a consistent theory of epistemic rationality that accommodates both requirements and doxastic commitments. If this is the case, then the Robust View is a viable and consistent theory.

4. The Normative Priority of Doxastic Commitments

According to the Robust View, theoretical rationality is comprised of both external rational rules and internal doxastic commitments. One question we might ask is the following: Are these two kinds of constraints on par with each other? That is to say, is one type of normative constraint more fundamental than the other? My aim in this section is to show (albeit briefly) that our doxastic commitments are normatively prior or more fundamental. Let us call this the "fundamentality thesis".

There are several reasons for endorsing the fundamentality thesis. First, doxastic commitments are constitutive elements of one's mental states and mental abilities. For instance, cognitive commitments partly individuate our attitudes of belief. They are decisive normative constraints that are engendered by the attitudes of beliefs themselves (see chapter 2). Similarly, one's semantic commitments are subjective constraints generated by one's beliefs, understanding, conceptual resources, and cognitive abilities (see chapter 3). Rules of rationality, however, are external constraints that agents should comply with. They are not dependent on one's beliefs or cognitive abilities. This suggests that our doxastic commitments are normatively prior.

Second, it is possible to view external rational requirements as extensions of our doxastic commitments. That is to say, one might plausibly contend that doxastic commitments serve as a normative basis for some of the rules of epistemic rationality. Consider the rule of Closure. According

to this rule, one is required to believe the propositions that follow via Modus Ponens from one's beliefs. This presupposes that a given agent is typically in a cognitive position to believe *p* when *p* follows via Modus Ponens from her current beliefs. This encapsulates the idea that Closure is a specific rule with a normative force stemming from the normativity of one's semantic commitments.

Similar considerations apply to Non-contradiction. This is a rule that can be broken. One might consciously believe *p* but have a dispositional belief that not-*p*. This rule could be based on the cognitive commitment of Doxastic Veridicality. First, our beliefs are conceptually tied to the truth (see chapter 1). If *S* consciously believes *p* at time, *t*, then she is disposed to mentally treat *p* as being true at *t*. Indeed, *S* cannot also consciously believe not-*p* at *t* unless she first revises her initial belief that *p*. This follows from Doxastic Veridicality. Second, an agent can combine her individual beliefs to form a more complex belief. Consequently, she will be in a position to regard the complex propositional content of this new belief as being true. However, agents can have conflicting beliefs. According to Coliva, an agent, *S*, might consciously believe *p* while having some background dispositional belief that not-*p* (Coliva 2019). Given that a dispositional belief can (in principle) be consciously recalled in one's thought, one should be able to combine the content of that belief with the content of one's other beliefs. If, however, one has a dispositional belief that not-*p*, and one consciously believes *p* at some later time while maintaining that dispositional belief, then one is not in a position to form a conscious conjunctive belief that (*p* & not-*p*). Doxastic Veridicality precludes the possibility of regarding a contradiction as being true (see chapter 2 and Coliva 2019). Non-contradiction can thus be understood as a requirement intimately connected to Doxastic Veridicality. The rule of Non-contradiction is an extension of Doxastic Veridicality that applies to a wider class of beliefs. This supports the claim that the cognitive commitment of Doxastic Veridicality is normatively prior.

The same might be said about Justification. Note that a given cognitive commitment of Doxastic Rationality will normatively constrain its possessor to be receptive to evidence and

information concerning the truth of a proposition. One specific way to satisfy this commitment and thus be receptive to this kind of epistemic information is to form a belief that *p* when one also believes that *p* is epistemically warranted. This is what the rule of Justification requires of rational agents. Hence, the rule of Justification can be understood as a specific requirement based on our commitments of Doxastic Rationality. If this is correct, then Doxastic Rationality would be more fundamental than the rule of Justification.

A third reason to accept the fundamentality thesis centers on the overriding force that our doxastic commitments seem to possess. To illustrate this idea, consider the Closure rule. According to this rule, any agent ought to believe *p* when *p* follows via Modus Ponens from her other beliefs. However, it is possible to conceive of a situation in which an agent is not able to make these simple kinds of deductive inferences. Perhaps this agent is under cognitive duress or lacks the cognitive resources to reliably follow the rule of Modus Ponens. In these situations, the agent is cognitively incapable of complying with Closure. However, if she is cognitively incapable of complying with this rule, then she is not in the cognitive position to employ Modus Ponens in her reasoning and belief-formation. Consequently, she lacks the semantic commitments to believe the propositions that follow via Modus Ponens from her current beliefs. In these special circumstances, it seems permitted or appropriate to excuse this agent from complying with Closure. This excuse is licensed given that there is a lack of a semantic commitment. Hence, our semantic commitments have a certain kind of normative overriding force. If this is right, then our doxastic commitments have a distinct normative priority that external rules of rationality do not enjoy.

The Robust View holds that theoretical rationality is comprised of doxastic commitments and external epistemic requirements. I have shown that we have good reason to think that the fundamentality thesis is correct. Hence, our doxastic commitments are normatively fundamental. This, however, does not mean that one is never obligated to abide by rules of epistemic rationality. One

ought to comply with both doxastic commitments and epistemic rules. Nevertheless, doxastic commitments seem to be normatively prior to epistemic requirements. Epistemic rationality is therefore best construed a hierarchy of normative constraints.

5. THE PLAUSIBILITY AND APPLICATION OF THE ROBUST VIEW

I have demonstrated that the Robust View is a viable and consistent theory. In this section I establish the claim that the Robust View is a more preferable theory of epistemic rationality. The argument for this claim can be expressed as follows.

(P1*) The Robust View has the theoretical power and resources to provide a more satisfactory and complete picture of epistemic rationality.
(P2*) The Robust View has the potential to provide some unique and plausible solutions to some problems in epistemic rationality.
(P3*) If (P1*) and (P2*) are true, then the Robust View is a more tenable and preferable theory of epistemic rationality.
(C*) The Robust View is a more preferable theory of epistemic rationality.

I take (P3*) to be fairly straightforward and obvious. In this section I attempt to defend both (P1*) and (P2*). I first support (P1*) by demonstrating that the Robust View (1) has the conceptual resources to accurately handle a wider range of cases than standard accounts of epistemic rationality, and (2) provides better normative guidance concerning how agents ought to maintain doxastic coherence and rational beliefs. In order to establish (P2*), I briefly discuss how the Robust View can provide some *prima facie* plausible solutions to certain problems facing standard accounts of rationality. If this is the case, then we have sufficient reason to accept the thesis that the Robust View is a more tenable theory of epistemic rationality.

5.1 DEFENDING (P1): ADVANCING A COMPLETE PICTURE OF THEORETICAL RATIONALITY*

I now turn to a defense of (P1*). It was shown that epistemic rules are not normatively redundant or superfluous. Indeed, the content of these rules are unique, and therefore differ from the contents of

doxastic commitments. It follows from this that countenancing both requirements and commitments affords us with a richer conceptual framework that adequately covers a wider range of situations.

For example, consider the fact that Closure only concerns those cases of believing the propositions that deductively follow via Modus Ponens from one's current beliefs. By also countenancing semantic commitments, epistemic rationality can accommodate other inference rules and methods of belief-formation (e.g., believing the inductive consequences of one's beliefs). The result of endorsing the Robust View is a more comprehensive picture of theoretical rationality.

Moreover, Doxastic Veridicality can be supplemented by the requirement of Non-contradiction. For Doxastic Veridicality is a commitment associated with consciously formed beliefs. Non-contradiction, however, extends the normative guidance of Doxastic Veridicality to all dispositional and non-occurrent beliefs that agents can accrue in tacit ways. An agent *S* might consciously come to believe *p*. On the Robust View, *S* can plausibly live up to her commitment of Doxastic Veridicality when she contemplates the target proposition of her consciously formed belief that *p*. In this case she would mentally regard *p* as obtaining. However, if *S* also has some unconscious dispositional belief that not-*p*, then *S* is not entirely rational. We can capture this intuition by appealing to the requirement of Non-contradiction. In this situation, *S* violates the Non-contradiction rule and should rationally abandon one (or both) of these beliefs. Appealing to both requirements and commitments can adequately handle these kinds of cases. This suggests that the Robust View provides a more satisfactory and richer picture of theoretical rationality.

The rule of Justification will similarly augment the cognitive commitment of Doxastic Rationality. To believe *p* and acquire a commitment of Doxastic Rationality is to be normatively compelled to be sensitive to the evidential basis one has for *p*. The rule of Justification specifies one way to satisfy this commitment: one's beliefs about the evidential basis for *p* should cohere with beliefs concerning the truth or falsity of *p*.

Finally, the cognitive commitment of Doxastic Integrity specifies a rational way for an agent to maintain her beliefs over time. Other things being equal, if an agent, *S*, believes *p*, then *S* takes on a commitment to be disposed and cognitively willing to regard *p* as being true until she rationally revises this belief. Doxastic Integrity is a normative constraint not covered by the standard epistemic requirements.[187]

These remarks support the view that appealing to both doxastic commitments and epistemic requirements will provide us with more normative guidance in a wider range of cases. Consequently, the Robust View provides a richer and more satisfactory picture of theoretical rationality. Therefore, (P1*) is true.

5.2 Defending (P2): Applying the Robust View*

The next two subsections provide a defense of (P2*). Some theorists have identified several worries facing wide scope views of rational rules. I focus on two of these worries. I do not discuss these issues in detail. Nor do I discuss all the relevant positions that philosophers have defended. To do so would exceed the scope of my project. Instead, I simply aim to show that the Robust View has the potential to resolve these two issues. I therefore contend that we have sufficient reasons to accept (P2*). While my treatment of these issues will be brief, it is my intention to build upon this work in the future.

[187] While working on cognitive commitments I discovered some literature by John Broome suggesting that there might be a rule similar to Doxastic Integrity (Broome 2013, pp. 185 – 186). Broome calls this rule, "*Persistence of Belief*", and formulates it as follows: "If t_1 is earlier than t_2, rationality requires of *N* that, if *N* believes at t_1 that *p*, and no cancelling event occurs between t_1 and t_2, and if *N* cares at t_2 whether *p*, then either *N* believes at t_2 that *p*, or *N* considers at t_2 whether *p*" (Broome 2013, pp. 185 – 186). Broome is not sure if this is the correct formulation (Broome 2013, p. 186). However, I am not convinced that there is an external wide scope rule like Persistence of Belief. Moreover, I am suspicious about the clause concerning one's caring whether *p* obtains. Broome endorses the clause because he thinks that one can rationally forget beliefs that they do not care about (Broome 2013, p. 186). But I disagree with this claim. For I take it that there is something always rational about retaining beliefs that one has not rationally revised. Hence, I prefer my account of Doxastic Integrity.

5.2.1 THE SYMMETRY PROBLEM

One issue facing wide scope theories of rational requirements is the "symmetry" problem (see Schroeder 2004; Way 2010; Way 2011; and Shpall 2013). There can be different ways to formulate and address this worry.[188] For the sake of brevity and simplicity, I will address the formulation that Jonathan Way provides. Moreover, I will focus on epistemic or theoretical rationality. My argument here does not concern the symmetry problem facing rules of pragmatic rationality. My contention is that the Robust View has the conceptual resources to resolve the symmetry problem.

First, wide scope rules are expressed in such a way that "rationality requires" ranges over a conditional statement. More precisely, these requirements typically take the following form: an agent, *S*, at time, *t*, is rationally required to see to it that if *S* has one mental state, M_1, then they also have or maintain a mental state, M_2. To comply with this rule, an agent must satisfy the conditional statement. However, as Way correctly points out, many theorists insist that wide scope rules are symmetrical: an agent can satisfy a wide scope rule by either (1) forming the mental state identified in the consequent of the conditional, or (2) abandoning the mental state identified in the antecedent of the conditional (Way 2010, p. 1060). The problem, however, is that there are cases in which one would not be rational simply by doing either one of these things – i.e., the proper satisfaction of the relevant conditional is not symmetrical (Way 2010, p. 1060). That is to say, in some cases it might be more rational to abandon the initial mental state identified in the antecedent. In other cases, the rational thing to do would be to form and maintain the mental state identified in the consequent of the conditional. But wide scope

[188] On formulation I will not consider is Shpall's (see Shpall 2013). As I understand Shpall, appealing to only wide scope rules leads to a problematic kind of symmetry: these rules (by themselves) cannot adequately account for the fact that there is an asymmetrical rational compulsion guiding an agent to satisfy the consequent of a requirement by forming the relevant mental state given that she already has the state specified in the antecedent of that requirement (Shpall 2013, pp. 722 -723). According to Shpall, this asymmetry is captured by certain commitments that mental states generate (Shpall 2013, pp. 727 – 728). I agree with Shpall to some extent: rationality includes rules and commitments. However, my view differs in that I defend and utilize different doxastic commitments.

requirements do not distinguish between which of these routes an agent should take. Way notes that, on the face of it, wide scope rules "are not good guides to escaping incoherence" (Way 2010, p. 1060).

We can illustrate this with an example. Suppose that an agent, *S*, irrationally believes *that my neighbor is the devil*. Call this belief, "B_1". Suppose further that *S* irrationally believes the conditional, *that if my neighbor is the devil, then my neighbor is not human*. Call this belief, "B_2" Now, according to Closure, *S* is rationally required to see to it that if they have these two beliefs, then they also have a further belief, *that my neighbor is not human*. Call this further belief, "B_3". The requirement is symmetrical. This means that *S* can rationally comply with this rule by either doing one of two things: she can drop one or both of the initial grounding beliefs (i.e., B_1 and B_2), or she can form B_3. However, this seems counterintuitive. *S*'s initial beliefs, B_1 and B_2, are irrational. Indeed, *S* should revise these beliefs. However, the wide scope rule of Closure is silent on this matter and does not provide satisfactory guidance. This worry extends to other wide scope rules.

According to Way, a proponent of wide scope rules might appeal to additional principles of rationality to provide the necessary guidance (Way 2010, p. 1061; Way 2011, pp. 232 – 234). However, the Robust View already has the resources to deal with the symmetry problem. Instead of appealing to additional epistemic principles, we simply comply with our doxastic commitments. In the case specified above, the agent should appeal to the commitment of Doxastic Rationality. Accordingly, if an agent believes *p*, then she is committed to believe in accordance with good evidence, and to be sensitive and receptive to information concerning the truth of *p*. Hence, if an agent's belief is based on weak or insufficient rational support, then she violates her commitment. In these cases, one should revise one's irrational belief. Given the case above, the agent, *S*, bases her beliefs, B_1 and B_2, on weak and insufficient rational grounds. She therefore violates her cognitive commitments. In order to maintain rationality and comply with Closure, she should satisfy her cognitive commitment of

Doxastic Rationality by dropping B_1 and B_2. By dropping these two beliefs she *ipso facto* complies with the wide scope rule of Closure.

My contention is that the proper way to satisfy wide scope rules of epistemic rationality is prescribed by one's doxastic commitments. In most cases, this normative guidance stems from the cognitive commitment of Doxastic Rationality. By satisfying one's commitments, one can appropriately comply with epistemic requirements. The Robust View of epistemic rationality countenances both commitments and requirements. Hence, the Robust View can satisfactorily handle the symmetry problem. Moreover, it provides the conceptual resources and normative guidance required for a tenable theory of theoretical rationality.

5.2.2 The Normativity Problem

Another important issue facing theories of rationality is the so-called, "normativity puzzle" (see Kolodny 2005; Broome 2005; Broome 2008; and Way 2010). In this section I briefly discuss some select arguments that Kolodny and Broome advance.[189] After discussing these particular arguments, I go on to show that the Robust View provides a unique solution to the problem. To be clear, my focus is on epistemic rationality. My proposed strategy may not work for the normativity problem associated with rules of practical rationality. Nevertheless, my argument here will help establish (P2*).

Some theorists have questioned whether or not rationality is actually normative (Kolodny 2005; Broome 2005). According to Broome, if a rational requirement, *R*, is normative, then *R* must provide agents with some reason to comply with it (Broome 2005, p. 325).[190] And according to Niko

[189] See also Broome (2005; 2008) for additional arguments.

[190] There are different ways to conceive of the reasons stemming from rational rules. Following Broome, we might say that the reason is "either sufficient or pro tanto" (Broome 2005, p. 325). A pro tanto reason for Broome is some reason that "plays some role in helping [you] determine whether or not you ought to" do some act, *x* (i.e., it need not be a sufficient or decisive reason to do *x*) (Broome 2005, p. 326). In what follows, I will merely argue that we have some instrumental pro tanto reason to comply with rational rules. In particular, I will attempt to establish the thesis that Broome refers to

Kolodny, there might be different kinds of reasons to comply with rules of rationality. One reason might be instrumental, in that complying with a rational rule is necessary to achieve some other end or goal that one should achieve (Kolodny 2005, p. 543). Another candidate reason Kolodny identifies is that complying with rules of rationality is an inherent aspect of being "a believer or agent" (Kolodny 2005, p. 544). However, Kolodny denies that we always have these kinds of reasons. First, he argues that there are cases in which satisfying a rational rule is not always instrumentally valuable (Kolodny 2005, p. 543). Moreover, he contends that a failure to live up to rational requirements does not automatically result in a forfeit of one's status as a believer and agent (Kolodny 2005, pp. 544 – 545). He concludes that we have reason to doubt that rational rules are normative (Kolodny 2005, p. 547).[191]

Broome has also advanced arguments that purport to show that rules of rationality might not always be normative. For example, he similarly contends that rational requirements will not "necessarily" provide agents with instrumental reasons to comply with them (Broome 2005, p. 332 – 333). Like Kolodny, he insists that there can be cases in which satisfying a rule of rationality will not lead you to achieve something you should achieve (Broome 2005, p. 333). If this is the case, then instrumental reasons and rules of rationality come apart. Consequently, rules of rationality do not always generate instrumental reasons and are thus not normative (Broome 2005, p. 333).

We can illustrate this with an example described by Kolodny (see Kolodny 2005, p. 543). First, Kolodny discusses a principle of epistemic rationality he refers to as "B-", which states: "Rationality requires one not to believe that p, if one believes that there is not sufficient evidence that *p*" (Kolodny 2005, p. 521). He goes on to give a case purporting to show that we do not always have instrumental reasons to comply with B- or other rules of rationality. He describes the case as follows:

---

as "*weak normativity*", which he expresses as follows: "Necessarily, if rationality requires *N* to *F*, there is a reason for *N* to *F*" (Broome 2008, p. 156).

[191] Kolodny ultimately defends the claim that "rationality is only apparently normative" (Kolodny 2005, pp. 513, 557-560).

> Suppose I believe that I lack sufficient reason that *p*, but believe that *p*. As things stand, I violate B-. Then I comply by not believing that *p*. It might be the case, however, that I have reason to believe that *p*. So being rational leads me to lose a belief that I have reason to have (Kolodny 2005, p. 543).

In this case, one does not have instrumental reason to comply with B-. Kolodny generalizes this and insists that we need not always have instrumental reasons to follow rules of rationality (Kolodny 2005, pp. 543 – 544).[192]

I want to restrict my focus to those arguments against the view that rational requirements provide agents with instrumental reasons. Specifically, I will focus on the epistemic rules of Non-contradiction, Justification, and Closure. Intuitively, it seems that epistemic rational requirements are normative. Indeed, we often think that people are incorrect or irrational if they fail to satisfy these rules. If, for example, an agent believes *p*, then we think that she is normatively compelled to refrain from believing not-*p*. We can preserve this intuitive idea if we endorse the Robust View. According to this view, epistemic rationality is composed of a system of epistemic rules and doxastic commitments. My contention is that epistemic rational requirements always provide some pro tanto instrumental reason to comply with them. I will show that complying with external rules of rationality will put one in a position where he or she will *not* automatically violate one or more doxastic commitments. Since we ought to satisfy our doxastic commitments, we always have an instrumental reason to comply with epistemic requirements. Thus, if we endorse the Robust View, then we can preserve the intuitive view that epistemic rules are normative. Let me explain this in more detail.

First, I argued that doxastic commitments are normatively prior or fundamental. Doxastic commitments are internal normative constraints engendered by our mental states. In some cases, our doxastic commitments can guide, direct, and override epistemic rules. Moreover, it has been demonstrated that agents always ought to live up to their doxastic commitments (see chapter 2).

[192] See Broome (2005) for similar examples.

Second, it can be shown that when an agent complies with a wide scope epistemic requirement, then they can put themselves into a position where they will not automatically violate one or more of their doxastic commitments. For instance, if an agent complies with the rule of Non-contradiction, then they *ipso facto* are in a cognitive position to live up to their cognitive commitment of Doxastic Veridicality. By not harboring multiple beliefs that have propositional contents that contradict each other, one can safely combine these beliefs in such a way to regard the resulting target proposition with conjunctive content as being true. Hence, they will satisfy Doxastic Veridicality. Moreover, by refraining from having individual contradictory beliefs, one is *ipso facto* in the cognitive position to satisfy Doxastic Rationality. From an individual agent's perspective, both *p* and not-*p* cannot be true simultaneously. Hence, agents will typically have sufficient epistemic reason not to believe the conjunction (*p* & not-*p*). If an agent does not have contradictory beliefs, then they are typically believing in accordance with their evidence. They are being appropriately sensitive and receptive to their evidence. Hence, they are living up to their commitments of Doxastic Rationality. Complying with Non-contradiction provides one with an instrumental reason: doing so puts them in a position where they will not automatically fail to live up to their doxastic commitments.

Similar considerations apply to the rule of Justification. By satisfying this rule, one automatically is in a position to satisfy one's commitment of Doxastic Rationality. Believing *p* when one believes that *p* is sufficiently supported is one way to be appropriately response to one's epistemic reasons. Thus, if one violated Justification, one would automatically fail to live up to Doxastic Rationality. In this situation, one would not be appropriately sensitive and receptive to evidence and information concerning the truth or falsity of the contents of one's beliefs. Consequently, satisfying Justification provides us with an instrumental reason to live up to Doxastic Rationality.[193]

[193] Recall that doxastic commitments are constitutive of beliefs. These commitments are inherent aspects of the attitudes themselves. Thus, to believe is to always have a constitutive reason to satisfy the associated doxastic commitments. However, my point here is that epistemic rational requirements

Finally, compliance with Closure will automatically safeguard one from violations of their semantic commitments. In many cases, agents typically have semantic commitments to believe the propositions that follow from their current beliefs via Modus Ponens. By complying with Closure, an agent will automatically fulfill the corresponding semantic commitments.

Now, in some rare cases, an agent, *S*, might not be in the cognitive position to reliably or correctly employ Modus Ponens in her reasoning. In this sort of case, *S* would lack a semantic commitment to believe in accordance with Modus Ponens. Nevertheless, if *S* comes to correctly believe a proposition, *p*, that does indeed follow from her current beliefs via Modus Ponens, then *S* will not violate any semantic commitment. This is because there is no semantic commitment to believe *p*. Hence, complying with Closure in these rare cases still puts one in the cognitive position where they will not automatically violate any of their doxastic commitments.

Adopting the Robust View seems to provide a solution to the normativity problem. Epistemic rationality is comprised of epistemic rules and doxastic commitments. By complying with these epistemic rules, an agent will put herself in a position where she will not automatically violate any of her normatively prior doxastic commitments. Being in this position is a valuable thing. For we ought to strive to live up to our doxastic commitments. Hence, if compliance with epistemic rules will ensure that one will not automatically violate a doxastic commitment, then this provides an instrumental reason to abide by these rules. Consequently, epistemic rules will be normative.

This is an intuitive result. Moreover, this view seems compatible with other accounts that purport to show that rules of rationality are normative.[194] I have merely shown that requirements of

---

provide a distinct instrumental reason to comply with them. For compliance with these rules will always partly facilitate successful satisfaction of one's doxastic commitments. This does not suggest that there cannot be other reasons (perhaps non-instrumental reasons) to comply with these rules.

[194] For one example, see Nicholas Southwood (2008). I should note that Southwood considers and denies the idea that commitments of belief might be involved in establishing the normativity of rationality (Southwood 2008, p. 16). However, the commitments he considers are different from the commitments discussed here. For further overview of this issue see Way (2010).

epistemic rationality provide pro tanto *instrumental* reasons to satisfy them. It is consistent with this view that these requirements might also generate some further, independent reasons (instrumental or intrinsic) to comply with them. At any rate, I contend that my arguments here establish premise (P2*).

It might be objected that I have only considered a small number of epistemic rules that partly constitute epistemic rationality. Once we specify all the requirements governing belief, it might be the case that some of these rules do not provide the same type of instrumental reason. That is to say, satisfying these additional rules do not necessarily put one in a position where they will not automatically violate their doxastic commitments. One such rule might be Harman's Clutter Avoidance principle (see section 1). According to this principle, we ought to avoid believing trivial things (Harman 1986, p. 12). It is difficult to see how complying with this rule generates the same kind of instrumental reason associated with the other epistemic rules.

I have two responses to this objection. First, I admit that there could be additional epistemic rules that I have not considered. These rules may not provide the same kind of instrumental reason that Non-contradiction, Justification, and Closure provide. Nevertheless, it does not follow that these additional epistemic rules cannot provide us with some further instrumental reason to satisfy them. Perhaps we have pragmatic reasons to comply with rules like the Clutter Avoidance principle. Second, given my current aim and goal here, it is enough to show that some of the standard or crucial epistemic requirements seem to always generate an instrumental reason. These rules are normatively connected to our doxastic commitments. This provides us with sufficient reason to think that the Robust View is plausible theory of rationality.

If my arguments are cogent, then premises (P1*) – (P2*) are true. If these premises are true, then the Robust View is a more tenable and preferable theory of epistemic rationality. It stands to reason, then, that we should accept the thesis that Robust View is a more tenable and preferable theory of epistemic rationality.

6. CONCLUSION

I have advanced the Robust View of epistemic rationality. On this view, epistemic rationality is comprised of external mind-independent epistemic rules and internal mind-dependent doxastic commitments. I argued that this account was a viable and consistent theory. I further argued that this theory provides a more tenable picture of epistemic rationality. By countenancing both rules and commitments we are afforded with the conceptual resources to accurately handle a wide variety of cases. The Robust View provides a more detailed and satisfactory picture of the normativity that guides proper belief-formation and doxastic maintenance. Moreover, the Robust View has the potential to provide plausible solutions to the symmetry problem and the normativity problem. Consequently, the Robust View presents itself as a more preferable account of epistemic rationality.

VI. Species of Understanding

Adam C. Sanders

0. Introduction

Understanding is an epistemically valuable cognitive state. But the nature of this state has not been fully explicated. One common view holds that understanding is a "cognitive achievement" that is essentially based on justified true beliefs (Kvanvig 2003; Grimm 2006; Pritchard 2014). Let us call this the "JTB" model of understanding. Recently, however, many theorists have challenged the JTB account. Some deny that understanding requires justification (Hills 2016; Dellsén 2016). Others reject the claim that truth is a necessary condition for understanding (Elgin 2017). And some contend that understanding does not require belief (Wilkenfeld 2017; Dellsén 2016; Elgin 2017). Intuitions differ here. But one important question worth considering is whether these alternative approaches (if successful) deliver the same *kind* of understanding one achieves on the JTB model. With this question in sharp focus, this chapter aims to further our understanding of *understanding*. I advance the view that understanding can come in different species: veridical understanding, conjectural understanding, and narrative understanding. Each species, I contend, requires a further cognitive ability of *epistemic discernment*. Let us refer to this as the "Multiple Species" (MS) view of understanding.

On the MS view, *veridical understanding* is based on justified true beliefs. It is a strong cognitive achievement, based on strong rational support, that is not typically bound to a context for pragmatic considerations. *Conjectural understanding* is generated by epistemically justified true (or approximately true) acceptance. It is properly construed as a moderate cognitive achievement that can be restricted to a context given certain pragmatic considerations. *Narrative understanding* is a weak cognitive achievement generated by belief, acceptance, or supposition. Narrative understanding need not be factual. Nor need it be epistemically justified. It is simply the understanding one achieves when one competently discerns how the elements of an account or theory, Φ, cohere together, regardless of

whether Φ is true or believed. Each species of understanding, I contend, requires a further ability of epistemic discernment. Roughly, epistemic discernment involves the ability to competently discern some epistemically significant factor concerning the relationship between pieces of information, propositions, or facts.

Marking the distinction between veridical, conjectural, and narrative forms of understanding is important for several reasons. First, it affords us with the conceptual machinery needed to accommodate a wide variety of cases involving understanding. Second, it helps clarify the relationship between understanding and knowledge. Third, it elucidates certain aspects of the current disagreement over the nature of understanding. It will be shown that much of the disagreement boils down to a semantic dispute over the meaning of "understanding". The MS model can countenance and make sense of these alternative views of understanding. Hence, the MS view is an explanatorily powerful theory that preserves many intuitive elements associated with other, alternative accounts of understanding.

The chapter is structured as follows. In the first section I briefly discuss the robust nature of understanding and the kind of cognitive abilities typically associated with this state. In the second section I advance the MS view. Section 3 compares the MS model to other accounts of understanding that similarly reject the JTB model. I raise several challenges to these views and argue that the MS model is a more plausible theory of understanding. In section 4 I conclude the chapter. I briefly discuss the role of cognitive commitments, the relationship between understanding and knowledge, and a taxonomy of cognitive achievements.

1. Understanding as a Robust Epistemic State

According to Kvanvig, we can draw a distinction between "objectual" understanding and "propositional" understanding (Kvanvig 2003, p. 189). *Objectual* understanding concerns an object like

a theory, view, or subject matter (Kvanvig 2003, pp. 189, 192). For example, one can objectually understand subject matters like molecular biology, trigonometry, and nautical navigation. On the other hand, *propositional* understanding, like understanding *why*, *how*, *what*, and so on, involves understanding a certain fact or how a proposition obtains (Kvanvig 2003, p. 189).

Many theorists working on propositional understanding typically focus on understanding *why* (Pritchard 2014, 2016; Grimm 2014; Hills 2016). One commonly shared view is that understanding why *p* involves understanding a causal relationship between *p* and its cause, *q* (i.e., *q* provides a true explanation of *p*'s occurrence) (see, for example, Pritchard 2014 and Grimm 2014). An agent might, for example, correctly believe on good epistemic grounds that the titanic sank because it struck an iceberg. And in optimal situations, this agent might be credited with understanding *why* the titanic sunk. (I say more about the conditions for understanding later.)

But understanding why need not be restricted to causal relations. We can also understand why something *ought* to be the case. In this sense, understanding why might involve *reasons*.[195] According to Lord and Maguire, if one has "decisive" reasons to perform *F*, then these reasons outweigh any competing reasons to not do *F* (Lord and Maguire 2016, p. 10).[196] And if one has decisive reasons to do *F*, then one is normatively required to bring about *F* (Lord and Maguire 2016, p. 9). One could therefore understand why *p*, where "*p*" specifies the performance of an act, *F*, such that one correctly believes on good rational grounds that there are decisive reasons for *p*.

Indeed, moral understanding seems to be intimately connected to reasons. Alison Hills, for instance, claims that if "you understand why X is morally right or wrong, you must have some

[195] I will follow Scanlon here and take a *reason* for performing some *F* to be "a consideration that counts in favor of" doing *F* (Scanlon 2000, p. 17). Reasons, as I understand them, can be epistemic considerations, pragmatic considerations, and moral considerations that factor into what we ought to believe or do.

[196] Following Raz, reasons for action can be "stronger or more weighty than others" (Raz 1975, p. 25).

appreciation of the reasons why it is wrong" (Hills 2009, p. 101).[197] For example, Alice might correctly believe that the fact that a child is drowning provides a strong reason to save that child. And if this reason is not outweighed by any other competing reasons or obligations, then that reason is decisive: it normatively determines what one ought to do, namely, save the drowning child. In this case, if Alice believes that there is a decisive reason to save a drowning child, then (in optimal cases) we might attribute Alice with understanding why. Alice could come to understand why something ought to be the case. If this is correct, then understanding why need not involve causal relations. For one can understand that there are decisive reasons to perform some action and yet fail to perform that action.

We can generalize this to cases involving practical reasons and epistemic reasons. Suppose that Zeke is planning a trip to hike Mt. Fuji. Zeke has decisive or conclusive epistemic reasons supporting the truth of the proposition that Mt. Fuji is southwest of Tokyo.[198] Call this proposition, "*p*". If Zeke recognizes or believes that he has good epistemic reasons for *p* that *sufficiently* outweigh any reasons for not-*p*, then Zeke can, all things being equal, be attributed with understanding why one ought to believe *p* if one is in a similar context with similar evidence. Similar considerations apply to cases involving pragmatic understanding. For example, while planning his trip, Zeke might believe on good rational grounds that it will rain during his hike up to the summit of Mt. Fuji. And if Zeke intends to stay dry, he has a pragmatic reason to bring a jacket. If Zeke recognizes that this reason is decisive, then Zeke could come to understand why he ought to take a jacket.

Now, understanding seems to require something more than having a justified true belief (or some belief-like attitude). Consider a case in which I come to possess a justified true belief via reliable

---

[197] According to Hills, moral understanding requires more than having a belief that there are moral reasons for performing an act, F: one must "grasp" those reasons, where grasping is constituted by a set of abilities to provide explanations, follow explanations, and draw correct conclusions regarding the content of what is understood (Hills 2009, p. 102). I say more about grasping later.

[198] Recall that decisive reasons to believe *p* will outweigh the reasons not to believe *p*. In these cases, decisive reasons for *p* sufficiently warrant a rational belief that *p*.

testimony that my plants in the garden generate chemical energy because they undergo the process of photosynthesis. I might know this explanation or fact. But even though I know this explanation or fact, I might still fail to genuinely understand it. In addition to belief, many theorists insist that understanding requires some special cognitive ability.

One popular view maintains that understanding involves an element of *grasping* (Grimm 2006, 2014; Hills 2016; Elgin 2017). According to Grimm, *grasping* the relationship between propositions that are causally or explanatorily related is an ability to correctly comprehend relations of "dependencies" that exist or obtain between those propositions (Grimm 2006, pp. 532-533). Moreover, he claims that grasping involves the "ability to answer "what-if-things-had-been-different?" questions" (Grimm 2006, p. 532). Hills takes this idea further. She claims that understanding is constituted by what she calls "cognitive control", which is a set of abilities to manipulate the relationship between propositions, provide explanations, and draw correct inferences (Hills 2016, p. 663). And Elgin argues that the fundamental abilities associated with understanding are not simply capacities or dispositions to provide correct explanations, but rather the ability to competently use the relevant information to achieve one's epistemic ends (Elgin 2017, pp. 46, 57). For instance, she claims that an agent's understanding of some account, Φ, must go further than merely recognizing how the elements of Φ are related: an understander "should also be both willing and able (and perhaps be aware that she is willing and able) to use that information – to profess it, to reason with it, to apply it, perhaps to use it as a source of working hypotheses about related matters when her ends are cognitive" (Elgin 2017, p. 46).

It is plausible that understanding is intimately related to the possession of certain cognitive abilities. If one genuinely understands some phenomena or account, then one likely has an ability to competently employ the content of what is understood in reasoning, inference, and explanation. For example, if I genuinely understand that my plants in the garden generate chemical energy because they

undergo the process of photosynthesis, then perhaps I must grasp how photosynthesis causally relates to the production of chemical energy in plants. But determining which set of abilities are necessary for understanding is a contentious matter.[199] I will not provide or defend a detailed account of all the abilities that are necessary and sufficient for understanding. Rather, I contend that understanding requires a minimum basic cognitive ability that I refer to as "epistemic discernment".

At a first approximation, *epistemic discernment* is based on Pritchard's idea that an agent must have an "epistemic grip" when she understands why *p* (Pritchard 2014, p. 316). He argues that if one truly understands why *p*, where *p* is caused by some *q* (or where *q* is the correct explanation of *p*) then one must be able to competently conceive of how *q* can bring about or explain *p* (Pritchard 2104, p. 316). This seems correct. However, understanding need not involve causal relations. One can understand why something ought to be the case, where this understanding is based on *reasons* and not *causes*. Moreover, one can have objectual understanding of some subject matter that makes no appeals to causes or explanations. Hence, we need a more robust notion of *epistemic grip*.[200]

What is minimally required for understanding is *epistemic discernment*, which is a general ability to adequately and competently discern how some epistemically significant factor holds between propositions, information, reasons, or facts. What counts as being *epistemically significant* will depend on the content of what is understood. In cases of understanding causal relations, one must be able to competently discern how causes are related to their effects. In cases of understanding why *p* ought to be the case, where *p* ought to obtain because of some decisive reason, *R*, one must competently discern how *R* normatively requires one to bring about *p*. And in cases of objectual understanding, where one

[199] See Pritchard (2014) for a critique of Grimm's notion of *grasping*. And see Elgin (2017) for a critique of accounts that appeal to abilities or dispositions to competently provide correct explanations.
[200] To be clear, my remarks here are not meant to be understood as a critique of Pritchard's account. On the contrary, Pritchard is concerned with a specific type of propositional understanding involving causal relations. I am merely extending Pritchard's notion of *epistemic grip* to make sense of the cognitive ability that might be required for a general account of understanding.

understands an account, theory, or subject matter, Φ, the kind of epistemic discernment required is an ability to comprehend how the elements of Φ cohere together and depend on one another. This fits with Kvanvig's account of grasping associated with objectual understanding. According to Kvanvig, *grasping* the elements of some theory, account, or subject matter requires a comprehension and recognition of the way those elements cohere together and represent or explain some phenomena or state-of-affairs (Kvanvig 2003, 192). If one can competently discern how the elements of an account cohere, and how the account represents some phenomena, then one exercises epistemic discernment.

Epistemic discernment is a necessary feature of understanding. Consequently, understanding is a robust epistemic state that goes beyond the mere possession of having certain representational states like belief. This does not suggest that there cannot be other cognitive abilities that are necessary for understanding. One might, for example, insist that understanding also requires certain abilities to provide correct explanations, and to use the content of what is understood in reasoning. For my purpose I will set this issue aside. I will proceed by characterizing and defending the MS model of understanding. On this view, understanding comes in different species.

2. Species of Understanding

According to the MS model that I advance here, understanding comes in different species. In this section I characterize veridical understanding, conjectural understanding, and narrative understanding. Both propositional understanding and objectual understanding can take veridical and conjectural forms. Narrative understanding is a special species of objectual understanding. For the sake of brevity, I first focus on understanding *why* (hereafter "understanding" unless otherwise noted). But my arguments apply *mutatis mutandis* to objectual understanding.

First, understanding requires some kind of cognitive endorsement. If, for example, an agent, *S*, genuinely understands-why *p* (where *p* is caused or correctly explained by some *q*), then *S* will

cognitively endorse some complex proposition of the form *that p because q*. Similar considerations apply to objectual understanding. For example, according to the JTB model, understanding requires a justified true belief. Belief, then, is the relevant cognitive endorsement. However, it can be shown that this model does not capture all the different kinds of understanding that agents can enjoy. Different cognitive endorsements can engender different species of understanding when one can also competently exercise the ability of *epistemic discernment*. The kind of understanding one possesses will therefore depend on the type of cognitive endorsement involved, the rational considerations involved, and the way in which epistemic discernment can be competently exercised.

Before sketching out the details of the MS model, I must first characterize the notions of *belief*, *acceptance*, and *supposition* that I will appeal to.

A belief is a representational state. To believe that *p* is to represent some state-of-affairs as obtaining. I take it that belief has the following features:

(B1) For an agent, *S*, to believe *p* at some time, *t*, is for *S* to be disposed to regard *p* as being true, or being the case, or being more likely than not-*p* at *t*, in such a way that *S* cannot be agnostic toward the truth of *p* at *t*.

(B2) An agent's rational belief that *p* is responsive and sensitive to evidence and epistemic considerations for or against the truth of *p*, regardless of nonevidential or pragmatic reasons that agent has in a specific context.

(B3) An agent's consciously formed belief that *p* is typically the product of a mental act of judging that *p* is true, or that *p* is the case, or that *p* is more likely than not-*p*, such that she cannot voluntarily will herself to rationally believe *p* by disregarding the evidence and epistemic reasons for the truth of *p*.[201]

In what follows I will understand belief as involving features (B1) – (B3). While (B1) – (B3) might not be jointly-sufficient for belief, I take it that they are at least necessary features.

[201] To be clear, I am not providing a novel account of belief. Many theorists commonly hold that belief involves some or all these features. For the sake of brevity, I will assume this conception defended by theorists elsewhere. For more information on feature (B1) and the connection between belief that *p* and considering *p* to be true, see Wayne Davis (1988) p. 170. For more detail concerning feature (B2) and beliefs sensitivity to rational support, see Duncan Pritchard (2016) p. 90. For more on feature (B3) and the involuntary judgments associated with belief, see Coliva (2016) pp. 33-34 and Bratman (1999) p. 18.

Acceptance can be distinguished from belief. There are different ways that philosophers have drawn this distinction.[202] My account of acceptance draws on the work by L. J. Cohen and Michael Bratman. Cohen defines *acceptance* as follows:

> …to accept that *p* is to have or adopt a policy of deeming, positing, or postulating that *p* – i.e., of including that proposition or rule among one's premisses for deciding what to do or think in a particular context, whether or not one feels it to be true that *p* (Cohen 1992, p. 4).[203]

Moreover, it is commonly thought that acceptance is voluntary: an agent voluntarily chooses or decides to accept a proposition (Cohen 1992, p. 22; Bratman 1999, p. 27). And one can accept a proposition for either evidential or pragmatic reasons (Cohen 1992, pp. 4-5, 12; Bratman 1999, pp. 20-21). This is an important point. Acceptance might involve a mixture of pragmatic and evidential considerations. But as Bratman correctly notes, acceptance is typically context-dependent: one accepts a proposition, *p*, based on *some* pragmatic considerations within a given context (Bratman 1999, p. 28).

On my view, for an agent to accept *p* involves (1) taking some positive mental stance toward *p*, and (2) having a mental resoluteness to posit and implement *p* in thought or reasoning, given her available reasons for *p*. Acceptance has the following features:

(A1) To accept *p* is to implement *p* in thought and to be disposed to adopt some positive stance toward *p*, such that one can still be agnostic toward the truth or falsity of *p*.

(A2) A rationally held attitude of accepting *p* is based on reasons. These reasons can be either epistemic or practical considerations in favor of *p*.

(A3) An acceptance that *p* can be the result of a voluntary choice or decision: an agent can will herself to accept *p* in such a way that she ignores the epistemic considerations for or against *p*.[204]

---

[202] On my own view, there are different species of acceptance. But for the sake of argument, I will only appeal to one species of acceptance. The arguments for the TS model should apply *mutatis mutandis* to other species of acceptance.

[203] It should be noted that Cohen defines *belief* that *p* as "a disposition…normally to feel it true that *p* and false that *not-p*, whether or not one is willing to act, speak, or reason accordingly" (Cohen 1992, p. 4). This differs from the characterization of *belief* that I provide. While it might be true that human beings typically have dispositions to feel as though propositions are true or false, it does not seem plausible to say that belief must always involve a disposition to feel a certain way.

[204] On this view, Cohen's account of acceptance would be one species of acceptance. For premising a proposition is one particular way of utilizing a proposition in thinking and reasoning.

Features (A1) – (A3) specify the basic characteristics of the notion of *acceptance* that I appeal to here.

Finally, *supposition* is a cognitive endorsement with similar features to that of acceptance. Following Cohen, supposing that *p* typically involves "an inherently temporary act of imagination" (Cohen 1992, pp. 12-13). One need not regard the target proposition of a supposition to be true. Moreover, supposing that *p* need not be grounded in good reason: one need not consider belief or acceptance in the target proposition to be warranted. As Nicholas Rescher puts the point, supposition can be "belief-contravening": a supposition that *p* is *belief-contravening* if *p*'s truth-value is not known, or if *p* is either believed or known to be false (Rescher 1964, pp. 3-4). Furthermore, Bratman correctly notes that suppositions can differ from acceptance considering that they typically do not "directly shape [one's] action" in a context: accepting that *p* in a context is usually "tied more directly to practical reasoning" and action than merely supposing that *p* in some context (Bratman 1999, p. 28). Finally, some suppositions, I contend, need not have *any* rational grounding (epistemic or pragmatic). These are *arational* suppositions, and they are more akin to states of imagining that simply occur.

For my purposes here, it will suffice to maintain that suppositions have the following essential features:

(S1) For an agent, *S*, to suppose a proposition, *p*, involves a weak and tentative fixing of *p* as being *given* in thought or reasoning, such that (1) *S* need not commit herself to implement *p* in her planning, thinking, and reasoning, and (2) *S* need not take a positive mental stance toward *p* – i.e., *S* need not regard *p* as being true, likely, possible, credible, useful, or practically beneficial for her decisions and plans.

(S2) Suppositions are sensitive and responsive to evidential and practical considerations. Rational suppositions can be based on either evidential or pragmatic reasons. Arational suppositions need not be based on any reasons.

(S3) An agent's supposition that *p* can be the product of a willful and voluntary decision (or intention) to weakly and tentatively adopt, posit, or fix *p* in thought within a specific context, such that she need not have to take into consideration the available epistemic reasons for or against *p*.

What is important to note here is that belief can fail to coincide with either acceptance or supposition. Cohen, for instance, correctly claims that one can fail to believe that *p* while nevertheless

coming to *accept p* (Cohen 1992, p. 5). I might, for example, accept that it will rain today without believing that this proposition is true. Perhaps I have some evidential reasons supporting the truth of this proposition that I do not consider to be sufficient to warrant belief. But given additional pragmatic considerations within this context, I can voluntarily accept that it will rain today. Similar considerations apply to suppositions. Consequently, belief can be distinguished from both acceptance and supposition.

I will argue that belief, acceptance, and supposition can generate different species of understanding. Possessing one species of understanding over another depends on the kind of cognitive endorsement involved, the rational support for that endorsement, and the ability to exercise epistemic discernment. Let us discuss each species in turn.

2.1 VERIDICAL UNDERSTANDING

Veridical understanding coheres with the core tenets of the JTB account. If one possesses veridical understanding, then one's understanding is generated by a justified true belief. Hence, if an agent veridically understands why *p*, where *p* is caused by some *q*, then that agent possesses a justified true belief *that p because q*. Moreover, she can epistemically discern how it is that *q* causally brings about *p*.

Several clarificatory points are in order. First, to believe *p* is to regard *p* as being true. If an agent veridically understands why *p* by forming a belief that *q* causes *p*, then she is committed to regard this causal relationship as being true.

Second, a belief that generates veridical understanding is rationally grounded by available evidence or epistemic reasons. As Pritchard puts the point, if *S* understands that *q* caused *p*, then *S* will typically possess "good reflectively accessible grounds in support" of her true belief that *q* caused *p* (Pritchard 2009, p. 39). Similar considerations apply to understanding why *p* ought to be the case.

One cannot veridically understand why *p* without having good evidence in support of the relevant belief. In this way, veridical understanding involves *strong rational grounding*.

Third, if one veridically understands why *p*, and this understanding is based on a belief supported by strong rational grounding, then this understanding is not primarily dependent on pragmatic considerations one has in a context.[205] Veridical understanding can persist throughout different contexts insofar as one retains the relevant belief that grounds the understanding.

Fourth, a belief generating veridical understanding must be true. Veridical understanding involves getting things right. Following Pritchard, the belief that generates this understanding is a "cognitive success" (Pritchard 2014, p. 319). One correctly understands how some phenomena occurs, or how things ought to be. One cannot possess veridical understanding if the associated belief is false.

Finally, when an agent possesses veridical understanding, she can competently exercise the ability referred to as *epistemic discernment*. If, for example, one understands why *p* is the case, then one can competently discern how *p* relates to its cause (or how some *q* correctly explains *p*).

Now, understanding is typically characterized as a sort of cognitive achievement (Kvanvig 2003; Elgin 2009, 2017; Pritchard 2014). According to Pritchard, when one's cognitive success (or true belief) is based on one's "cognitive abilities" and sensitivity to good, reflectively accessible epistemic reasons in favor of the proposition, *that q caused p*, then the resulting understanding is inherently a type of "cognitive achievement" (Pritchard 2014, p. 319). Considering that veridical understanding is an epistemic state based on a true belief supported by strong rational grounding, it is appropriate to characterize this species of understanding as a *strong cognitive achievement*. It is a cognitive achievement in the sense that the relative cognitive endorsement (i.e., belief) is correct and gets things right. Considering that belief is involved, one correctly commits to the truth of what is understood.

[205] One might argue that understanding why an action should be performed will crucially depend on pragmatic considerations for that action. This is true. However, the point here is that the justification of the underlying belief should be *rationally* connected to good epistemic reasons.

Moreover, one appropriately bases the relevant belief on good rational grounds while possessing the cognitive ability of epistemic discernment. Hence, the cognitive achievement is a strong one.

2.1 CONJECTURAL UNDERSTANDING

Conjectural understanding is a distinct species of understanding. Like veridical understanding, conjectural understanding requires a cognitive endorsement. But the cognitive endorsement is not a belief. Typically, conjectural understanding involves an attitude of acceptance.

Consider the following case:

> ***Tortoise:*** Brooke is investigating the tortoise population on a remote island. She correctly believes that the population is increasing. She decides to test a hypothesis claiming that the increase in the island's tortoise population is due to an increase in annual rainfall. Call this hypothesis, "*h*". Brooke adequately comprehends how an increase in rainfall can cause an increase in the tortoise population. However, she has insufficient evidence to rationally warrant a belief regarding the truth or falsity of *h*. Nevertheless, she still has some evidence for *h*'s truth. Moreover, she correctly recognizes that *h* is a plausible and likely explanation for the increase in tortoise reproduction. Based on these considerations and her goal of inquiry, Brooke accepts *h*. In this case, *h* is true.

In the *Tortoise* case, Brooke has some sort of understanding. She accepts a true hypothesis, *h*, and she exercises epistemic discernment. I contend that the best way to make sense of this case is to attribute Brooke with *conjectural understanding*. Unlike veridical understanding, the possessor of conjectural understanding need not regard the content of what is understood to be true. This is due to the fact that acceptance can come apart from belief. One can accept things one does not believe. An acceptance generating conjectural understanding functions in such a way that one posits the content or object of what is understood in one's thought, inquiry, or reasoning.

Moreover, conjectural understanding typically involves a true acceptance.[206] In the *Tortoise* case, Brooke accepts that the increase in tortoises was caused by an increase in rainfall. Her acceptance is

[206] It can also involve approximately correct acceptances. I explain this when discussing Elgin's view.

correct because its content is true. Similar considerations apply to understanding involving reasons. If one conjecturally understands why some *p* ought to be the case given decisive reasons, *R*, then one must form a true acceptance that *R* normatively requires one to bring about *p*. Understanding why typically involves getting things right. It is incorrect to insist that understanding why can be false.[207] Hence, conjectural understanding why is factive.

Furthermore, conjectural understanding must be appropriately based on epistemic reasons, even if these epistemic reasons are not sufficient to warrant rational belief. Otherwise we get the counter-intuitive result that one could understand why *p* is the case by correctly accepting *p because q* for purely pragmatic reasons. Indeed, if acceptance can come apart from belief, then one could accept that *p because q* while simultaneously believing on sufficient epistemic reasons that it is *not the case that p because q*. This, however, does not suggest that one has understanding why. Understanding why based on purely pragmatic reasons is more plausibly characterized as a kind *objectual understanding*: one merely understands a story or theory about the cause of some *p* that one has insufficient evidential reasons to accept or believe (I say more about this in the next subsection). Conjectural understanding requires the corresponding acceptance to be based on good epistemic reasons. At minimum, conjectural understanding requires rational doxastic agnosticism.

Even though conjectural understanding is rationally based on epistemic reasons, pragmatic considerations can play some role in accepting that something is the case (or that something ought to be the case). Unlike belief, acceptance is a voluntary affair that results from a decision to posit a proposition in thought, reasoning, or inquiry. If one has epistemic reasons for *p* that are not decisive to warrant rational belief that *p* is true, then one can decide to accept *p*. The decision is voluntary. And in many cases, the decision to accept *p* will be partly informed by pragmatic considerations. Typically,

[207] It might be plausible to say that some kinds of *objectual* understanding can contain falsehoods (see Elgin 2017). I will say more about the factivity of objectual understanding in section 3.

pragmatic considerations are goal directed. According to Elgin, we might accept something to promote our "cognitive ends", such as pursuing some line of inquiry (Elgin 2017, p. 19). Consider again the *Tortoise* case above. Brooke accepts a hypothesis, *h*, that specifies a causal relationship (e.g., *p* is caused by *q*). Brooke accepts *h* based on her epistemic reasons in conjunction with certain pragmatic considerations she has in her present context, *C*. These pragmatic reasons might be directed at her goal of conducting inquiry concerning the relationship between rainfall and tortoise births. If pragmatic reasons in *C* factor into her decision to accept *h*, then her acceptance can be restricted to *C* (and similar contexts where she has these same reasons for accepting *h*). Outside of these contexts Brooke might not accept that the increase of tortoises is caused by an increase in rainfall. Unlike veridical understanding based on belief, conjectural understanding based on acceptance can be bound to a specific set of contexts for pragmatic reasons.

The foregoing remarks suggest that conjectural understanding involves *moderate rational grounding*: an attitude of acceptance that generates conjectural understanding is appropriately based on good epistemic reasons. These epistemic reasons need not be decisive to warrant rational belief. But conjectural understanding based on acceptance cannot be achieved when one has decisive epistemic reasons against the truth of what is accepted. And in many cases the relevant acceptance is partly dependent on pragmatic considerations one has in certain contexts. Hence, conjectural understanding lacks the evidential support that veridical understanding enjoys.

Finally, conjectural understanding seems to involve the same kind of cognitive abilities associated with veridical understanding. If one forms a true acceptance that *p because q* in some context, and if this acceptance can generate conjectural understanding, then one ought to be able to competently and accurately comprehend how *q* causes or explains *p*. Conjectural understanding requires epistemic discernment.

In light of these remarks, conjectural understanding is a type of cognitive achievement. This point may conflict with other views concerning the nature of cognitive achievements. However, characterizing conjectural understanding as a weaker form of achievement is plausible considering that conjectural understanding is a kind of cognitive success based on epistemic considerations. First, it is typically generated by a true acceptance. Indeed, one gets things right. Second, it is (weakly) rationally grounded in epistemic reasons. Third, it is the result of an agent's cognitive abilities. Specifically, it requires agents to be sensitive and responsive to the epistemic reasons for or against the target propositions that they accept. Moreover, conjectural understanding requires *epistemic discernment*. When the abilities of epistemic discernment and rational sensitivity are appropriately exercised, one's attitude of acceptance can generate understanding. But given the type of attitude and degree of epistemic support involved, conjectural understanding is properly construed as a *moderate cognitive achievement*. Unlike veridical understanding, conjectural understanding may be restricted to certain contexts for pragmatic reasons. Nor is one categorically committed to the truth of the content of what is understood. Hence, the cognitive achievement associated with objectual understanding is weaker than the cognitive achievement associated with veridical understanding.

2.3 NARRATIVE UNDERSTANDING

Narrative understanding differs from both veridical and conjectural understanding. It will be shown that narrative understanding is a special kind of objectual understanding. Consider the following case.

> ***Conspiracy:*** Alex believes on good epistemic grounds that human beings landed on the moon in 1969 because NASA's Apollo program placed them there. His evidence for his belief is based on textbooks, pictures, and video footage depicting American astronauts on the moon. Alex comes to learn about a conspiracy theory claiming that the Apollo moon landing never occurred. According to this theory, the moon landing in 1969 was a hoax propagated by NASA and the U.S. government. Alex momentarily supposes that the theory is true and can comprehend how the moon landing could be faked if the government manufactured the available evidence. Nevertheless, Alex does not believe that the moon landing was faked. He does not think there is any good evidence to support this conspiracy theory.

In the *Conspiracy* case, it seems that Alex possesses some form of understanding regarding the theory that the moon landing was a hoax. His understanding is not based on belief. And his understanding is not based on acceptance that is rationally grounded in good epistemic reasons. Nevertheless, he understands a theory that states that the moon landing never occurred because NASA and the U.S. government faked the landing and manufactured misleading evidence. Alex has a type of *objective* understanding I refer to as "narrative understanding".

Narrative understanding can take the form of *p because q* or *p ought to be the case because of q*. But even if it takes these forms, narrative understanding is not a type of understanding *why*. Rather, it is more appropriately classified as a type of *objectual* understanding. What is understood is a theory, account, or subject matter. In the *Conspiracy* case, Alex does not understand why the moon landing was faked. He understands a story or account claiming that the moon landing was a hoax.

Now, there seems to be two distinct forms of narrative understanding: hypothetical understanding and misguided understanding. *Hypothetical* understanding can be characterized as the understanding of a theory or account that one does not think is true. *Misguided* understanding is a kind of understanding one might possess when one falsely believes a theory or account. Let me explain.

One can hypothetically understand a theory when one *knows* that the theory is false. Consequently, hypothetical understanding need not be factive. Moreover, hypothetical understanding need not be based on good epistemic reasons. One can know, for example, that there are no good reasons supporting the theory of phlogiston and still understand this theory.

The attitudes most appropriate for generating hypothetical understanding are epistemically unjustified acceptances and suppositions. For example, a consequentialist might accept a deontological theory for purely pragmatic reasons. Or she might momentarily suppose this deontological theory to explain, contemplate, or reason with that theory in a given context. In this case, she might believe that deontology is false. Hence, her acceptance or supposition of deontology can be belief-contravening.

Misguided understanding is a different form of narrative understanding. If one has misguided understanding of why *p* obtains, then one falsely believes that *p* obtains because of some *q*. For example, I might come to believe on reliable testimonial grounds that the Sixth street bridge collapsed because it was blown up in a terrorist attack. Even though my belief is justified, it is false. The Sixth street bridge was demolished to construct a new bridge. Nevertheless, I could be attributed with some sort of understanding. I believe and understand a false account of the collapse of the Sixth street bridge. Hence, I have misguided understanding.

Narrative understanding is understood as a species of objectual understanding. Standard views of objectual understanding typically aim at getting things right. However, narrative understanding lacks a tight connection between cognitive success and true beliefs. One can narratively understand a theory even when one does not think that that theory is true or plausible (hypothetical understanding). Moreover, one can narratively understand a theory even when one falsely believes that the theory is true (misguided understanding).

Now, it might be objected here that narrative understanding is not really understanding at all. Hypothetical understanding does not require its possessor to regard the target theory or account to be true or plausible. Moreover, hypothetical understanding need not be factive, nor does it need to be based on epistemic reasons. And misguided understanding, as I have construed it here, simply gets things wrong: one falsely believes that something is the case (or that something ought to be the case). But narrative understanding can still be a form of understanding. It is weak cognitive achievement. This seems plausible given that if one has narrative understanding of some theory or account, Φ, then one can competently discern how the elements or information that constitute Φ cohere and relate to one another. One epistemically discerns the internal consistency and dependency relations that hold between the elements of Φ. And, other things being equal, one could be in a position to adequately explain Φ or employ Φ in one's reasoning. This suggests that an agent that has narrative understanding

is in a better cognitive position than one who merely endorses several propositions or pieces of information but cannot discern how those elements fit together. I contend, then, that narrative understanding is indeed a kind of cognitive achievement. It is simply an objectual form of understanding constituting a *weak* cognitive achievement.

3. The Tenability of the Three Species Model of Understanding

I advanced the MS model of understanding. On this model, understanding comes in different species: veridical understanding, conjectural understanding, and narrative understanding. Propositional understanding (e.g., understanding why) can be veridical or conjectural. Objectual understanding can be veridical, conjectural, or narrative. The main advantage of the MS model is that it provides an intuitive picture of the robust nature of understanding. Agents can possess different species of understanding depending on the cognitive endorsements they have, and the reasons they based those endorsements on. Consequently, the MS model provides the conceptual resources needed to adequately handle a wide variety of cases. In this section I expand on this line of thought. I address several alternative theories of understanding that similarly reject the JTB account. I argue that the MS model can accommodate these views. Moreover, it avoids some of the challenges these views face. I conclude that the MS model is a more tenable theory of understanding.

3.1 Unjustified Belief and Understanding

One possible alternative to the JTB model maintains that agents can possess understanding based on unjustified beliefs. Both Alison Hills (2016) and Finnur Dellsén (2016) have defended some version of this view. Hills focuses on understanding why, whereas Dellsén focuses on objectual understanding.

In this section I focus only on the account that Hills provides. Her account faces several challenges that the MS model overcomes. Hence, the MS model is a preferable theory.[208]

According to Hills, it is possible for an agent to possess understanding why *p* is the case even when that agent lacks adequate epistemic reasons that support her belief that *p* is caused by (or correctly explained by) some *q* (Hills 2016, p. 672). She motivates this claim with the following thought experiment:

> Suppose that you read in your book that Napoleon was tactically astute, and so on, and on [that] basis [you] conclude that he was a great leader. But now your history teacher, whom you regard as extremely trustworthy, tells you that Napoleon was not a great leader. Your teacher is not basing this judgment on other information or on a different interpretation of what it takes to be a great general: he simply irrationally dislikes Napoleon. You have no idea about any of this, but even so, you ignore your teacher and continue to maintain your conclusion (Hills 2016, p. 672).

In this case, the testimony from your teacher would defeat or undermine your epistemic justification for your true belief. But Hills contends that, if you have certain cognitive abilities she refers to as "cognitive control", then your true belief *that Napoleon was a great leader because he was tactically astute* can amount to understanding (Hills 2016, p. 672). As it was noted in section 1, *cognitive control* involves having certain abilities to provide explanations, follow explanations, and make correct inferences.

The main worry with Hills' account is that it allows for irrational understanding.[209] This is problematic. For it seems counterintuitive to maintain that understanding can be the product of irrationality. Daniel Wilkenfeld has also raised this worry (Wilkenfeld 2017). He points out that if understanding can be the product of unjustified true beliefs, then this would lead to "abominable

---

[208] It should be noted that Dellsén ultimately describes and endorses a view claiming that epistemically unjustified *acceptance* can generate objectual understanding (Dellsén 2016, pp. 247-251). This view escapes some of the worries facing Hills' account. Nevertheless, Dellsén's view is problematic. I address these issues in the next sub-section.

[209] Beliefs can be normatively assessed on their evidential basis: an agent's belief that *p* is *rational* when it is based on sufficient evidence or epistemic reasons in favor of *p* being the case, where these reasons are not outweighed by one's reasons for not-*p* (see, for example, Pritchard 2016, p. 90).

conjunctions" of the form: "S understands why p in virtue of believing that p *because* q, but S is irrational to believe that p *because* q" (Wilkenfeld 2017, p. 320). Indeed, it is difficult to make sense of this statement. And it seems implausible to say that one could genuinely understand why *p* is the case, where one's understanding is the result of irrational beliefs.

Nevertheless, even if one could gain understanding from an irrational true belief, this does not seem like the same kind of understanding one gets from veridical understanding based on good epistemic reasons. Veridical understanding is a strong cognitive achievement considering that its possessor is rationally sensitive and receptive to the kinds of epistemic reasons she has. Irrational understanding militates against the intuition that understanding is a cognitive achievement. And this seems implausible.

One way to avoid these challenges is to adopt the MS view and maintain that one can have conjectural understanding based on partially defeated justification if, and only if, the relevant cognitive endorsement is an attitude of acceptance. Consider again the case that Hills provides. When confronted with the testimony that Napoleon is not a great leader, you become aware of a normative defeater (albeit a misleading defeater in this case) that significantly weighs against your original evidence in support of the truth of the proposition *that Napoleon was a great leader because he was tactically astute and clever*. But if you voluntarily *accept* this true explanation, and if you grasp or epistemically discern how being astute and clever explains why Napoleon was a great leader, then you could come to possess conjectural understanding. In this case, your acceptance is not irrational. Your acceptance is based on the epistemic reasons you do have, namely, the true statements contained in your textbook. Hence, your acceptance is weakly grounded in evidential support. Consequently, your weakly justified acceptance in this case could only amount to a conjectural kind of understanding.

The MS model can handle the challenges that face Hills' account of understanding. On the MS model, understanding cannot be based on irrational beliefs. But one could have weakly justified

acceptances that generate conjectural understanding. This is a much more satisfactory account that tracks well with our intuitions. Hence, the MS model is a more tenable theory of understanding.

3.2 Unjustified Acceptance and Understanding

Finnur Dellsén similarly rejects the JTB account of understanding. Like Hills, he denies that objectual understanding requires justification (Dellsén 2016, pp. 240-247). Moreover, he rejects the idea that understanding requires belief (Dellsén 2016, pp. 247-251). He defends the view that, in some cases, unjustified acceptance, rather than belief, can generate objectual understanding (Dellsén 2016, pp. 247-251).[210] According to Dellsén, "understanding something may merely involve treating certain propositions or theories as given in the context of explaining something, as opposed to being disposed to feel that the propositions or theories are true" (Dellsén 2016, p. 248). Moreover, he thinks it is possible for one to objectually understand some theory, φ, by simply treating φ as given in a context, while simultaneously possessing a justified false belief that φ does not obtain (Dellsén 2016, p. 250).

He motivates this account of understanding with a thought experiment. He describes the case as follows:

> Carrie is a theoretical physicist in a nearby possible world (perhaps this one) in which string theory is true. Carrie has built her career around using string theory to explain various known phenomena about the natural world, and has become one of the world's leading contributors in the field because of her unmatched insight into the theory and its applications. Moreover, she has adopted the policy of treating string theory as given in her scientific endeavours – using it in explanations of various natural phenomena – and thus accepts string theory for explanatory purposes. However, like many other physicists, Carrie has significant methodological reservations about string theory in its current form, and therefore is not disposed to feel that string theory is even approximately true. In other words, Carrie does not believe that string theory is even approximately true (Dellsén 2016, p. 249).

Dellsén contends that Carrie has objectual understanding: she understands the subject matter of string theory given her unjustified, belief-contravening acceptance of this theory (Dellsén 2016, p. 249).

[210] Dellsén's also appeals to the notion of *acceptance* articulated by L. J. Cohen (Dellsén 2016, p. 248).

There are several worries with this view. First, Dellsén's view holds that unjustified acceptance can generate objectual understanding. This seems right. But it is not plausible to think that this kind of understanding is the same as veridical objectual understanding. Nor is it correct to claim that the understanding Carrie has is the kind of understanding one gets on an epistemically justified acceptance. In the case that he provides, Carrie's understanding is generated by an acceptance based primarily on pragmatic considerations. To claim that this kind of understanding is like veridical or conjectural understanding leads to problematic results. More precisely, his view is too inclusive: it allows for there to be veridic-like understanding in too many cases that, intuitively, we would not allow.

The proper diagnosis of Dellsén's case is that Carrie has *hypothetical* objectual understanding. Her understanding of string theory is based on an unjustified acceptance that conflicts with her belief that string theory is false. Even though her acceptance of the theory is true, she gets things *doxastically* wrong. She incorrectly believes that what she accepts is false. But if she has some understanding of string theory, then clearly this is not the same kind of cognitive achievement one has when they get things right if they possess veridic understanding of string theory. It is more appropriate to attribute Carrie with hypothetical understanding.

Moreover, Dellsén's view does not differentiate between understanding that is context-independent (veridic) and understanding that is context-dependent (conjectural and narrative). Veridic understanding generated by justified true belief is context-independent: one considers the object or content of understanding to be true regardless of pragmatic considerations one has in a context. But if acceptance can generate understanding, and acceptance is characteristically based on some pragmatic considerations within a context, then this kind of understanding will be contextually constrained.[211]

Consider again the case involving Carrie. In this case, Carrie accepts string theory for pragmatic reasons, namely, that Carrie is a physicist that works on string theory. Hence, her acceptance

[211] Bratman argues convincingly for this point about attitudes of acceptance (see Bratman 1999).

is context-dependent. Outside her role of being a physicist or engaging with other scientists in her field, she would not accept this theory. For she does not think the theory is plausible. Carrie would not treat string theory as given in all contexts. Hence, she ought to be credited with narrative understanding, namely, hypothetical understanding. This is a *weak* kind of cognitive achievement.

Dellsén could amend his view by adopting the MS view. The MS model better accommodates our intuitions about cases like Carrie's. Hence, it is a more tenable theory of understanding.

3.3 Low Credences and Understanding

Another view that similarly rejects the JTB model of understanding has been defended by Daniel Wilkenfeld (2017). On Wilkenfeld's view, understanding why *p* is construed as a type of mental "representation" an agent can have toward some causal or explanatory relationship between *p* and some *q* (Wilkenfeld 2017, p. 321). He argues that one can understand why *p* is the case by merely having a low justified credence (or subjective probability) that *p* is caused (or correctly explained by) some *q* which does not amount to having a full belief (Wilkenfeld 2017, pp. 321-324).

Wilkenfeld motivates this idea with the following thought experiment:

> *Challenger*: Richard is an established and skilled scientist tasked with investigating why the Challenger space shuttle exploded. His investigative instincts are excellent, and he is immediately suspicious of the O-Rings' ability to operate at cold temperatures. In fact, the reason the Challenger exploded was because of O-Ring failure. Sadly, before he goes public he is subjected to a deliberate cover-up, casting doubt on whether the O-Rings failed in this case. Arbitrarily large amounts of evidence can be presented that there really was no O-Ring failure, to the point that Richard loses confidence in his belief that the Challenger exploded because of O-Ring failure. His subjective credence […] indicate that he puts the probability of "the Challenger exploded (= *p*) because the O-Rings failed (= *q*)" at about 30 percent (hereafter the proposition that "The Challenger exploded because of O-Ring failure" will be referred to as τ). Being probabilistically coherent, he puts the probability of not-τ at 70 percent. However, Richard still has a detailed model of how O-Ring failure would cause the explosion […]. Other scientists, who are not subject to the same cover-up, consult Richard (in the guise of asking about unlikely but possible scenarios) on how the O-Ring failure destroyed the Challenger. Richard is, in fact, the world's preeminent expert on the cause of the Challenger explosion – he just fails to believe his own conclusions (Wilkenfeld 2017, pp. 321-322).

Wilkenfeld claims that since Richard still has accurate knowledge of O-Rings along with a working model of O-Ring failure causing the Challenger explosion, he still has a "justified credence" after the cover-up that the Challenger exploded due to O-Ring failure (Wilkenfeld 2017, p. 330). And even though Richard's credence toward this causal relation is too low for full belief, Wilkenfeld contends that Richard still understands why the Challenger exploded (Wilkenfeld 2017, 322).

Call the time before Richard is subjected to the misleading cover-up "t_1". And call the time after Richard is subjected to the coverup "t_2". Let us suppose that we have such credences that ought to obey the probability calculus. Let us further suppose that a credence toward *p* that falls below, say, .5 does not amount to having a full belief that *p*. If we admit these assumptions, then we can safely say that Richard justifiably believes at t_1 that the Challenger exploded due to O-Ring failure. And we can plausibly say that at time, t_2 (after Richard is subjected to the misleading cover-up) he loses his belief and only possesses a low credence that the Challenger space shuttle exploded due to O-Ring failure. Moreover, since Richard's credences at t_2 obey the probability calculus, he is rationally coherent. It seems uncontroversial that Richard has *veridic* understanding *why* the Challenger space shuttle exploded at t_1. The problem, however, is that it is implausible to claim that Richard possesses this same kind of understanding at t_2.

What is crucial to note here is that at time, t_2, Richard does not think that it is true, likely, or plausible that the Challenger exploded due to O-Ring failure. This is because Richard no longer has sufficient evidence to support the claim that the Challenger exploded due to O-Ring failure. It therefore seems mistaken to claim that Richard has genuine understanding *why* at t_2. To claim otherwise would lead to another sort of abominable conjunction (A) associated with understanding:

(A) *S* understands why *p* is the case, where *q* caused (or correctly explains) *p*, but *S* has a low degree of confidence that *q* caused (or correctly explains) *p*, and thus *S* does not think it is true, likely, or even plausible that *q* caused (or correctly explains) *p*.

(A) is false. If one possesses genuine understanding why *p* is the case, where *p* is because of some *q*, then one ought to think it is true, likely, or plausible that *q* is the cause of *p*. On Wilkenfeld's view, for example, I might have a .25 credence that the Titanic sank because it struck an iceberg. Being probabilistically coherent, I would have a .75 credence that it was not the case that the Titanic sank because it struck an iceberg. Consequently, I would think it was highly likely or probable that *it is not the case* that the Titanic sank because it struck an iceberg. Even if my low credence is based on some epistemic reasons, it is incorrect to claim that I understand why this certain state-of-affairs obtains given that I do not think it is likely that the Titanic sank because it struck an iceberg. If one understands that *q* caused *p*, then one must consider *p because q* to be a genuine possibility such that one cannot be credited with believing the negation of *p because q*.

It is more plausible to maintain that in situations resembling the *Challenger* case above, agents have a kind of *narrative* understanding. In Richard's case, he has *hypothetical* understanding at t_2. But narrative understanding is properly characterized as a kind of *objectual* understanding. Indeed, as Wilkenfeld himself notes, "objectual understanding (i.e., "understanding-of") […] does not require believing that the thing which is understood accurately represents the world" (Wilkenfeld 2017, p. 319). He claims that one could, for instance, possess understanding of a theory, *T*, by merely understanding the content of the theory while not believing that *T* (or the claims that comprise *T*) is true (Wilkenfeld 2017, p. 319). By his own lights, one might plausibly interpret the *Challenger* case as involving two types of understanding: understanding why and a hypothetical form of objectual understanding. Richard at time, t_1, understands why the Challenger exploded. But at t_2, Richard is aware of a misleading normative defeater and comes to lack the relevant understanding why. Instead, at t_2 Richard comes to only possess a type of objectual understanding: Richard *hypothetically* understands a *theory* of the Challenger explosion that he does not consider to be a likely explanation. On this

alternative interpretation of the *Challenger* case, there is a shift or equivocation in the meaning of *understanding*. But Richard's understanding at t_2 does not amount to genuine understanding why.

The MS model can adequately handle these issues. On the MS model, understanding why need not require full belief. An agent, *S*, could have a low credence that *p because q* that does not amount to full belief. And *S*'s credence could generate *conjectural* understanding why *p* if (1) it is true that *q* is the cause of *p*, (2) *S*'s credence does not amount to a full belief in the negation of *p because q*, (3) *S* voluntarily accepts *p because q*, (4) *S* properly grasps the explanatory relations between *p* and *q*, and (5) *S* basis her acceptance on epistemic reasons. However, if *S* has an epistemically justified low credence that amounts to a full belief in the negation of *p because q*, and if *S* forms a true acceptance that *p because of q* that conflicts with her belief, then *S*'s acceptance only amounts to a kind of *narrative* understanding.

The MS model presents itself as an attractive account that better accords with our intuitions concerning situations like the *Challenger* case. Of course, Wilkenfeld could incorporate this account into his own. The explanatory power of the MS model, I contend, provides us with good reason to prefer this theory of understanding over Wilkenfeld's account.

3.4 THE FACTIVITY OF UNDERSTANDING

Up until now it has been taken for granted that understanding why is factive. This seems correct. It does not seem plausible to claim that one can understand why *p* is true when *p* is false. But is objectual understanding essentially factive? Kvanvig, among others, insists that the kind of objectual understanding that is important for epistemology is essentially factive (Kvanvig 2003, p. 191). To objectually understand a subject matter, he claims, is to possess true beliefs about that subject matter (Kvanvig 2003, p. 191). But this intuition has been challenged by Catherine Elgin (Elgin 2009; 2017). Elgin argues that much of our objectual understanding in areas like science is not straightforwardly factive (Elgin 2009; 2017). Elgin affirms that "understanding somehow answers to facts" (Elgin 2017,

p. 37). But she claims that, while our understanding of some of our best theories are appropriately connected to the facts, these accounts often contain falsehoods (Elgin 2017, p. 15). Elgin's view of non-factive objectual understanding is sophisticated, and I will not spend much time discussing the finer details of it here. Rather, in this section I aim to do two things. First, I will briefly sketch the fundamental aspects of Elgin's account of objectual understanding. Second, I argue that the MS account can adequately accommodate Elgin's position while also preserving certain intuitions concerning the factivity of objectual understanding. I contend that Elgin's account would be strengthened if it incorporated the core idea of the MS model: understanding comes in multiple species.

According to Elgin, if an agent understands a given theory or account, then she "must also grasp how the various truths relate to each other and to other elements of the account" (Elgin 2017, p. 46). And she insists that if one understands some account, then one "should also be both willing and able (and perhaps be aware that she is willing and able) to use that information – to profess it, to reason with it, to apply it, perhaps to use it as a source of working hypotheses about related matters when her ends are cognitive" (Elgin 2017, p. 46).[212]

On this view, many cases of understanding will be factive. However, Elgin points out that much of our scientific understanding expands, such that the initial understanding of a scientific subject matter will contain some false claims that, over time, are either replaced with truths or modified in a way that they become "closer to the truth" (Elgin 2017, p. 59). Elgin contends that even though many of our scientific theories contain falsehoods, our endorsements of them constitute genuine understanding, and that this understanding broadens when we expand and improve these theories (Elgin 2017, p. 61).

[212] For Elgin, understanding involves a type of know-how (see Elgin 2017, Ch. 3, pp. 46-50).

In addition to her claims about scientific progress, Elgin notes that many of our scientific theories contain idealizations that are strictly speaking false representations of how things are (Elgin 2017, p. 61). She notes, for example, that the ideal gas law is a fundamental principle employed in many scientific accounts (Elgin 2017, p. 61). But the ideal gas law does not accurately represent real gases. She points out that the ideal gas law stipulates that gases are "dimensionless, spherical molecules that are not subject to friction and exhibit no intermolecular attraction" (Elgin 2017, p. 61). But even though the ideal gas law and other idealized models employed in science are strictly speaking false, Elgin maintains that they are epistemically useful, and that it is still appropriate to credit scientists with genuine understanding of the accounts that contain these idealizations (Elgin 2017, p 62). Consequently, objectual understanding need not be factive.

What is important to note here is that Elgin contends that we are not warranted in believing these idealizations and falsehoods. She asserts the following:

> I do not then claim that it is epistemically acceptable to *believe* what is false or that it is linguistically acceptable to *assert* what is false. Rather, I suggest that epistemic acceptance is not restricted to, and does not always involve, belief (Elgin 2017, p. 18).

When one endorses and employs an idealization, one can *accept* this idealization without simultaneously believing that the idealization is a veridical or accurate representation of how things stand in the world. For Elgin, then, acceptance plays a significant role in our objectual understanding.[213] When idealizations factor into a body of information or account that we endorse, Elgin contends that we can justifiably accept these falsehoods when they are "true enough" – that is, when their "divergence from truth is negligible" (Elgin 2017, p. 28). Hence, objectual understanding that is non-factive will typically be grounded in attitudes of acceptance that do not coincide with belief.

[213] Elgin appeals to a modified version of Cohen's account of *acceptance*. She claims that "to accept that *p* is to be willing and able to take *p* as an assertoric premise, epistemic norm or rule of inference in one's reasoning or as a basis for action when one's ends are cognitive" (Elgin 2017, p. 19).

It seems plausible that objectual understanding in the scientific domains can be non-factive. And one could rationally accept these non-factive accounts. But on the MS model, understanding generated by acceptance differs from the kind of understanding grounded in justified true belief. Elgin, however, does not clearly draw this distinction. On her account, it seems that non-factive understanding based on approximately true acceptances can deliver the same kind of understanding generated by justified true beliefs. But this does not seem quite right. It is incorrect to conflate the understanding generated by true belief with the kind of non-factive understanding generated by acceptance. The former kind of understanding is veridical. It is a stable and true representation of how things obtain that is context-independent and based on strong epistemic support. Hence, veridical understanding is a strong cognitive achievement. But objectual understanding generated by approximately true acceptance within a context is a weaker cognitive achievement that is essentially less stable. This is a form of conjectural objectual understanding: it is typically context-dependent, and it lacks the strong rational grounding enjoyed by veridical objectual understanding. Understanding is properly characterized as coming in different species. And different cognitive endorsements deliver different kinds of understanding.

If, for example, a group of scientists accept and employ an idealization into a coherent theory or account that they endorse and properly comprehend, and if this overall account or theory is approximately true, then these individuals could have conjectural objectual understanding. Such understanding would differ from veridical objectual understanding. Considering that these agents recognize that these idealizations only approximate the truth, the kind of objectual understanding they possess would be contextually bound to those cognitive situations in which they have pragmatic reason to accept and employ these theories and accounts when performing inquiry. This seems to be a more plausible interpretation of the kind of understanding that factors into scientific understanding that is non-factive.

Adopting the MS model would strengthen Elgin's position. It provides the conceptual resources needed to accurately describe our scientific understanding. A further benefit of the MS account is that it offers a middle ground between non-factivist and factivist conceptions of understanding. On the MS model, veridical understanding is essentially factive. On the other hand, conjectural objectual understanding can be non-factive or approximately true. This accords well with our intuitions. The MS model therefore provides a cogent theory of understanding.

4. Conclusion: Understanding, Knowledge, Commitments, and Achievements

In this chapter I articulated the MS model of understanding. On this account, understanding comes in different species: veridical, conjectural, and narrative. The MS model, I argued, provides an intuitive and accurate picture of understanding. Moreover, this model has the conceptual machinery needed to successfully handle other, rival accounts of understanding. I therefore conclude that the MS model is a tenable and explanatorily powerful theory that advances our understanding of understanding.

The MS model also provides a new framework for characterizing cognitive achievements. Cognitive achievements can have different degrees of strength. Veridical understanding engenders a strong cognitive achievement. This is due to the fact that veridical understanding involves a justified true belief. Hence, this understanding gets things right. Moreover, veridical understanding is based on an agent's ability to exercise certain cognitive abilities. First, she must be receptive and sensitive to her available reasons. Second, she must competently exercise epistemic discernment. Conjectural understanding, however, naturally presents itself as a moderate form of cognitive achievement. For conjectural understanding need not always involve a correct acceptance. In some cases, one can conjecturally understand things that are merely approximately true. Hence, this form of understanding lacks a tight connection to the truth of what is understood. Nevertheless, conjectural understanding requires one to be rationally sensitive to one's available reasons, and to be able to competently exercise

epistemic discernment. Finally, narrative understanding (in cases of objectual understanding) is a weak form of cognitive achievement. Unlike veridical and conjectural understanding, one need not get things right to have narrative understanding. Nor does one need to base one's narrative understanding on some epistemic aim or goal to get at the truth. Hence, narrative understanding is not always linked to a kind of success or accomplishment. However, narrative understanding still requires an ability to adequately exercise epistemic discernment. I have argued that this leads to a weak kind of cognitive achievement. Consequently, we have a taxonomy or hierarchy of cognitive achievements associated with understanding.

Moreover, the MS model further explicates the relationship between knowledge and understanding. One common view maintains that understanding is simply a kind of knowledge (Kitcher 2002; Lipton 2004; Grimm 2006, 2014). Peter Lipton claims that "[u]nderstanding is not some sort of super-knowledge, but simply more knowledge: knowledge of causes" (Lipton 2004, p. 30). However, some theorists, including advocates of the JTB model of understanding, reject this claim. These theorists argue that understanding is a distinct kind of cognitive achievement. They typically argue that, unlike knowledge, understanding is not undermined by luck (Kvanvig 2003; Pritchard 2009, 2014; Hills 2016).[214] If this is right, then understanding and knowledge can come apart.

Adopting the MS model provides another way to sever the connection between knowledge and understanding. For if knowledge entails belief, and if understanding can be generated by

[214] Some theorists contend that understanding is only compatible with some kinds of luck. Pritchard, for example, identifies two kinds of luck that undermine propositional knowledge (i.e., knowledge-that): "Gettier-style" luck and "environmental epistemic" luck (Pritchard 2014, p. 317). As Pritchard explains elsewhere, in cases of Gettier-style luck, an agent forms a justified, true belief, and yet "something intervenes 'betwixt belief and fact'" (Pritchard 2009, p. 21). Environmental epistemic luck, however, occurs when an agent forms a justified, true belief in a reliable and appropriate way, but only does so luckily given that they were in a lucky epistemic situation (Pritchard 2009, pp. 35-36). For Pritchard, one can fail to know that *p* is caused by *q* due to environmental luck, and yet, in this case, one could still have understanding why *p* (Pritchard 2009, p. 36; 2014 p. 317).

acceptance and supposition, then it follows that understanding *simpliciter* is not a kind of knowledge. Understanding is a robust epistemic state that manifests different species of cognitive achievements.

Finally, it is important to note that understanding is deeply connected to rationally held cognitive endorsements. Understanding is based on rational cognitive endorsements. A necessary condition for possessing a rational cognitive endorsement is the proper satisfaction of its corresponding cognitive commitments (see chapter 2). Hence, our cognitive commitments play an important role in facilitating cognitive achievements of understanding. It follows, then, that understanding is also an inherently normative matter.

CONCLUSION: NEW FOUNDATIONS OF RATIONALITY AND UNDERSTANDING

Adam C. Sanders

This dissertation has explored several issues concerning the nature and normativity of our cognitive endorsements. Moreover, my research has explicated the nature of understanding and rationality. Each of the contained essays have argued for several theses. Taken together, these individual theses establish a general, overarching thesis, (T):

> **(T)** Beliefs, acceptances, suppositions, presuppositions, and hypothesizing are distinct cognitive endorsements generating unique normative commitments that govern rational thinking and facilitate understanding.

Beliefs, acceptances, suppositions, presuppositions, and hypothesizing are unique mental states and processes. Each of these states differ in how they are conceptually tied to the truth, how they are rationally supported by one's reasons, and how they are rationally brought about in thought.

It was further shown that conscious cognitive endorsements generate unique cognitive commitments. These commitments are constitutive elements of the attitudes themselves. Each cognitive commitment normatively governs how one ought to mentally treat the propositional content of the associated state.

Moreover, our beliefs partly engender semantic commitments that normatively constrain our diachronic belief-forming processes. If a proposition, *p*, follows from an agent's current beliefs via some inference rule, then that agent can be semantically committed to believe *p* insofar as they possess the required cognitive resources to endorse *p*. One's cognitive resources to endorse *p* is constituted by one's current belief-set, understanding, conceptual repertoire, and cognitive abilities.

Considering that our cognitive commitments and semantic commitments constrain our thinking and reasoning, it is therefore no surprise that these commitments will contribute to epistemic rationality. Indeed, I have argued for the Robust View of epistemic rationality. According to this view,

the rationality governing our doxastic states is constituted by both external, mind-independent rules and internal, mind-dependent commitments.

Finally, I advanced the "Multiple Species" view of understanding. On this view, understanding involves a cognitive endorsement and an ability of epistemic discernment. Rationally held attitudes of belief, acceptance, and supposition can each generate unique species of understanding. This view provides a more plausible picture of our cognitive achievements of understanding.

Establishing (T) is theoretically important for several reasons. First, it helps provide the conceptual resources needed to clearly demarcate our individual cognitive endorsements that factor into our thinking and reasoning. Second, it helps explicate the normative structure of thought. Our conscious cognitive endorsements generate unique commitments. Consequently, these states are inherently normative. This further illuminates the proper way to conceive of epistemic rationality. Epistemic rationality is not merely a system of external rules that agents ought to comply with. Rather, rationality is best construed as a system of rules and commitments. Finally, we are left with a more intuitive picture of cognitive achievements. It is uncontroversial that belief is an important state for generating knowledge and strong cognitive achievements of veridical understanding. However, our cognitive endorsements of acceptance and supposition can also facilitate understanding. Hence, these states can also facilitate cognitive achievements and promote rational agency.

It is my intention to build upon this important work. Doing so, I contend, can lead to important contributions to epistemology and the philosophy of mind. Moreover, I believe that this project provides foundations for future research concerning epistemic rationality and understanding. Before concluding the dissertation, I want to briefly discuss some of this potential research.

A portion of the work contained in this dissertation has centered on doxastic rationality. However, it is an interesting and important task to explicate the rationality of our other cognitive endorsements. I am particularly interested in constructing a complete theory of rationality governing

acceptance. My contention is that acceptance will involve its own theory of rationality. For these attitudes engender unique normative cognitive commitments. Consequently, agents ought to satisfy these commitments. To do otherwise would result in erroneously held states of acceptance. In addition to the fulfillment of these commitments, it is conceivable that attitudes of acceptance might be governed by further external rules and norms. However, the rational rules governing attitudes of acceptance will be significantly different from the requirements governing belief.

What might a complete theory of the rationality of acceptance look like? First, let us recall that accepting a proposition, *p*, is different from believing *p*. Accepting *p* need not involve thinking that *p* is true. Moreover, one can accept *p* for pragmatic reasons. Consequently, it is possible to pragmatically accept contradictions. If there is a set of further rules governing acceptance, then this set would not include a rule of Non-contradiction. However, accepting *p* is rationally held for reasons an agent has. It was previously argued that these reasons can be epistemic or pragmatic. As such, it is plausible to say that there could be further rational norms and rules governing the rational basis of attitudes of acceptance. These rules, however, should be broad enough to countenance both epistemic and pragmatic considerations. If this is correct, then the system of rationality governing acceptance would fundamentally differ from the robust view of doxastic rationality defended in chapter 5.

Another important topic worth pursuing concerns issues in collective epistemology. There has been significant work done on the epistemology of groups. One important question is the following: Can a group of agents be credited with epistemic states of belief? According to Margaret Gilbert, we typically ascribe beliefs to groups (see Gilbert 2004). She explicates these ascriptions by appealing to "joint commitments" (Gilbert 2004). She claims that a group, *G*, believes *p* if and only if the members that constitute *G* "are jointly committed to believe as a body that *p*" (Gilbert 2004, p. 100). Roughly, the members of *G* commit jointly to believing *p* as a body when they convey their willingness to believe *p* with each other, and this expression is common knowledge among the members (Gilbert 2004, p.

100).[215] Other theorists, however, insist that the epistemic attitudes that groups have are not beliefs, but some other type of state. Brad Wray, for instance, argues that the epistemic state that a group can be credited with is properly characterized as *acceptance* (Wray 2001, p. 325). He insists that belief involves an involuntary feeling that something is true, whereas acceptance is typically brought about in a voluntary way for pragmatic ends (acceptance can, in some cases, conflict with belief) (Wray 2001, p. 325). According to Wray, groups seem to voluntarily and collectively endorse things to satisfy some goal, and thus the resulting attitude we commonly attribute to groups is best characterized as acceptance (Wray 2001, p. 325). In response, Gilbert and Pilchman defend the view that groups can be attributed with beliefs (Gilbert and Pilchman 2014, pp. 189 – 212). They also note that the notions of *belief* and *acceptance* that are commonly deployed in the above debate are developed out of individual epistemology (Gilbert and Pilchman 2014, pp. 207 – 209). But *collective belief*, they contend, could differ from our notions of *individual* states of *belief* and *acceptance* (Gilbert and Pilchman 2014, pp. 207 – 209).

The debates and views noted above are sophisticated. I have only given a rough and incomplete presentation of the aforementioned positions concerning group beliefs. For my purposes, I will not discuss the finer details. Instead, I want to briefly discuss another important issue in collective epistemology that directly relates to my current project: collective understanding. This issue deserves more attention. I have previously argued that individual understanding can be generated by a cognitive endorsement and an ability of epistemic discernment. *Epistemic discernment* is an ability to recognize how propositions and pieces of information fit together or relate to one another in some epistemically significant way. I am also sympathetic to the idea that groups of agents can manifest understanding. However, this type of understanding would differ from the understanding individual agents can possess. For one thing, groups do not seem to possess the same kinds of mental states that individual

[215] This is a rough picture of Gilbert's view. Her view differs from what she calls "rejectionism" (I will discuss Wray's position, which is a version of rejectionism). See Gilbert (2004) and Gilbert and Pilchman (2016) for further discussion of joint commitment and collective beliefs.

agents possess. Moreover, it is unlikely that groups of people would possess the same kind of epistemic discernment ability that individual agents can manifest.

Nevertheless, it is plausible that groups can possess understanding. One possible way to characterize collective understanding is as follows. We might say that a group, *G*, understands some phenomenon, *p*, when (1) *G* appropriately acquires some state of belief or acceptance concerning *p*, and (2) *G* can competently manifest the ability of collective epistemic discernment. *Collective epistemic discernment* is plausibly understood to be an ability that members of a group have to collectively employ and use information in collective reasoning and inquiry.[216] It seems that if a group, *G*, can competently satisfy conditions (1) and (2), then we can correctly attribute *G* with some form of understanding.

This is a rough suggestion of what an account of group understanding might look like. There is no doubt that more work is needed to articulate a plausible and complete theory of collective understanding. However, my current research provides the necessary conceptual resources for pursuing this important project. It is my intention to continue this important line of work.

Once we conceive of cognitive endorsements as inherently normative states, then we begin to accurately understand the complex normative structure of thought and reasoning. Moreover, we are able to illuminate several important aspects of rational thinking. This has interesting consequences for our views of understanding, our theories of rationality, and our accounts of the nature and normativity of the mind. It is my hope that the research I have produced in this dissertation will continue to promote fruitful research in the fields of epistemology, philosophy of mind, and theories of action.

[216] This collective ability is inspired by Catherine Elgin's account of individual understanding (see Elgin 2017). For Elgin, an individual agent's understanding of Φ includes an ability and readiness "to use that information [i.e., the information constituting Φ] – to profess it, to reason with it, to apply it, perhaps to use it as a source of working hypotheses about related matters when her ends are cognitive" (Elgin 2017, p. 46). On my proposed account, collective understanding involves a collective ability to use the content of what is understood in collective reasoning and inquiry to satisfy a collective goal.

REFERENCES

Adler, Jonathan. 1996. "An Overlooked Argument for Epistemic Conservatism." *Analysis* 56, 2: 80 – 84.

Alston, William. 2007. "Audi on Nondoxastic Faith." In *Rationality and the Good*. Edited by Mark Timmons, John Greco, and Alfred R. Mele, pp. 123 – 139. Oxford: Oxford University Press.

Audi, Robert. 1991. "Intention, Cognitive Commitment, and Planning." *Synthese* 86, 3: 361 – 378.

Audi, Robert. 1994. "Dispositional Beliefs and Dispositions to Believe." *Nous* 28, 1: 419 – 434.

Balcerak Jackson, Magdalena and Brendan Balcerak Jackson. 2013. "Reasoning as a Source of Justification." *Philosophical Studies* 164: 113-126.

Bilgrami, Akeel. 2006. *Self-knowledge and Resentment.* Cambridge: Harvard University Press.

Bilgrami, Akeel. 2017. "Norm and Failure in Mind and Meaning." *Argumenta* 3, 1: 69 – 85.

Boghossian, Paul. 2014. "What is Inference?" *Philosophical Studies* 169: 1 – 14.

Bratman, Michael. 1992. "Practical Reasoning and Acceptance in a Context." *Mind*. 101, 401: 1 – 15.

Bratman, Michael. 1999. *Faces of Intention: Selected Essays on Intention and Agency*. Cambridge: Cambridge University Press.

Bratman, Michael. 2009. "Intention Rationality." *Philosophical Explorations* 12, 3: 227 – 241.

Broome, John. 1999. "Normative Requirements." *Ratio* XII, 4: 398 – 419.

Broome, John. 2000. "Normative Requirements." In *Normativity*. Edited by Jonathan Dancy, pp. 78 – 99. Malden: Blackwell Publishers.

Broome, John. 2002. "Practical Reasoning." In *Reason and Nature: Essays in the Theory of Rationality*. Edited by Jose Bermudez and Alan Millar, pp. 85-112. Oxford: Oxford University Press.

Broome, John. 2005. "Does Rationality Give Us Reasons?" *Philosophical Issues* 15: 321 – 337.

Broome, John. 2007. "Wide or Narrow Scope?" *Mind* 116, 462: 359 – 370.

Broome, John. 2008. "Is Rationality Normative?" *Disputatio* II, 24: 153 – 170.

Broome, John. 2013. *Rationality Through Reasoning*. Malden: Wiley Blackwell.

Broome, John. 2014. "Normativity in Reasoning." *Pacific Philosophical Quarterly* 95: 622 – 633.

Brunero, John. 2010. "The Scope of Rational Requirements." *The Philosophical Quarterly* 60, 238: 28 – 49.

Calhoun, Cheshire. 2009. "What Good is Commitment?" *Ethics* 119, 4: 613 – 641.

Cassam, Quassim. 2010. "Judging, Believing and Thinking." *Philosophical Issues* 20: 80 – 95.

Christensen, David. 1994. "Conservatism in Epistemology." *Nous* 28, 1: 69 – 89.

Christensen, David. 2004. *Putting Logic in its Place: Formal Constraints on Rational Belief.* Oxford: Oxford University Press.

Cohen, L. Jonathan. 1989. "Belief and Acceptance." *Mind* 98, 391: 367 – 389.

Cohen, L. Jonathan. 1992. *An Essay on Belief and Acceptance.* Oxford: Oxford University Press.

Coliva, Annalisa. 2015. *Extended Rationality: A Hinge Epistemology*. London: Palgrave Macmillan.

Coliva, Annalisa. 2016. *The Varieties of Self-knowledge.* London: Palgrave Macmillan.

Coliva, Annalisa. 2019. "Disagreeing with Myself: Doxastic Commitments and Intrapersonal Disagreement." *American Philosophical Quarterly* 56, 1: 1 – 13.

Dancy, Jonathan. 2000. *Practical Reality*. Oxford: Oxford University Press.

Davidson, Donald. 1985. "Incoherence and Irrationality." *Dialectica* 39, 4: 345 – 354.

Davis, Wayne. 1988. "Knowledge, Acceptance, and Belief." *The Southern Journal of Philosophy* XXVI*,* 2: 169 – 178.

Dellsén, Finnur. 2016. "Understanding Without Justification or Belief." *Ratio* XXX, 3: 239 – 254.

Dretske, Fred. 2000. *Perception, Knowledge and Belief: Selected Essays.* Cambridge: Cambridge University Press.

Elgin, Catherine. 2009. "Is Understanding Factive?" In *Epistemic Value*. Edited by Alan Millar, Adrian Haddock, and Duncan Pritchard, 322-330. Oxford: Oxford University Press.

Elgin, Catherine. 2017. *True Enough*. Cambridge: MIT Press.

Feldman, Richard. 2000. "The Ethics of Belief." *Philosophy and Phenomenological Research* 60, 3: 667 – 695.

Fisher, Alec. 1989. "Suppositions in Argumentation." *Argumentation* 3: 401 – 413.

Foley, Richard. 1993. *Working Without a Net: A Study of Egocentric Epistemology*. New York: Oxford University Press.

Frankish, Keith. 1998. "Natural Language and Virtual Belief." In *Language and Thought*. Edited by Peter Carruthers and Jill Boucher, pp. 248 – 269. Cambridge: Cambridge University Press.

Frege, Gottlob. 1964. *The Basic Laws of Arithmetic*. Translated and edited by Montgomery Furth. Berkeley: University of California Press.

Frege, Gottlob. 1997. "Negation." In *The Frege Reader*. Edited by Michael Beaney, pp. 346 – 361. Malden: Blackwell Publishing.

Frege, Gottlob. 1997. "On *Sinn* and *Bedeutung*." In *The Frege Reader*. Edited by Michael Beaney, pp. 151 – 171. Malden: Blackwell Publishing.

Frege, Gottlob. 1997. "Thought". In *The Frege Reader*. Edited by Michael Beaney, pp. 325 – 345. Malden: Blackwell Publishing.

Friedman, Jane. 2013. "Suspended Judgment." *Philosophical Studies* 162: 165 – 181.

Gilbert, Margaret. 2004. "Collective Epistemology." *Episteme* 1, 2: 95 – 107.

Gilbert, Margaret. 2013. "Commitment." In *The International Encyclopedia of Ethics*. Edited by Hugh LaFollette, pp. 899 – 905. Hoboken: Wiley-Blackwell.

Gilbert, Margaret. 2014. *Joint Commitment: How We Make the Social World.* Oxford: Oxford University Press.

Gilbert, Margaret. 2018. *Rights and Demands: A Foundational Inquiry*. Oxford: Oxford University Press.

Gilbert, Margaret and Daniel Pilchman. 2016. "Belief, Acceptance, and What Happens in Groups: Some Methodological Considerations." In *Essays in Collective Epistemology*. Edited by Jennifer Lackey, pp. 189 – 212. Oxford: Oxford University Press.

Grimm, Stephen. 2006. "Is Understanding a Species of Knowledge?" *British Journal for the Philosophy of Science* 57: 515 – 535.

Grimm, Stephen. 2014. "Understanding as Knowledge of Causes." In *Virtue Epistemology Naturalized: Bridges Between Virtue Epistemology and Philosophy of Science*. Edited by Abrol Fairweather, 329 – 345. Cham: Springer International Publishing.

Harman, Gilbert. 1983. "Rational Action and the Extent of Intentions." *Social Theory and Practice* 123 – 141.

Harman, Gilbert. 1986. *Change in View*. Cambridge: MIT Press.

Hills, Alison. 2009. "Moral Testimony and Moral Epistemology." *Ethics* 120, 4: 94 – 127.

Hills, Alison. 2016. "Understanding Why." *Nous* 50, 4: 661 – 688.

Hume, David. 1975. *Enquiries Concerning Human Understanding and Concerning the Principles of Morals*. Edited by L. A. Selby-Bigge. Oxford: Oxford University Press.

Hume, David. 1978. *A Treatise of Human Nature*. Edited by L. A. Selby-Bigge. Oxford: Oxford University Press.

Jackson, Frank. 2000. "Non-cognitivism, Normativity, Belief." In *Normativity*. Edited by Jonathan Dancy, pp. 100 – 115. Oxford: Blackwell Publishers.

Kitcher, Philip. 2002. "Scientific Knowledge." In *The Oxford Handbook of Epistemology*. Edited by Paul Moser, 385 – 408. New York: Oxford University Press.

Kolodny, Niko. 2005. "Why Be Rational?" *Mind* 114: 509 – 563.

Kvanvig, Jonathan. 2003. *The Value of Knowledge and the Pursuit of Understanding*. New York: Cambridge University Press.

Laudan, Larry. 1977. *Progress and Its Problems*. Berkeley: University of California Press.

Lehrer, Keith. 1974. *Knowledge*. Oxford: Clarendon Press.

Lehrer, Keith. 1990. *Theory of Knowledge*. Boulder: Westview Press.

Levi, Isaac. 1996. *For the Sake of Argument: Ramsey Test Conditionals, Inductive Inference, and Nonmonotonic Reasoning*. New York: Cambridge University Press.

Levi, Isaac. 1997. *The Covenant of Reason: Rationality and the Commitments of Thought*. Cambridge: Cambridge University Press.

Liberman, Alida and Mark Schroeder. 2016. "Commitment: Worth the Weight." In *Weighing Reasons*. Edited by Errol Lord and Barry Maguire, pp. 104 – 120. Oxford: Oxford University Press.

Lipton, Peter. 2004. *Inference to the Best Explanation*. New York: Routledge.

Lord, Errol and Barry Maguire. 2016. "An Opinionated Guide to the Weight of Reasons." In *Weighing Reasons*. Edited by Errol Lord and Barry Maguire, pp. 3 – 24. Oxford: Oxford University Press.

Lynch, Michael. 2009. *Truth as One and Many*. Oxford: Oxford University Press.

Lynch, Michael. 2013. "Epistemic Commitments, Epistemic Agency, and Practical Reasons." *Philosophical Issues* 23: 343 – 371.

McCann, Hugh. 1986. "Intrinsic Intentionality." In *Action, Decision, and Intention: Studies in the Foundations of Action Theory as an Approach to Understanding Rationality and Decision*. Edited by Robert Audi, pp. 247 – 274. Dordrecht: Reidel Publishing Company.

Millar, Alan. 2004. *Understanding People: Normativity and Rationalizing Explanation*. Oxford: Oxford University Press.

Murphy, Peter. 2013. "Another Blow to Knowledge From Knowledge." *Logos & Episteme* IV, 3: 311 – 317.

Peacocke, Christopher. 1998. "Conscious Attitudes, Attention, and Self-Knowledge." In *Knowing Our Own Minds*. Edited by Barry C. Smith, Cynthia Macdonald, and Crispin Wright, pp. 63 – 98. Oxford: Oxford University Press.

Pietz, Andreas and Umberto Rivieccio. 2013. "Nothing but the Truth." *Journal of Philosophical Logic* 42, 1: 125 – 135.

Pritchard, Duncan. 2014. "Knowledge and Understanding." In *Virtue Epistemology Naturalized: Bridges Between Virtue Epistemology and Philosophy of Science*. Edited by Abrol Fairweather, 315 – 327. Cham: Springer International Publishing.

Pritchard, Duncan. 2016. *Epistemic Angst: Radical Skepticism and the Groundlessness of our Believing*. New Jersey: Princeton University Press.

Quine, W. V. 1987. *Quiddities*. Cambridge: Harvard University Press.

Quine, W. V. 2004. "Two Dogmas of Empiricism." In *Quintessence*. Edited by Roger F. Gibson, Jr., pp. 31 – 53. Cambridge: The Belknap Press of Harvard University Press.

Quine, W. V. 2008. "Mind and Verbal Dispositions." In *Quintessence*. Edited by Roger F. Gibson, Jr., pp. 313 – 326. Cambridge: Harvard University Press.

Quine, W. V. and J. S. Ullian. 1970. *The Web of Belief*. New York: Random House.

Raz, Joseph. 1975. *Practical Reason and Norms*. London: Hutchinson & Co.

Rescher, Nicholas. 1964. *Hypothetical Reasoning*. Amsterdam: North-Holland Publishing.

Scanlon, T. M. 2000. *What We Owe To Each Other*. Cambridge: Harvard University Press.

Schroeder, Mark. 2004. "The Scope of Instrumental Reason." *Philosophical Perspectives* 18: 337 – 364.

Searle, John. 1970. *Speech Acts*. London: Cambridge University Press.

Searle, John. 2001. *Rationality in Action*. Cambridge: MIT Press.

Shpall, Sam. 2013. "Wide and Narrow Scope." *Philosophical Studies* 163: 717 – 736.

Shpall, Sam. 2014. "Moral and Rational Commitment." *Philosophy and Phenomenological Research* LXXXVIII, 1: 146 – 172.

Southwood, Nicholas. 2008. "Vindicating the Normativity of Rationality." *Ethics* 119: 9 – 30.

Stalnaker, Robert. 1987. *Inquiry*. Cambridge: The MIT Press.

Stevenson, Leslie. 2002. "Six Levels of Mentality." *Philosophical Explorations* 5, 2: 105 – 124.

Swinburne, Richard. 2001. *Epistemic Justification*. Oxford: Oxford University Press.

Ullmann-Margalit, Edna and Avishai Margalit. 1992. "Holding True and Holding as True." *Synthese* 92: 167 – 187.

van Fraassen, Bas. 1980. *The Scientific Image*. Oxford: Oxford University Press.

Velleman, David. 2000. *The Possibility of Practical Reason*. Oxford: Oxford University Press.

Way, Jonathan. 2010. "The Normativity of Rationality." *Philosophy Compass* 5, 2: 1057 – 1068.

Way, Jonathan. 2011. "The Symmetry of Rational Requirements." *Philosophical Studies* 155, 2: 227 – 239.

Wray, K. Brad. 2001. "Collective Belief and Acceptance." *Synthese* 129, 3: 319 – 333.

Wedgwood, Ralph. 2002. "The Aim of Belief." *Philosophical Perspectives* 16: 267 – 297.

Whiting, Daniel. 2013. "Truth: the Aim and Norm of Belief." *Teorema* 32, 3: 121 – 135.

Wilkenfeld, Daniel. 2017. "Understanding Without Believing." In *Explaining Understanding: New Perspectives from Epistemology and Philosophy of Science*. Edited by Stephen Grimm, Christopher Baumberger, and Sabine Ammon, 318 – 333. New York: Routledge.

Williams, Bernard. 1973. *Problems of the Self*. Cambridge: Cambridge University Press.

Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press. pp. 313 – 326. Cambridge: Harvard University Press.

Wittgenstein, Ludwig. 1969. *On Certainty*. Malden: Blackwell Publishing.

APPENDIX A: A DIAGRAM OF ACCEPTANCE

The following figure characterizes some of the main parts and properties of the attitude of *acceptance*.

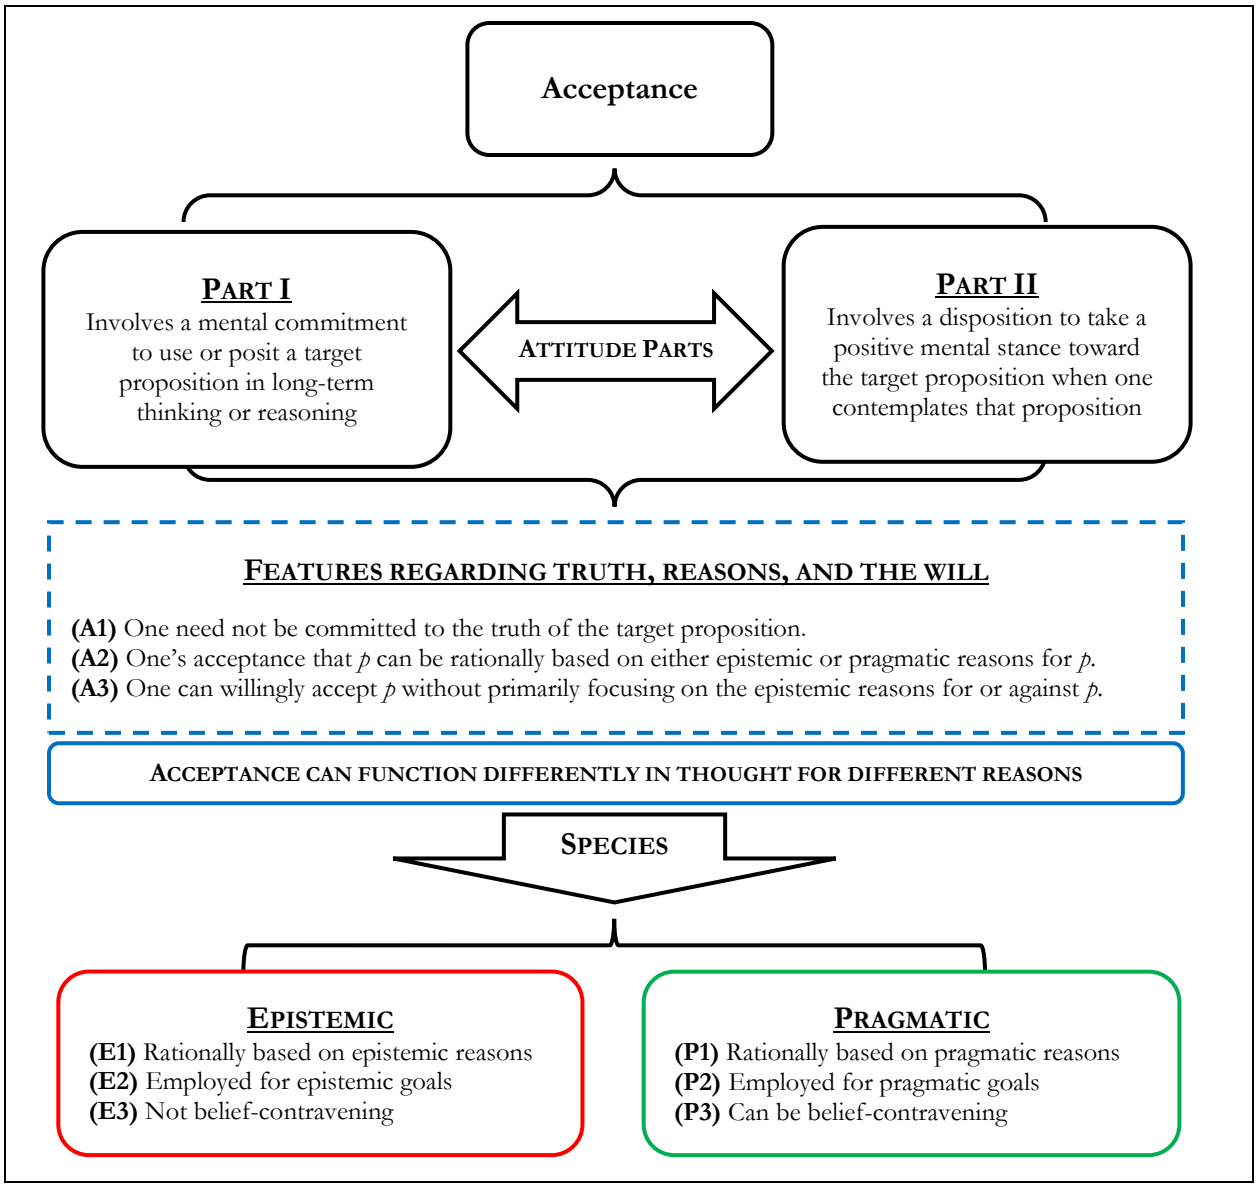

Figure 3: A Diagram of Acceptance

NOTES	
Attitude Parts	Part I of *acceptance* conflicts with Cohen's (1992) view. However, Part II coheres with Alston's (2007) account.
Species	*Epistemic acceptance* (i) coheres with Alston's (2007) view, and (ii) is typically employed in inquiry and science. *Pragmatic acceptance* (i) coheres with elements of Bratman (1999) and Cohen's (1992) account, and (ii) is sometimes synonymous with taking something for granted for pragmatic ends.

Table 10: Notes on Acceptance

APPENDIX B: A DIAGRAM OF SUPPOSITION

The figure below characterizes some of the main parts and properties of the attitude of *supposition*.

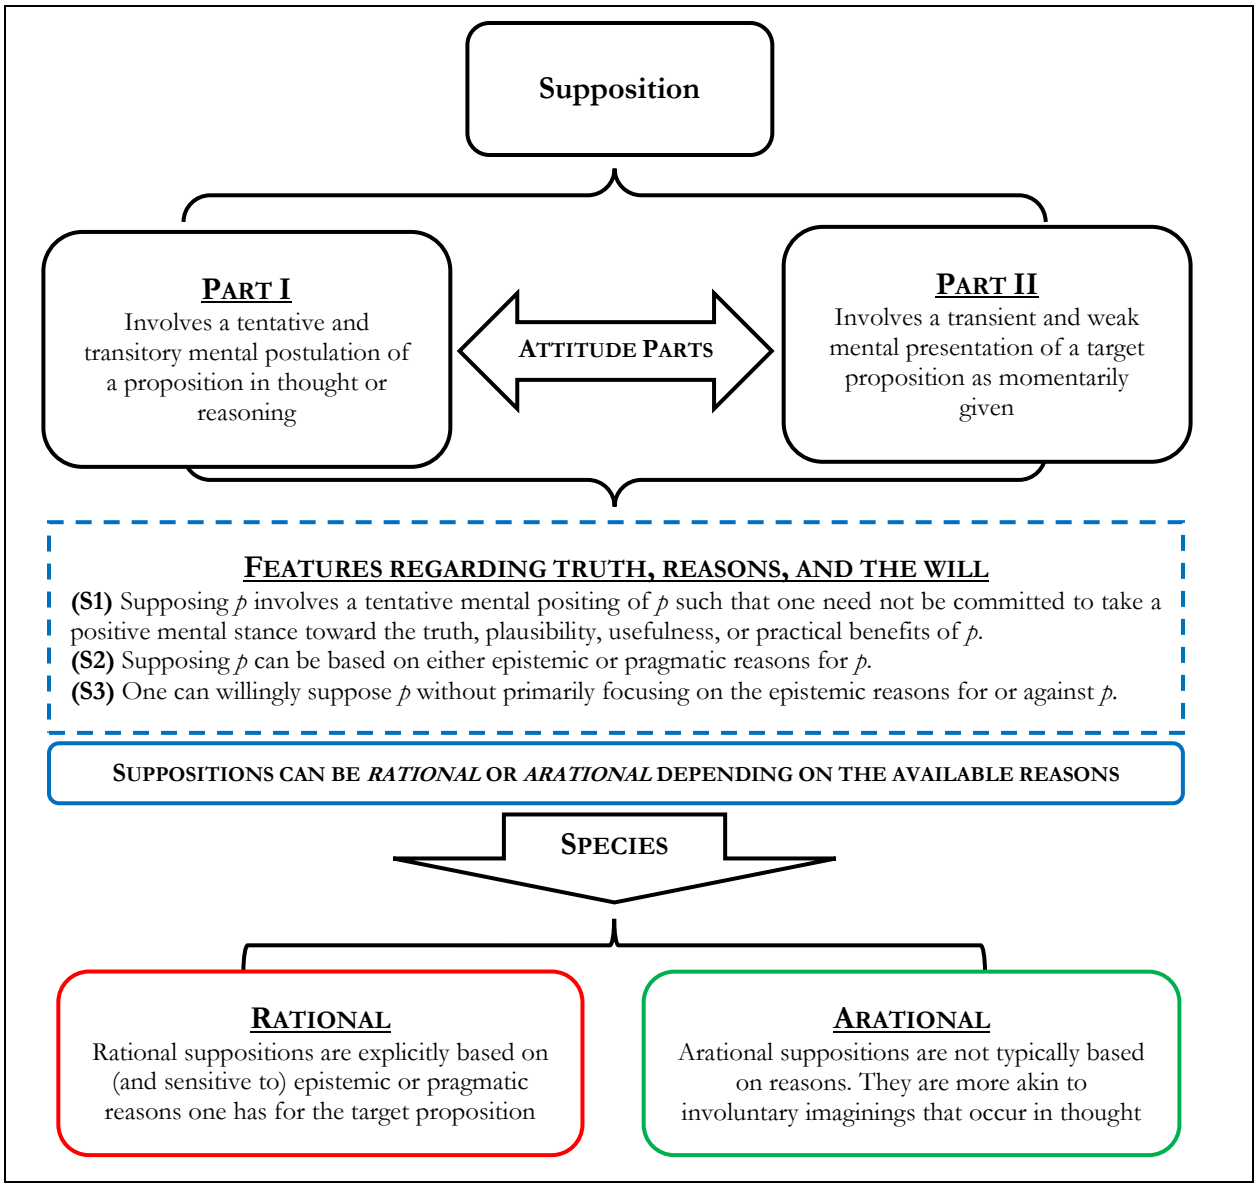

Figure 4: A Diagram of Supposition

NOTES	
Attitude Parts	Parts I and II cohere with Cohen's (1992) view. Part I distinguishes supposition from acceptance. Acceptance involves a long-term commitment to mentally use a proposition. Part II distinguishes supposition from belief and acceptance. Supposing *p* need not involve a committed positive mental stance toward *p*.
Species	*Rational* suppositions can overlap with rational acceptances. *Arational* suppositions are typically an involuntary affair that need not be rational or irrational.

Table 11: Notes on Supposition