

A STUDY OF BELIEF

By

ELISABETH K. I. SHORTSLEEVE

A DISSERTATION PRESENTED TO THE GRADUATE SCHOOL
OF THE UNIVERSITY OF FLORIDA IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

UNIVERSITY OF FLORIDA

2009

© 2009 Elisabeth K. I. Shortsleeve

For Kevin and for Olga

ACKNOWLEDGMENTS

I wish to thank the members of my committee, Drs. John Biro, Kirk Ludwig and Gene Witmer, as well as Drs. Robert D'Amico and Greg Ray, for their teaching, patient guidance and friendship. They have nurtured my love of philosophy. I wish to thank also the many graduate students in philosophy I met while at the University of Florida and with whom I have had many a fascinating, enlightening, and sometimes plain silly conversation, especially Bill Butchard, Cathy Galko, Jon Hendrix and Kelly Trogdon. Philosophizing in their company was a delight. I wish to thank the Philosophy Department at the University of Florida, where I have learned much of what I know about analytic philosophy, and more. I deeply regret that the University of Florida administration has failed to appreciate the excellence and importance of its Philosophy Department. I wish to thank Virginia Dampier, our departmental Secretary, for her kindness, warmth and help. Finally, I wish to thank my husband, Kevin, and my mother, Olga, for their support during these long disserting years.

TABLE OF CONTENTS

	<u>page</u>
ACKNOWLEDGMENTS.....	4
ABSTRACT	7
CHAPTER	
1 SYNOPSIS	9
2 THE AIM AND FUNCTION OF BELIEF	19
Constitutive Aims	19
<i>Prima Facie</i> Case for the Thesis that Belief Constitutively Aims at Truth	24
Belief is Necessarily Governed by the Norm of Truth	25
Persuading others	27
Belief attribution	28
Testimony	29
Memory	35
Doxastic Evaluation and Practical Interests	38
The Transparency of Belief	43
Belief and Believer Evaluability	47
Theoretical Refinements	48
Should the Thesis be Understood Literally?	48
Concern that belief is not intentional	49
Concern that belief is not an event, but a state	57
Believing and Guessing	59
The Norm of Truth and the Concept of Belief	62
3 THE COGNITIVE CONTENT OF BELIEF	65
Concepts as Types of Thought Content	65
Complex Concepts and Simple Concepts	66
Informational Content, Linguistic Content and Cognitive Content	68
Directly Referring Concepts, Predicative Concepts and Logical Concepts	74
Synchronic Cognitive Content vs. Individuating Cognitive Content	76
Beliefs About Non-Existent Particulars	81
Beliefs about Misidentified Particulars	84
Concepts as Individuated by their Possession Conditions	85
Requirements on Concept Possession Conditions	88
Vagueness	93
Belief Possession Conditions vs. Belief Attribution Conditions	105
The Concept of Belief	109

4	RATIONAL BELIEF	114
	Rational Evaluability and Control	115
	Reflective Control: Motivating vs. Regulating.....	122
	Regulating Control: Practical vs. Doxastic	125
	Rationally Evaluable Belief: Minimal Requirements	134
	The Deliberative Stance	140
5	TWO APPLICATIONS	147
	<i>De Re, De Dicto</i> and <i>De Se</i> Attitude Reports.....	147
	Some Definitions	147
	Desiderata	148
	Theoretical Tools	149
	Account	152
	<i>De Re</i> and <i>De Dicto</i> belief reports	152
	<i>De Se</i> belief reports	156
	Assessment	160
	Gettier's Cases	162
	Conclusion.....	169
	LIST OF REFERENCES	170
	BIOGRAPHICAL SKETCH.....	173

Abstract of Dissertation Presented to the Graduate School
of the University of Florida in Partial Fulfillment of the
Requirements for the Degree of Doctor of Philosophy

A STUDY OF BELIEF

By

Elisabeth K. I. Shortsleeve

August 2009

Chair: Kirk Ludwig
Cochair: D. Gene Witmer
Major: Philosophy

This study proposes that belief is a psychological state individuated by its aim, function and individuating cognitive content, and that the authority of the doxastic norms to which belief is subject is underwritten by the nature of belief. According to this account, belief aims at truth constitutively, in the sense that the norm of truth necessarily governs believing, understood as an event, in virtue of the proper cognitive function of believing, i.e., to generate and sustain belief, and in virtue of the proper cognitive function of belief, i.e., to sanction and preserve information as such in conceptualized form. The cognitive content of belief is wholly constituted by concepts in given modes of combination, where concepts are conceived as types of thought content wholly individuated by their possession conditions. Concepts can be categorized into distinct types—directly referring, predicative, and logical concepts—depending on the type of contribution they make to the cognitive content of thoughts containing them. Since simple directly referring concepts do not contribute any information about their referent to the cognitive content of thought, they are typically tokened as constituents of complex directly referring concepts, together with predicative concepts tokened in an identifying role—i.e., identifying concepts—used to track the referent of one’s thought. What identifying concepts one uses to track the

referent of one's thought might change over time, while one's thought about that referent might remain unchanged. Accordingly, identifying concepts do not belong to the cognitive content that individuates a belief over time (i.e., its individuating cognitive content), but only to the cognitive content this belief contributes at a time to the cognitive processes of which it is a part (i.e., its synchronic cognitive content). A belief is evaluable as rational or irrational only if the believer has the capacity to take a deliberative stance toward it, which minimally requires the capacity to entertain the deliberative question, 'Do I believe that p ?' This requires the capacity to entertain a question and possession of the first-person concept and the concept of belief, in addition to the concepts that constitute the individuating and current synchronic cognitive contents of that belief.

CHAPTER 1 SYNOPSIS

The aim of this study is to articulate and defend the thesis that belief is a psychological state individuated by its aim, function and cognitive content, and to argue that the authority of the doxastic norms to which belief is subject is underwritten by the very nature of belief.

Chapter 2 articulates and defends the thesis that belief aims at truth constitutively. I start by explicating the notion of a constitutive aim in relation to the practical realm, reflecting on the paradigmatic case of shooting at the bull's eye. This reflection uncovers the following features as characteristic of any action-description formulated in terms of a constitutive aim:

Let A be an agent, d a description, ϕ an action and φ an aim such that d describes ϕ in terms of φ and describes ϕ as done by A .

Then:

- (i) Necessarily, d is true only if A intended to φ (though A need not have intended to φ consciously).
- (ii) ϕ , described as d , is subject to evaluation along dimensions of success and performance independently of any other aim and desire A might have.
- (iii) A is subject to corresponding evaluation and to praise or blame as a result of this evaluation, at least *prima facie*.

I propose that the claim that belief constitutively aims at truth should be understood along similar lines. To say that belief aims at truth constitutively is to say that belief is necessarily governed by truth as a regulating principle (which need not be consciously endorsed). In other words, to say that belief aims at truth constitutively is to say that belief is necessarily governed by the norm of truth. Further, this claim casts belief as subject to evaluation relative to truth along the dimensions of success and performance independently of any other aim and desire the believer might have. Finally, this claim casts the believer as subject to corresponding evaluations and attendant praise or blame, at least *prima facie*.

I then offer a *prima facie* case for the thesis that belief aims at truth constitutively, which underscores its explanatory power. I argue that this thesis explains:

- (i) why belief, under normal conditions, is responsive to (the appearance of) truth —belief is so responsive because it is necessarily governed by the norm of truth; in turn, the thesis that belief, under normal conditions, is responsive to (the appearance of) truth explains:
 - (a) our success at securing interpersonal agreement as to what to believe and do;
 - (b) our success at attributing beliefs to others while having obtained few or no relevant verbal reports from them;
 - (c) why utterances that are properly taken to express belief (defeasibly) entitle their audiences to believe that these utterances are true and to form beliefs on their basis;
 - (d) why individuals can remain entitled to their memory beliefs while having forgotten their reasons for these beliefs;
- (ii) why beliefs, unlike, say, imaginings, are liable to criticism as such for failing to respond to (the appearance of) truth under normal conditions;
- (iii) why the deliberative questions, ‘Do I believe that *p*?’ and ‘Shall/should I believe that *p*?’ are transparent to the question, ‘Is it the case that *p*?’ in the sense that the first two questions typically are, and should be, answered by way of answering the third question;¹
- (iv) why assertions of the forms, ‘I believe that *p*, but *p* is false’ and ‘*p*, but I don’t believe that *p*,’ strike us as paradoxical (Moore’s paradox).

I then argue that, interpreted literally, this thesis commits two category mistakes, since no state, as such, can properly be said to have, or be regulated by, any aim, and since belief, as such, cannot be purposive. I then propose an interpretation of the thesis that is free of these mistakes and according to which the norm of truth governs believing, understood as an event, in virtue of the proper cognitive function of believing, i.e., to generate and sustain belief, and in virtue of the proper cognitive function of belief, i.e., to sanction and preserve information as such in conceptualized form.

¹ Richard Moran introduces and discusses the notion that belief is so transparent in Moran 2001: 60 ff.

Chapter 3 argues that the cognitive content of belief (i.e., the content that belief, as such, contributes to the cognitive processes of which it is a part) is wholly constituted by concepts in given modes of combination, where concepts are conceived as types of thought content, wholly individuated by their possession conditions, which specify the conditions under which a thinker who possesses the relevant concept finds transitions to and from thoughts containing this concept primitively compelling.² Concepts can be simple or complex, and the possession conditions of complex concepts are a function of the possession conditions of their simple component concepts and their modes of combination. Further, concepts can be categorized into distinct types depending on the type of contribution they make to the cognitive content of thoughts containing them. The cognitive role of a directly referring concept is to contribute a referent for predicative thought. The cognitive role of a predicative concept is to contribute information concerning the instancing of a given property. The cognitive role of a logical concept is to operate on elements of thought. The distinction among directly referring concepts, predicative concepts and logical concepts provides an exhaustive taxonomy of these types.

Simple directly referring concepts, as such, do not contribute any information about their referent to the cognitive content of thought (contributing information to the cognitive content of thought is the role of predicative concepts). For this reason, they are typically tokened as constituents of complex directly referring concepts, together with predicative concepts tokened in an identifying role—i.e., identifying concepts—used to identify and recognize the referent of one's thought. What identifying concepts one uses to identify and recognize the referent of one's

² With some modifications, I am endorsing the account of concepts Christopher Peacocke proposes in *A Study of Concepts*.

thought might change over time, while one's thought about that referent might nevertheless remain unchanged. This suggests the following set of theses:³

Let there be two cognitive contents whose constituents are:

(i) **a[the X] is F**

(ii) **a[the Y] is F**

where 'a' stands for a simple directly referring concept,
'F', 'X' and 'Y' stand for distinct predicative concepts,
and 'X' and 'Y' are tokened in an identifying role;

let (i) be the cognitive content of some belief state, b_1 , at some time, t_i ,

let (ii) be the cognitive content of some other belief state, b_2 , at some time, t_j ,

Then:

- (a) b_1 and b_2 have the same individuating cognitive content, namely, **a is F**, where 'individuating cognitive content' stands for the cognitive content that individuates a belief over time;
- (b) if some believer, S , has b_1 at t_i , then, at t_i , the canonical way for S to verify that a is F is by way of verifying that the X is F (the same is true, *mutatis mutandis*, for b_2);
- (c) having b_1 and b_2 need not place S in a position to judge that the $X =$ the Y .

I propose that identifying concepts do not belong to the cognitive content that individuates a belief over time (i.e., its individuating cognitive content), but only to the cognitive content this belief contributes at a given time to the cognitive processes of which it is a part (i.e., its synchronic cognitive content). In set-theoretic terms, the difference between the individuating and synchronic cognitive contents of belief is fully constituted by identifying concepts—if the synchronic cognitive content of a belief includes no identifying concepts, its synchronic and individuating cognitive contents coincide. I define the individuating and synchronic cognitive contents of belief as follows:

³ Henceforth, I will use a distinct font when using schematic letters to stand for concepts. Unless otherwise indicated, corresponding schematic letters written in normal font stand for the objects to which these concepts refer or the properties or relations which these concepts predicate (as appropriate). Likewise, expressions written in a distinct font stand for the concepts which these expressions express.

Individuating Cognitive Content:

Let b be a particular belief state such that being in b requires tokening concepts c_1-c_n , and let k be the knowledge that anyone who possesses c_1-c_n has just by virtue of possessing those concepts.

Belief b is individuated by the content that p just in case from k and the supposition that b is true, it is *a priori* derivable that p .

Any two particular belief states, b_1 and b_2 , that meet the above condition count as being of the same individuating type (i.e., as being the same belief, where ‘belief’ is used to refer to a type of state rather than to a particular state) by virtue of having the same individuating cognitive content.

Synchronic Cognitive Content:

Let b be a particular belief state individuated by the content that p , and let S have b at t .

If the individuating cognitive content of b can be represented as **... a is F ...**, then the synchronic cognitive content of b at t can be represented as **... a [the X] is F ...** just in case the canonical way for S to verify at t that a is F is for S to verify that the X is F .

Whenever two particular belief states, b_1 and b_2 , meet the above condition, b_1 and b_2 have the same synchronic cognitive content (i.e., b_1 and b_2 make the same type of contribution, while they are undergone, to the cognitive processes of which they are a part).

I stated at the outset that I intended to articulate and defend the thesis that belief is a psychological state individuated by its aim, function and cognitive content. Given the foregoing, this thesis should now be revised as follows: belief is a psychological state individuated by its aim, function and individuating cognitive content.

On the account I am proposing, then, the synchronic and individuating cognitive contents of belief are complex concepts, wholly individuated by their possession conditions, which are a function of the possession conditions of their simple component concepts and their mode of combination. Given that my account posits the existence of directly referring concepts, it faces a challenge akin to the challenge faced by semantic theories that posit the existence of directly referring expressions. Further, my account must also answer the question, what are the truth conditions of beliefs whose synchronic cognitive content includes a complex directly referring concept in cases when the referent of their component simple directly referring concept does not

fall under their component identifying concept? Chapter 3 addresses both challenges, and shows that neither is fatal to the account I am advocating.

Concerning the nature of concepts generally, I endorse the general approach Peacocke proposes in *A Study of Concepts*, modifying his account whenever doing so is necessary for my purposes. On Peacocke's view, concepts are wholly individuated by their possession conditions, which specify the conditions under which a thinker who possesses the relevant concept finds transitions to and from thoughts containing this concept primitively compelling. Concerning this account, I raise the following two questions:

- (i) What does it take for a set of possession conditions to be the possession conditions of a concept?
- (ii) Does there correspond a single concept, or a multiplicity of concepts, to a vague expression, and is(are) the relevant concept(s) also vague?

Peacocke offers one requirement on concepthood, which he leaves rather underspecified. I interpret it as a requirement that the possession conditions of any concept must determine consistent deployment conditions. Reflection on this requirement reveals that Peacocke is committed to the view that there is no concept truth. On this matter, I endorse Peacocke's commitment but coin the term, 'notion,' for sets of possession conditions that fail to determine consistent deployment conditions and argue that notions can nevertheless play a significant role in our cognitive lives. Further reflection reveals that Peacocke's requirement, interpreted as a requirement on deployment conditions, should be restricted to simple concepts.

Peacocke's own formulation of his requirement on concepthood is as follows:

The determination theory for a given concept (together with the world in empirical cases) assigns semantic values in such a way that the belief-forming practices mentioned in the concept's possession condition are correct. That is, in the case of belief formation, the practices result in true beliefs, and in the case of principles of inference, they result in truth-preserving inferences, when semantic values are assigned in accordance with the determination theory. (Peacocke 1992: 19)

Reflection on vague concepts reveals that this requirement should be reformulated as follows:

The determination theory for a given concept (together with the world in empirical cases) assigns semantic values in such a way that the belief-forming practices mentioned in the concept's possession condition are correct. That is, in the case of belief formation, the practices result in information-carrying beliefs, and in the case of principles of inference, they result in information-preserving inferences, when semantic values are assigned in accordance with the determination theory.

Finally, Chapter 3 offers a conceptual argument against the view that the content of belief is indeterminate.

Chapter 4 is devoted to the question, what is required for a belief to be subject to evaluation as rational or irrational? I propose that a belief is evaluable as rational or irrational only if the believer has the capacity to take a deliberative stance toward it. Taking a deliberative stance toward the belief that *p* minimally requires the capacity to entertain the deliberative question, 'Do I believe that *p*?'⁴ This requires the capacity to entertain a question and possession of the first-person concept and the concept of belief, in addition to the concepts that constitute the individuating and current synchronic cognitive contents of that belief.

Intuitively, liability to normative praise and blame, which attends rational evaluability, requires a dual form of control, i.e., it requires causal responsibility and the capacity to abide by the norm relative to which one is evaluated. I argue that this intuition is correct, and that the type of control exercised in deliberation (both practical and doxastic), i.e., regulative control, suffices to secure the required control. Doxastic deliberation, I argue, minimally requires the capacity to take a deliberative stance toward one's belief. By virtue of possessing the concepts that constitute the individuating and current synchronic cognitive contents of her belief, by virtue of being

⁴ I am borrowing Moran's terminology here. According to Moran, the deliberative question, 'Do I believe that *p*?' is equivalent to the question, 'Should I believe that *p*?' and does not aim merely to yield an accurate attitude report; it aims to yield belief—i.e., an attitude that conforms to what Moran calls 'the condition of transparency,' i.e., "that I can report on my *belief* about X by considering (nothing but) X itself" (Moran 2001: 84; italics in the original; see also 63).

capable of conceptualizing it as her own belief, and by virtue of being capable of wondering whether to believe it, the believer has the capacity to recognize that a given set of doxastic norms governs her belief, in the sense of having the capacity to recognize what it would take for her belief to be true and to recognize her belief as beholden to these truth conditions. By virtue of being a conscious, purposive activity, deliberation also secures the capacity to refrain from abiding by these norms by making it possible to subject belief to intentional governance. Accordingly, satisfaction of the minimal requirements for the capacity for doxastic deliberation secures the control intuitively required for rational evaluability.

Chapter 5 applies the foregoing account to two issues, namely, the issue raised for semantics by *de re*, *de dicto* and *de se* attitude reports, and the issue raised for epistemology by Gettier's cases.

I propose that a purely *de re* belief report states the individuating cognitive content, while a purely *de dicto* belief report states the synchronic cognitive content, of the belief it reports. Further, I propose that reflexive beliefs are unusual in the following respect: their synchronic cognitive content may refer to the believer via simple directly referring concepts tokened independently of any identifying concepts. Generally, genuinely *de se* beliefs are reflexive beliefs that do not identify the believer via any predicative concept. I propose that *de se* belief reports are offered *de re* with regard to any reflexive expression included in their complement clause.⁵

My answer to Gettier's cases is as follows:

I propose that having knowledge requires satisfying the norm of truth as it applies to one's belief. Now, satisfying the norm of truth as it applies to any belief is a matter of satisfying the deployment conditions of the simple concepts, in the relevant modes of combination, that

⁵ For precise definitions of what I mean by *de re* and *de se* attitude reports, see Chapter 5.

constitute the synchronic and individuating cognitive contents of that belief. The deployment conditions of any concept mention the types of circumstance under which this concept is correctly deployed. It is highly plausible that the norm of truth sets a dual requirement on deployment of any concept: First, it sets a correctness requirement, by requiring that circumstances of the relevant type obtain and, second, it sets an evidence requirement, by requiring that one who deploys this concept has evidence that circumstances of the relevant type obtain. I propose that having knowledge requires satisfying the correctness and evidence requirements for the simple concepts, in the relevant modes of combination, that constitute the synchronic and individuating cognitive contents of that belief.

I propose that Gettier's first case⁶ involves a misidentified belief, i.e., a belief whose synchronic cognitive content has the form, ***a*[the *X*] is *F***, though *a* is not the *X*. Misidentification and falsehood are epistemic defeaters for analogous reasons. To evaluate a belief as false is to evaluate it negatively along the dimension of correctness because its individuating cognitive content fails to satisfy its correctness requirement by virtue of being constituted by some concept that fails to satisfy its correctness requirement. To evaluate a belief as subject to an error of misidentification is to evaluate it negatively along the dimension of correctness because its synchronic cognitive content fails to satisfy its correctness requirement by virtue of being constituted by some predicative concept, tokened in an identifying role, that fails to satisfy its correctness requirement.

I propose that Gettier's second case⁷ involves a belief that is epistemically flawed because it fails to satisfy the correctness and evidence requirements for the simple concepts, in the relevant modes of combination, that constitute individuating cognitive contents of that belief. The

⁶ For a detailed account of Gettier's first case, see Chapter 5.

⁷ For a detailed account of Gettier's second case, see Chapter 5.

relevant belief is a disjunctive belief. In Gettier's second case, the believer satisfies the correctness requirement relative to one component disjunct, and the evidence requirement relative to the other component disjunct. Satisfying the correctness and evidence requirements for the component concepts, in the relevant mode of combination, however, requires satisfying both the correctness and the evidence requirement for at least one of the conjuncts.⁸

The account that follows means to be intuitive and as fine grained as our pre-theoretic thoughts about our beliefs. Although some of the terminology and apparatus it introduces is complex and at times technical, the view it means to capture is simple.

⁸ This answer does not apply to tautological truths such as ' $p \vee \sim p$.'

CHAPTER 2 THE AIM AND FUNCTION OF BELIEF

That belief aims at truth is well nigh a philosophical platitude. In “Deciding to Believe,” Bernard Williams isolates five features that he takes to be characteristic of belief. The first of these is that belief aims at truth. Williams explicates this claim as follows:

When I say that beliefs aim at truth, I have particularly in mind three things. First: that truth and falsehood are a dimension of an assessment of beliefs as opposed to many other psychological states or dispositions ... [Second]: to believe that *p* is to believe that *p* is true ... [Third]: to say ‘I believe that *p*’ itself carries, in general, a claim that *p* is true.
(Williams 1973: 136-37)

Williams’ claim that belief aims at truth is frequently echoed in the philosophical literature, usually in a strengthened form—truth, it is thought, is the constitutive aim of belief. In this chapter, I explicate the notion of a constitutive aim and canvass a *prima facie* case for endorsing the stronger thesis. I then argue that, interpreted literally, this thesis commits two category mistakes, since no state, as such, can properly be said to have, or be regulated by, any aim, and since belief, as such, cannot be purposive. I then propose an interpretation of the thesis that is free of these mistakes and according to which the norm of truth governs believing, understood as an event, in virtue of the proper cognitive function of believing, i.e., to generate/sustain belief, and in virtue of the proper cognitive function of belief, i.e., to sanction and preserve information as such in conceptualized form.

Constitutive Aims

To start, something must be said about what the expression, ‘constitutive,’ adds to the thesis that belief aims at truth. The difference between constitutive and non-constitutive aims is most conspicuous in the realm of action. So, consider my shooting for the bull’s eye while playing at darts. My aim in doing so may be to win the game. This aim, however, is not constitutive of my action described as a shot for the bull’s eye, in the sense that I could have performed

this action without aiming to win the game—having scored a bull’s eye already and knowing that I will get no further points for hitting another, I may nevertheless shoot for another to show off my skill and lack of concern with winning. What is required for my action to count as a shot for the bull’s eye is that I performed it, *inter alia*, with the purpose of hitting the bull’s eye. Hitting the bull’s eye, then, is a constitutive aim of shooting for the bull’s eye, in the sense that I cannot shoot for the bull’s eye without having this aim. More precisely, in the realm of action, constitutive aims are subject to the following condition:

For any agent, A , action-type, Φ , action ϕ such that ϕ is a token of Φ , and aim, ψ , ψ is constitutive of Φ just in case, necessarily, if $A \phi$ s, then $A \phi$ s with the purpose of ψ ing.

On most accounts, what distinguishes an action from a mere bodily movement is that there is a description of the former, but not of the latter, under which it is performed intentionally. Any given action can be described in many different ways. Borrowing Donald Davidson’s example (Davidson 1980: 4), flipping the switch can sometimes also be described as turning on the light, or as alerting a prowler to one’s presence. Different types of action-descriptions have different types of logical profiles. A description of an action formulated in terms of a consequence of that action is factive relative to that consequence. So, I cannot truly be said to be turning on the lights unless the lights are thereby turned on, nor can I truly be said to be alerting a prowler to my presence unless a prowler is thereby so alerted. Such a description, however, has no specific implications concerning the agent’s intention in undertaking the action (although given that it describes the event as an action, it does imply that the agent undertook it with some intention). From the knowledge that Amy turned on the lights, I may not properly infer that Amy intended to turn on the lights—she might have meant to turn the lights on, but she might also have meant merely to test the switch. The converse holds for action-descriptions formulated in terms of a constitutive aim. Such action-descriptions need not be factive relative to what the

action achieves, but they do entail that the action was intended as described. So, my shot counts as a shot for the bull's eye even if my dart misses the board by a wide margin, but does not count as such if I hit the bull's eye while intending to hit a different target.

Although all actions are intentional under some description, the etiology of some actions may fail to involve any conscious intention to perform them. So, pacing is often done spontaneously—while puzzling over some particularly absorbing issue, I can be quite surprised to find myself marching around my room, having no memory of getting up from my seat. Nevertheless, unlike twitching and digesting, pacing is something that the agent does, rather than something that happens to him/her. I will use the term, 'deliberate action,' to refer to actions whose etiology involves a conscious intention to perform them (under some description), and the term, 'spontaneous action,' to refer to actions that involve no such conscious intention. Note that a chain of actions may be deliberate even though some of the acts that constitute it are spontaneous. So, I may deliberately ride my bicycle to grandma's while being incapable of entertaining separate conscious intentions concerning the myriad momentary acts in which such riding consists (my trying to form such conscious intentions would likely not bode well for my riding).

An action that is properly describable in terms of a constitutive aim need not be deliberate. Just as I can pace spontaneously, so too can I spontaneously shoot for the bull's eye. If shooting darts were a behavior in which I tended to engage while deep in thought, my shots would not be aimless for being undertaken absent-mindedly, nor would it be incorrect to describe a shot I had absent-mindedly aimed at the bull's eye as a shot for the bull's eye. An action-description formulated in terms of a constitutive aim does not, as such, attribute a conscious intention to the agent. Rather, it describes the action as governed by that aim as a regulating principle, which may, but need not, be consciously endorsed.

An action-description formulated in terms of a constitutive aim *ipso facto* casts the action it describes as subject to evaluation along two dimensions. First, the action can be evaluated as carrying out this aim either successfully or unsuccessfully (I will call this dimension of evaluation ‘the success dimension’). So, my shot for the bull’s eye is successful if and only if my dart hits the bull’s eye. Second, the action can be evaluated as carrying out this aim either well or poorly (I will call this dimension of evaluation ‘the performance dimension’). So, even if my dart reaches its target, my shot might nevertheless be poor if the manner in which I threw the dart lent me but little control over its course, or if the success of my shot depended crucially on the occurrence of some event over which I had no control, such as a twitch or a nudge. Conversely, my shot might be good even if it misses its target. Such would be the case if the manner in which I threw the dart lent me good control over its course and I missed the target by a short distance only. Further, such an action-description casts the action it describes as subject to evaluation along these two dimensions independently of any other aims and desires the agent might have. If I shot for the bull’s eye while knowing that my doing so would result in the frustration of all of my other aims and desires, my action would be profoundly irrational, but it would not be any less successful, or any worse, as a shot for the bull’s eye for that.

Agents are responsible for their actions—though of course this responsibility can be mitigated—presumably because agents are in a unique position to exert control over their actions via their intentions, and over their intentions via their practical judgments. Because an action-description formulated in terms of a constitutive aim attributes an intention to the agent, it casts not only the action it describes as subject to evaluation along the dimensions of success and performance mentioned above, but also the agent as liable to corresponding evaluation and attendant praise or blame. An evaluation of an action described as a shot for the bull’s eye

reflects on the agent—a successful, flawless shot for the bull’s eye is, as such, a successful, flawless performance of the agent and, as such, entitles her to praise, at least *prima facie*. By contrast, an agent’s alerting of a prowler to her presence may be said to be especially effective without the agent being thereby painted as effective or praiseworthy—the effectiveness of the agent’s action reflects on the agent only if the agent intended her action under this description.

Although the foregoing reflection has focused on the perhaps rather overly tailored example of shooting at the bull’s eye, I believe that it nevertheless uncovered features characteristic of any action-description formulated in terms of a constitutive aim. Any action-description that is not formulated in terms of a constitutive aim can be turned into one that is formulated in terms of a constitutive aim by prefacing it with the phrase, ‘trying to’.^{1,2} Consider the action-description, Amy turned on the light. As I noted earlier, this action-description is formulated in terms of a consequence of the action. Now consider the action-description, Amy tried to turn on the lights. Clearly, necessarily, this description is true only if Amy intended to turn on the lights. Further, her action is subject to evaluation both in terms of success and performance independently of any other aims and desires Amy might have. Furthermore, however we evaluate Amy’s action along these dimensions reflects on Amy, at least *prima facie*—whatever praise or blame there might be for succeeding or failing to turn on the lights, or for doing it well or poorly, as the case might be, belongs by right to Amy.

¹ Although uses of the phrase, ‘trying to’ often carry in English a conversational implicature of failure (by virtue of the maxim of relevance, I suspect), the phrase does not entail failure, as indicated by Grice’s cancellability test. To wit, the sentence, ‘Elka tried to finish her dissertation and succeeded.’

² Action-descriptions using the phrase ‘trying to’ may at times pick out mental acts involving no bodily component. ‘Trying to,’ however, is not synonymous with ‘intending to,’ or with any other mental-act verb. If Amy tried to turn on the lights and failed upon being struck by complete paralysis, it might very well be that her trying involved no further event than her intending (and perhaps exerting mental effort) to turn on the light. By contrast, the description of Amy’s action as ‘Amy tried to turn on the lights’ in the case when Amy succeeded applies to her entire action, bodily movements included.

Let us set aside for later consideration the worries that belief is not an action but a state and that the etiology of belief often fails to involve any conscious or unconscious intention to believe truly, as well as the related question of whether the claim that belief constitutively aims at truth should be understood metaphorically or literally. For now, the foregoing discussion suggests that this claim should be explicated as follows: To say that belief aims at truth constitutively is to say that belief is necessarily governed by truth as a regulating principle (which may, but need not, be consciously endorsed). In other words, to say that belief aims at truth constitutively is to say that belief is necessarily governed by the norm of truth.³ Further, this claim casts belief as subject to evaluation relative to truth along the dimensions of success and performance independently of any other aims and desires the believer might have. Finally, this claim casts the believer as subject to corresponding evaluations and attendant praise or blame.

Prima Facie Case for the Thesis that Belief Constitutively Aims at Truth

My case for this thesis, explicated along the lines delineated above, underscores its explanatory power. In what follows, I argue that:

The thesis that belief aims at truth constitutively explains:

- (i) why belief, under normal conditions, is responsive to (the appearance of) truth—belief is so responsive because it is necessarily governed by the norm of truth;
in turn, the thesis that belief, under normal conditions, is responsive to (the appearance of) truth explains:
 - (a) our success at securing interpersonal agreement as to what to believe and do;
 - (b) our success at attributing beliefs to others while having obtained few or no relevant verbal reports from them;
 - (c) why utterances that are properly taken to express belief (defeasibly) entitle their audiences to believe that these utterances are true and to form beliefs on their basis;
 - (d) why individuals can remain entitled to their memory beliefs while having forgotten their reasons for these beliefs;

³ Henceforth, I will use ‘is governed by the norm of truth’ as interchangeable with ‘is governed by truth as a regulating principle.’

- (ii) why beliefs, unlike, say, imaginings, are liable to criticism as such for failing to respond to (the appearance of) truth under normal conditions;
- (iii) why the deliberative questions, ‘Do I believe that p ?’ and ‘Shall/should I believe that p ?’ are transparent to the question, ‘Is it the case that p ?’, in the sense that the first two questions typically are, and should be, answered by way of answering the third question;
- (iv) why assertions of the forms, ‘I believe that p , but p is false’ and ‘ p , but I don’t believe that p ,’ strike us as paradoxical (Moore’s paradox).

Note, however, that I am not claiming that truth is the exclusive aim of belief; the thesis that belief constitutively aims at truth is compatible with believers having other aims in addition to truth while forming beliefs. The thesis that belief constitutively aims at truth is compatible also with belief having other constitutive aims in addition to truth.

Arguments that appeal to explanatory power are generally tricky, since they do not have compelling persuasive force unless they are supported by an argument to the effect that the proffered explanation is better than any possible alternative explanation of the relevant domain. It is typically difficult to show that some explanation is better than any possible alternative, and I am not in a position to show that there can be no better explanation of the phenomena listed above than the explanation I am proposing—in part because I find it difficult to envision suitable alternatives. Accordingly, I am not offering the following as a *bona fide* argument to the *best* explanation. Rather, I am arguing that the thesis that belief constitutively aims at truth figures in a good explanation of a wide range of phenomena. I am nevertheless hoping that the persuasive force of the argument will not be insignificant, given the wide range of the phenomena involved in the explanation.

Belief is Necessarily Governed by the Norm of Truth

Suppose Basil asked Amy, ‘Are there any bears in the yard?’ and Amy responded, ‘Yes.’ Under normal conditions, that Amy so responded would lead and warrant Basil to attribute to Amy the belief that there are bears in the yard and to believe that there are bears in the yard.

Now suppose Basil discovered that, although Amy understood his question and gave him a sincere response, she had no indication at all that it was true. Clearly, this discovery would rob Basil of any warrant to believe on this basis that there are bears in the yard. Further, suppose Basil also discovered that Amy would not change her mind about the matter in the face of any indication to the contrary she might encounter. It is not clear that, having made these discoveries, Basil would remain warranted in attributing to Amy the belief that there are bears in the yard. What prompted Amy's 'Yes,' it seems, is more akin to imagining or fearing than believing that there bears in the yard.⁴

That belief aims at truth constitutively explains why beliefs, under normal conditions, are responsive to (the appearance of) truth. Under normal conditions, mere awareness of my seeing a bear before me suffices for me to form, immediately and without any willful effort, the belief that there is a bear before me (no matter how strong my desire that there be no bear before me). Similarly, mere awareness of my failing to see a bear before me suffices for me to relinquish,

⁴ In an argument meant to cast doubt on the claim that believers are motivated by the norm of truth as a matter of conceptual necessity, Eric Funkhauser adduces three clinical cases of belief that is allegedly impervious to evidence (Funkhauser 2003: 186-187). Funkhauser mentions these cases under the heading, "Believing with complete indifference to the evidence." In doing so, however, Funkhauser misrepresents the assessment of the authors of the study he cites, who write:

Together, these three cases—FE, TH and DB—provide convincing evidence that an anomalous perceptual and/or affective experience, which the patient feels compelled to explain, underlies the formation of some [Delusions of Misidentification (DM)] ... We are not suggesting that perceptual or affective deficits of the kind that we observed in FE, TH and DB, are sufficient explanation of their delusions of misidentification. Rather, we are proposing that such deficits may provide necessary antecedents to the formation of DM in some patients. (Breen 2000: 102)

The authors do note, as Funkhauser reports, that these beliefs are pathologically resistant to contrary evidence:

[O]ur observations suggest that ... there is also a failure to be influenced by, or give weight to, evidence which for a normal person would override the influence of the anomalous percept or affective experience in the production and maintenance of belief. It is this failure which enables the DM to be accepted. As a result, all three patients (FE, TH and DB) were able to give untroubled expression, in a single narrative or conversation, to ideas which would be completely contradictory, and therefore troubling, to any cognitively intact person. (*Ibid.*)

Yet Funkhauser's claim that these are instances of "believing with complete indifference to the evidence" is an overstatement. Even in these pathological cases, it seems, the believers are motivated by the norm of truth, albeit very imperfectly. Further, it should be noted that these beliefs are viewed as pathological precisely because they are impervious to the corrective influence of contrary evidence. In this respect, these believers fail to be motivated as they should by the norm of truth. If believing were, as Funkhauser proposes, "simply regarding as true" (Funkhauser 2003: 193), it is not clear why these beliefs should be viewed as pathological.

immediately and without any willful effort, my belief to the contrary. If beliefs merely represented their content as true, but were not governed by the norm of truth, this would be a mystery that wants explaining.

Persuading others

The assumption that our interlocutors' beliefs are responsive to (the appearance of) truth underlies many of our communicative practices and interactions with other agents. When we seek to convince someone that p is true, we assume that our task consists in offering her sufficient doxastic reasons for thinking that p true. The assumption is that mere reflection on these reasons will suffice to elicit her belief. Barring misunderstanding, failure to believe on the part of our interlocutor is rationally acceptable to us only if it is motivated either by her thinking that the reasons we have offered are not true (in which case we seek to convince her that they are true by offering her reasons for thinking that they are) or by her thinking that these are not reasons, or insufficient reasons, for thinking that p is true (in which case we seek to show to her that not- p is inconsistent with, or unlikely given, the reasons we have offered). Similarly, when we seek to motivate someone to do something, we assume that our task consists in offering her reasons for thinking that the proposed action is an effective means to some end that is, all things considered, the most desirable end for her. In this case, the assumption is that mere reflection on reasons for thinking that the end is the most desirable end for her and that the proposed action is an effective means to this end will suffice to elicit in her the appropriate judgments and intention. The general assumption that underlies both strategies is that agreement concerning reasons for judging that p should, and typically does, suffice to secure agreement in judging that p . Our routine success at securing interpersonal agreement as to what to believe and do speaks in favor of this assumption.

Belief attribution

Our interactions with other agents often require that we attribute beliefs to them while having obtained few or no relevant verbal reports from them. In such cases, our task of deciphering their beliefs is not unlike the task facing Davidson's Radical Interpreter, albeit without the luxury of plentiful time and observational data and on a much-reduced scale. In such cases, it seems, we deploy something like Davidson's Principle of Charity. That is, we tend to maximize agreement between us and those to whom we attribute beliefs. In cases when we take ourselves to share the circumstances of, or evidence available to, those to whom we are to attribute beliefs, we simply ascribe to them the beliefs we hold concerning the relevant topics, unless we have some reason to expect that they do not accept these beliefs. When we do not take ourselves to share their circumstances or evidence, we try to imagine what it would be like to be in their circumstances, or what evidence they are likely to have, and ascribe to them the beliefs we would hold, were we in their situation (unless, once again, we have reason to expect they do not accept these beliefs). Concerning perceptual beliefs especially, we take it that common experience of given environmental stimuli warrants attribution of beliefs that these stimuli elicit in us. In doing so, we seem to endorse something akin to Davidson's Principle of Correspondence. Finally, we seem to take coherence to be an exceptionally strong requirement on correct belief attribution.^{5,6}

Although our track record in belief attribution is far from perfect, it is nevertheless quite re-

⁵ Of course, although the Radical Interpreter's triangulation procedure requires her to attribute psychological attitudes, her primary interest is linguistic. Language, by contrast, is typically not at issue in mundane belief attributions. I should note that my endorsement of something like Davidson's Principles of Charity, Correspondence and Coherence as governing mundane belief attribution is not accompanied by an endorsement of Davidson's indeterminacy thesis concerning belief content. For more on Davidson's Principle of Charity, see "Radical Interpretation" (1973; reprinted in 1984); on Charity interpreted as involving Correspondence and Coherence, see "A Coherence Theory of Truth and Knowledge" (1983; reprinted in 2001) and "Three Varieties of Knowledge" (1991; reprinted in 2001). See also Ludwig and Lepore (2005, Ch. 12).

⁶ The belief-attribution procedure sketched above does not mean to favor the Simulation Theory over the Theory Theory of Folk Psychology. On one reading, the procedure sketched above involves taking myself as a model, i.e., assigning the beliefs that I would endorse were I in the other's shoes; on another, equally appropriate, reading, the procedure involves a reference to an internalized folk-theory that requires me to assign beliefs that are true (by my lights).

markably good. Once again, if beliefs were not governed by the norm of truth, this would be a mystery that wants explaining. Further, the very principles that seem to guide our belief-attribution procedures, i.e., what Davidson calls Charity, Correspondence and Coherence, as well as the assumptions that underlie the communicative practices discussed earlier, strongly suggest that we tacitly endorse the view that beliefs are responsive to (the appearance of) truth by virtue of their very nature—in other words, that belief is necessarily governed by the norm of truth.

Testimony

Testimony is another important communicative and doxastic/epistemic practice that speaks in favor of the view that belief is necessarily governed by the norm of truth. Jennifer Lackey offers the following as a general definition of testimony:

T: S testifies that *p* by making an act of communication *a* if and only if (in part) in virtue of *a*'s communicable content, (1) S reasonably intends to convey the information that *p*, or (2) *a* is reasonably taken as conveying the information that *p*.⁷
(Lackey 2006: 193, Lackey and Sosa 2006: 3)

The practice of giving and taking testimony pervades our lives. A moment's introspection reveals that a very significant portion of our beliefs has a testimonial basis. Our practical lives, too, manifest a persistent tendency toward heavy reliance on testimony. I am driving through a foreign city. I need to find a gas station. I see a stranger walking by. I stop and ask him for directions. He obliges, I take him for his word and follow his guidance. This is a normal situation. But why is it reasonable for me to trust his word? This is one of the central questions in

⁷ I have some reservations concerning Lackey's definition. I take 'information' to be factive (a view shared by Dretske, for example). That is, if *a* conveys the information that *p*, then *p*. In Lackey, 2006: n. 11, Lackey argues that lying testimony is possible. If 'information' is factive, however, it is not clear that lying testimony is possible on Lackey's definition (at least in cases when clause (2) is not fulfilled). Indeed, it is questionable whether S can reasonably intend to convey the information that *p* by making an act of communication *a* when S knows that *p* is false (and, so, when S knows that *a* fails to convey the information that *p*)—can S reasonably intend to do something she knows she cannot do? My intuition tends toward a negative answer. Further, it seems to me that clause (2) unfairly slants the table in favor of non-reductive accounts of testimony. On the view that 'information' is factive, if it is reasonable to take *a* as conveying the information that *p*, it is reasonable to trust S's testimony that *p* (the reasonableness at issue here is quite clearly doxastic/epistemic reasonableness). Nevertheless, I will bracket the issue of the facticity of 'information' and take (T) as a working definition of testimony for my limited purposes here.

the epistemology of testimony. Tyler Burge offers his Acceptance Principle as an answer to this question:

A person is *apriori* entitled to accept a proposition that is presented as true and that is intelligible to him, unless there are stronger reasons not to do so, because it is *prima facie* preserved (received) from a rational source, or resource for reason; reliance on rational sources—or resources for reason—is, other things equal, necessary to the function of reason. (Burge 1993: 469)

Burge takes his Acceptance Principle to be underwritten by a connection between linguistic competence, rationality and truth:

A seemingly understood putative assertion is a *prima facie* sign of rational backing for the assertion and rational commitment to truth; and rational backing together with rational commitment constitute a *prima facie* sign of truth. (*Ibid*: 33)

Clearly, Burge recognizes that there is an important connection between belief (as expressed in testimonial assertion) and truth. What I want to ask here is, how strong a connection between belief and truth does Burge's Acceptance Principle require?

Imagine a population of beings just like us except that they have beliefs* instead of beliefs, i.e., states that are just like beliefs except that they are governed by the norms of practical rationality (I will call such beings 'believers*').⁸ To avoid issues concerning language learning, let us assume that their linguistic competence is innate. Given that true beliefs* are typically required for successful action, it is unlikely that their beliefs* would exhibit a far greater tendency toward falsehood than beliefs do (assuming that beliefs are, as I claim, necessarily governed by the norm of truth). Would entitlement to testimonial belief, as defined by the Acceptance Principle, extend to such beings? I believe that the answer is 'no,' for the following reason. Although any testimonial assertion from a believer* could properly be taken as a *prima*

⁸ That is, a reason to believe* that *p* casts believing* that *p* as desirable. Undefeated evidence that *p* is true might be a reason for believing* that *p*, but it need not. Likewise, undefeated, overriding evidence that *p* is true might be a decisive, all-things-considered reason for believing* that *p*, but it need not. Generally, alethic reasons are reasons for the belief* that *p* only if having a true belief* that *p* is desirable. Regarding my use of 'alethic' see n. 9 below.

facie sign of rational backing for the assertion, and even though the belief* it expressed would fairly likely be true, the connection between the rational commitment this belief* expressed and truth would remain an open question—this rational commitment could be of a sort that did not constitute a *prima facie* sign of truth. To determine whether the rational commitment expressed by this belief* was of a sort that did constitute a *prima facie* sign of truth, one would have to consider why the testifier acquired/sustained this belief*—specifically, her relevant aims, desires and circumstances. If the investigation revealed that her belief* was acquired/sustained in light of alethic⁹ considerations, trust in her testimony might be appropriate. Quite clearly, entitlement to testimonial belief, as defined by the Acceptance Principle, does not extend to such beings—for such beings, entitlement to testimonial belief* is certainly not an *a priori* matter.¹⁰ What this thought-experiment shows, I think, is that the *a priori* entitlement to testimonial belief Burge advocates requires an *a priori* entitlement to viewing belief as necessarily governed by the norm of truth.^{11, 12, 13}

Burge's account of testimony is non-reductive, in the sense that it denies that entitlement to testimonial belief reduces to other sources of entitlement, such as sense perception, memory or inference. I focused on Burge's account, but I believe that the argument outlined above general-

⁹ Here and in the remainder of this dissertation, I am using 'alethic' to mean 'having to do with truth' rather than 'designating modalities of truth.' Although my usage is not sanctioned by dictionaries, it is sanctioned by etymology and very convenient for my purposes.

¹⁰ I am ignoring the complicating consideration that such beings might have other interests than truth in seeking testimony and forming testimonial beliefs. What adding this consideration to the above reflection suggests is that believers* might have little use altogether for the practice of taking and giving testimony as we know it. If so, then there is some reason to think that the very existence of our social practice of testimony requires a community of beings whose beliefs are governed by the norm of truth.

¹¹ For the sake of readability, I am using 'assertion' instead of the phrase, 'seemingly understood putative assertion,' that figures in Burge's explanatory principle (*Ibid.*: 33, quoted above). Correspondingly, I am using 'belief^(*)' instead of 'putative belief^(*)' and 'rational commitment' instead of 'putative rational commitment.'

¹² I expect that Burge would agree with the claim that the *a priori* entitlement to testimonial belief he advocates requires an *a priori* entitlement to viewing belief as necessarily governed by the norm of truth. My aim here is to make this requirement obvious.

¹³ Of course, if I received testimony from a being who, unbeknownst to me, happened to be a believer*, I would be *a priori* entitled to believe it on Burge's account (provided I had no overriding reasons to doubt it).

izes to other non-reductive accounts of testimony. So, consider accounts according to which testimony transmits reasons for belief (e.g., Owens 2000). Let us call the principle according to which one is *prima facie* entitled to testimonial belief in virtue of testimony's power to transmit reasons for belief 'the Transmission Principle.'¹⁴ Clearly, entitlement to testimonial belief, as defined by the Transmission Principle, does not extend to our population of believers*, since what one believer* counts as a practical reason for a given belief* might not be counted as such by another. Let us suppose that testimony transmits reasons for belief(*). Then, in order to determine whether to trust some instance of testimony, a believer* would have to learn the reasons why the testifier acquired/sustained the belief* she expressed in her testimony. Only when the testifier's reasons for her belief* are ones that the hearer would also count as reasons to accept this belief* is it appropriate for the hearer to trust this instance of testimony. What the *prima facie* epistemic entitlement posited by the Transmission Principle requires is that nothing but alethic reasons count as reasons for belief. Although this requirement does not straightforwardly entail that belief is necessarily governed by the norm of truth, it strongly suggests it—how else would one explain that nothing but alethic reasons count as reasons for belief?

Now let us see whether an argument can be formulated along similar lines to show that reductive accounts of testimony, too, presuppose (or at least strongly suggest) that belief is governed by the norm of truth.¹⁵ Reductive accounts of testimony, unsurprisingly, do not present a uniform approach to the topic. Fortunately, I will not need to canvass every type of reductive

¹⁴ For the sake of consistency, I am continuing to use the term, 'entitlement,' which Burge carefully distinguishes from 'justification':

The distinction between justification and entitlement is this: Although both have positive force in rationally supporting a propositional attitude or cognitive practice, and in constituting an epistemic right to it, entitlements are epistemic rights or warrants that need not be understood by or even accessible to the subject.

Other authors formulate their accounts in different normative terms, justification included. Since 'entitlement' is logically weaker than 'justification,' I believe that the use of the former in place of the latter in this context is innocuous.

¹⁵ Reductionist accounts of testimony cannot be shown to presuppose that belief is *necessarily* governed by the norm of truth, since the entitlement they posit is *a posteriori*.

approach to testimony here. One very general distinction will suffice, that between global and local reductionism, which Lackey defines as follows:

The first answer, and the stronger version of reductionism—a view sometimes called global reductionism—is that the justification of testimony as a source of belief reduces to the justification of sense perception, memory, and inductive inference. In particular, global reductionists maintain that in order to justifiedly accept a speaker's report, a hearer must have non-testimonially based positive reasons for believing that testimony is generally reliable. The second, and weaker, version of reductionism—often called local reductionism—is that the justification of each particular report or instance of testimony reduces to the justification of instances of sense perception, memory, and inductive inference. Specifically, local reductionists claim that in order to justifiedly accept a speaker's testimony, a hearer must have non-testimonially based positive reasons for accepting the particular report in question. (Lackey and Sosa 2006: 13)

Let us start by considering global reductionism. Although I noted earlier that beliefs* were unlikely to exhibit an extraordinary tendency toward falsehood, let us remember that the entitlements to testimonial belief posited by the Acceptance and Transmission Principles were shown not to extend to believers* because believers* could not assume that their testifiers' reasons for belief* were alethic. With regard to global reductionism, the point can be brought to bear as follows: The evidence concerning the general reliability of testimony available to a believer* includes introspection, which reveals to her that she occasionally has beliefs* that are not grounded in alethic reasons. Further, introspection reveals to her that she only grounds her beliefs* in alethic reasons when having true beliefs* concerning some topic is required by her practical interests. Since she has no epistemic reason to think that testifiers form their beliefs* any differently, she has no epistemic reason to trust any particular instance of testimony unless she has sufficient epistemic reason to think it expresses a belief* whose truth is required by the testifier's practical interests. So, even if she has formed true testimonial beliefs* on many past occasions, she does not have sufficient epistemic reason for believing* that testimony is generally reliable. If she has sufficient epistemic reason for believing* a claim in this neighborhood at

all, it is that testimony expressing beliefs* whose truth is required by the testifier's practical interests is generally reliable.¹⁶ But this falls short of global reductionism. Entitlement to testimonial belief, as defined by global reductionism, presupposes that belief is governed by the norm of truth.

Let us now turn to local reductionism. The argument pertaining to global reductionism entails that in order to be entitled to trust any particular instance of testimony, a believer* would need have sufficient epistemic reason to believe* that it expresses a belief* whose truth is required by the testifier's practical interests. Clearly, this fails to show that on local-reductionist accounts, believers* are never entitled to testimonial belief*. However, the attitude toward testimony that local-reductionist accounts would advocate as epistemically appropriate for believers* is quite different from the attitude that we intuitively think is—and that local-reductionist accounts presumably advocate as—epistemically appropriate for us. This difference in attitude is so significant that it casts some doubt on the identity of the believers*'s testimonial practice with our own practice. Concern with whether a belief expressed in testimony was formed/sustained in light of alethic considerations belies actual testimonial practice. We simply take this for granted, and are presumably epistemically correct in doing so. What worries us is the possibility that this belief might nevertheless be false. It is the appropriateness of this worry that local-reductionist accounts mean to capture,¹⁷ rather than the appropriateness of the additional concern that burdens believers*. Entitlement to testimonial belief, as defined by local reductionism, does not require that belief be governed by the norm of truth. But actual testimonial practice does seem to presuppose it.

¹⁶ I am once again ignoring the consideration that she might have other interests than truth in seeking testimony.

¹⁷ Elizabeth Fricker, a local reductionist of internalist stripe, for example, argues that any view according to which we are *prima facie* entitled to trust any instance of testimony by default offers “an epistemic charter for the gullible and indiscriminating” (Fricker 1994: 126).

Summing up: My discussion of testimony has revealed that non-reductive accounts of testimony, which hold that entitlement to testimonial belief is *a priori*, presuppose, or at least strongly suggest, that belief is necessarily governed by the norm of truth. Reductive accounts of testimony, by contrast, hold that entitlement to testimonial belief is *a posteriori*. Global-reductionist accounts were shown to presuppose that belief is actually governed by the norm of truth. Local-reductionist accounts were not shown to presuppose that belief is governed by the norm of truth. However, they were shown to advocate an attitude toward testimony as epistemically appropriate for beings whose beliefs are not governed by the norm of truth that is vastly different from the attitude we intuitively think is epistemically appropriate for us.

Memory

An argument can be formulated along similar general lines to show that the epistemic appropriateness of trusting our memory beliefs even while having forgotten our reasons for these beliefs presupposes (or at least strongly suggests) that belief is governed by the norm of truth. Many of our beliefs belong in this category. I believe that the Swiss Confederation was founded in 1291, on the lovely Grütli pasture, by representatives of the Cantons of Uri, Schwyz, and Unterwalden, much to the chagrin of the Habsburgs. I believe that hydrogen is the lightest chemical element, and that its nucleus has a single proton and is orbited by a single electron. I no longer remember my reasons for these beliefs. I acquired them during my formal education (among many other beliefs). I certainly consider myself entitled to trust them. Arguably, our tendency to retain our confidence in our memory beliefs while forgetting our reasons for them (until we encounter some reason to doubt them) is highly beneficial by virtue of being economical of our cognitive resources—after all, our capacity to store information in memory is limited:

[T]here is a practical reason to avoid too much clutter in one's beliefs. There is a limit to what one can remember, a limit to the number of things one can put into long-term storage, and a limit to what one can retrieve. It is important to save room for important things and

not clutter one's mind with a lot of unimportant matters... When one reaches a significant conclusion from one's other beliefs, one needs to remember the conclusion but does not normally need to remember all the intermediate steps involved in reaching that conclusion. (Harman 1986: 41-42)

Our relation to our memory beliefs for which we have forgotten our reasons is the intrapersonal counterpart of our relation to the beliefs of those who testify to us. Of course, there are some significant differences between such memory beliefs and testimony. We are likely to have a far better idea of our own general reliability as believers than we do concerning the reliability of those who testify to us, especially when they are strangers. Further, in the case of memory belief, there is a single retained belief, while in the case of testimony, once a testimonial belief is acquired, there are two beliefs—the testifier's and the hearer's—which are token distinct albeit (roughly) type identical. Finally, in the case of testimony but not memory belief, language (or at least some form of communication) plays an essential mediating role. However, none of these differences invalidate the applicability of my earlier argumentative approach to memory beliefs for which one's reasons are forgotten.

On all accounts, entitlement to one's memory beliefs has something to do with the normative history of these beliefs.¹⁸ Further, on all accounts, preservative memory can preserve, but cannot alter, the normative status of a belief.¹⁹ Now, we are self-governing epistemic beings; we

¹⁸ This is so even according to reliabilism. Since memory, as it pertains to beliefs, is a belief-dependent process, the relevant clause in Alvin Goldman's recursive definition of justified belief is:

(6B) If S's belief in *p* at *t* results ("immediately") from a belief-dependent process that is (at least) conditionally reliable, and if the beliefs (if any) on which this process operates in producing S's belief in *p* at *t* are themselves justified, then S's belief in *p* at *t* is justified. (Goldman 1979: 347)

The notion of conditional reliability is explicitly introduced to apply to reasoning and memory, and defined as follows:

A process is conditionally reliable when a sufficient proportion of its output-beliefs are true *given that its input-beliefs are true*. (*Ibid*; italics in the original.)

Setting aside issues concerning how belief-producing/sustaining processes are to be individuated, and assuming that memory is a conditionally reliable process, a memory belief is justified on this view only if it was justified when originally formed (or perhaps reassessed).

¹⁹ I am adopting Burge's distinction between preservative and substantive memory, which he characterizes as follows:

tend to trust our memory beliefs even after we have forgotten our reasons for them, unless we encounter some reason to doubt them; this tendency might even be highly beneficial to us; but this does not settle the question, ‘Is it epistemically appropriate for us to do so and, if so, why?’ Once again, reflection on the epistemic situation of believers* should prove enlightening. Imagine a believer* whose current practical interests require that she have a true belief* concerning p . She has a memory belief* that p , but she does not remember her reasons for believing* that p . Clearly, she is in no position to trust her belief* that p —after all, the reasons for which she believed* that p , or in light of which she revised her belief* that p along the way, might not have been alethic. Given her current alethic interest, and given that she does not remember her reasons for believing* that p , her remembered belief* that p is now useless to her. By contrast, consider my relation to my belief that q , for which I have forgotten my reasons. I can trust my past self—I can trust that her reasons for believing that q were alethic and that she did her epistemic best to acquire/maintain a true belief that q in light of these reasons (given her resources, the time available to her, and her interest in the matter). Had she come across (undefeated) defeating evidence, she would have revised her belief that q . My not remembering my reasons for believing that q , by itself, is not a reason to relinquish my belief that q . It is, however, for the believer* I just considered, once her interest in having a belief* concerning p becomes alethic. Since the only difference between this believer* and me is that I, unlike the believer*, can assume that my state was properly formed in light of alethic reasons, this argument strongly suggests that belief is governed by the norm of truth.²⁰

Substantive memory is an element in a justification; it imports subject matter or objects into reasoning. Purely preservative memory introduces no subject matter, constitutes no element in a justification, and adds no force to a justification or entitlement. It simply maintains in justificational space a cognitive content with its judgmental force. (Burge 1993: 465)

²⁰ For the sake of brevity, I am confining the above argument to the intuitive level. Some authors (Burge 1993, 1997) hold that preservative memory, unlike testimony, does not confer, but merely preserves, epistemic entitlement. On such views, entitlement to memory belief, as such, is neither *a priori*, nor *a posteriori*. My argument only

Doxastic Evaluation and Practical Interests

One might accept the foregoing arguments insofar as they show that belief is actually governed by the norm of truth, but deny that belief is so governed necessarily, or in virtue of the nature of belief. Funkhouser, for example, writes:

The view advocated here minimizes the role that the norm of truth plays for believers as a matter of conceptual necessity. When an agent believes that *p*, she believes that *p* is true. But, she can take *p* to be true all the while knowing that she does not have any reasons or evidence that point to the truth of *p*. I would further suggest that believing is simply regarding as true, regardless of whether the believer is motivated to get the truth-value right. Of course, as a matter of fact, human believers are very much concerned with possessing truth-conducive reasons and this might partially explain our (contingent) inability to believe at will. (Funkhouser 2003: 192-3)

One might claim that although belief is typically responsive to truth, it is so responsive not because of anything peculiar to the nature of belief, but because of something peculiar to the nature of actual human agents. On this view, although belief is not necessarily governed by truth and is not necessarily responsive to truth (or the appearance of truth) under normal conditions, it typically responds to truth (or the appearance of truth) because of its function as a spring for action. It is in the nature of agents to seek to satisfy their desires by forming suitable intentions and carrying them out successfully. Since agency requires belief, and successful action typically requires true beliefs, agents typically acquire and revise their beliefs in the light of truth (or the

shows that such views strongly suggest that belief is actually governed by the norm of truth. Other authors view the epistemic role of preservative memory as analogous to testimony (Owens 2000). Owens endorses a memorial analogue of the Transmission Principle, which can be shown strongly to suggest that belief is necessarily governed by the norm of truth. His view is a version of epistemic foundationalism. Since memory beliefs for which one's reasons are forgotten are typically presented as a difficulty for coherentist accounts of memory, I am ignoring these accounts here. Finally, my argument does not speak to reliabilist accounts. However, reliabilist accounts do not speak directly to the question I am considering here, namely, 'Should we trust our memory beliefs once we have forgotten our reasons for them and, if so, why?' Note, however, that an application of the reliabilist criterion (offered in n. 19 above) to our memory beliefs for which our reasons are forgotten could yield no positive answer to my question unless we are able to assume that the beliefs of our past selves are generally justified. Further, my argumentative approach indicates that, in the context of this answer, the term 'justified' that figures in this criterion should not be construed merely as 'reliable,' but as 'properly formed in the light of alethic reasons.' A positive, reliabilist answer to my question, then, would strongly suggest that belief is actually governed by the norm of truth.

appearance of truth) so that they may satisfy their desires.²¹

Although the concept of belief is correctly characterized in terms of the function of belief as a guide to intentional behavior, the account adumbrated above is missing an important part of the story. It is because belief is responsive to truth that it can play this action-guiding function. More importantly, however, this account fails to explain why beliefs, unlike imaginings for example, are liable to criticism *as such* for failing to respond to truth (or the appearance of truth) under normal conditions. Amy's desire to become a better musician might very well be furthered by her believing that she is a more talented musician than her evidence indicates, as it might spur her to practice more often and learn harder pieces. Basil's desire to recover from a serious illness might very well be furthered by his believing that his chances of recovery are better than his evidence indicates, as it might motivate him to take actions that will ultimately lead to his recovering despite the odds. People often seek to instill such beliefs in their loved ones for reasons just like these. Yet regardless of the practical advantages one might reap from having such beliefs, they are nevertheless liable to criticism as such for having inadequate evidential support. Beliefs, it seems, are subject to evaluation in regard to their responsiveness to truth (or the appearance of truth) independently of whether having beliefs that are so responsive serves the practical interests of the agent who holds them. If beliefs did not aim at truth constitutively, but merely inherited their alethic aim as a byproduct of the practical aims of agents, it

²¹ Such a view need not thereby claim that we are believers*. Beliefs* were supposed to be subject to the norms of practical rationality directly, much as our intentions can be said to be subject to the norms of practical rationality directly. By contrast, a proponent of the view I am now considering might claim that beliefs are subject to the norms of practical rationality, but only indirectly. On this view, our beliefs are normatively beholden to our success as agents (and, so, to the satisfaction of our desires indirectly), rather than directly to the satisfaction of our desires. On such a view, it would be rational for me to believe that I am more talented than I have evidence to believe only if having this belief was likely to result in my being more successful as an agent. For a believer*, such a belief* would be rational if it were one that she desired, all things considered. Funkhouser, however, seems to hold that we are believers* who happen to have a keen interest in having true beliefs. This strikes me as highly implausible.

would be difficult to avoid the conclusion that the intuition that beliefs are so evaluable is misguided.

Consider language and etiquette. Language use and public behavior are subject to evaluation as correct or incorrect independently of the evaluated agent's practical interests. Clearly, I can be charged with speaking ungrammatically, or with behaving rudely, even when abiding by the rules of grammar or etiquette runs counter to my best interest. In the case of language use and public behavior, the evaluation is conducted relative to a convention, which sets the relevant standard of correctness. Similarly, we evaluate beliefs relative to truth regardless of whether having true beliefs serves the interests of the believer. In the case of belief, however, there is no convention to underwrite our evaluation. For belief, the standard of correctness that underwrites our evaluation seems to be set internally, by the very nature of belief. Our very evaluative practices, then, seem to commit us to the view that the relevant standard of correctness for belief evaluation—namely, truth—is set internally, by the content of belief.

That belief ought to respond to truth (or the appearance of truth) independently of practical considerations is a deep-seated, pre-theoretic intuition. It manifests itself in the indignation children typically feel when, having asked for reasons to believe that something is the case, they are given a threatening answer along the lines of 'Because I tell you to, or else...' Intuitively, nothing but an indication that a belief is true counts as a proper reason for accepting that belief. The natural, pre-theoretic view, it seems, is that beliefs, as such, are beholden to truth, not to the practical interests of agents. It is a strong mark in favor of the view that beliefs constitutively aim at truth that it can both explain why beliefs are responsive to truth (or the appearance of truth) and account for the intuition that they ought to be so responsive irrespective of any non-alethic considerations.

This is not to say that practical considerations have no import in what one believes or ought to believe. Practical considerations typically play a significant role in shaping one's doxastic interests. David Owens (2000: Ch. 2) argues convincingly that pragmatic considerations²² play a decisive role in determining whether a given body of evidence is sufficient for rational belief:

[P]urely evidential considerations *under determine* what we ought to believe until they receive pragmatic supplementation. The evidentialist may be right to insist that whether I believe p rather than not- p is something that should be fixed purely by the balance of evidence for and against p . But where and when I form a view as to whether p is true will be determined by my sense of how important the issue is, what the consequences of having a certain belief on the matter would be and how much of my limited cognitive resources I ought to devote to it before reaching a conclusion. (*Ibid*: 26-27)

Practical considerations can also give rise to doxastic duties. It is plausible that agents have the doxastic duty to ascertain those facts that are accessible to them and that are relevant to the moral status of the actions they are contemplating to undertake. If my bank account suddenly registered a far larger balance than I had reason to anticipate, it seems that I would have the duty to investigate whether the unexpected bounty was rightfully mine before I spent it on an expensive vacation. Agents might have the doxastic duty to identify the responsibilities they bear, e.g., toward those in their care. Ignorance of the responsibilities parents bear toward their children provides no excuse for neglecting one's children—parents ought to know their responsibilities as parents.

It would seem that positive doxastic duties, i.e., duties to believe rather than duties to refrain from believing, are typically grounded in the practical interests and practical duties of agents. Independently of such practical interests and practical duties, believers might merely have the negative doxastic duty to refrain from holding beliefs of whose truth they have insuffi-

²² Owens (2000: 24) distinguishes between pragmatic considerations (“non-evidential considerations that are relevant to belief”) and practical considerations (considerations of the sort invoked by practical norms). I call both types of considerations ‘practical,’ as it is not entirely clear to me that there is a sharp line to be drawn here.

cient indication. The corresponding positive duty, i.e., the duty to believe everything of whose truth one has sufficient indication, seems extravagant and difficult, if not impossible, to meet by finite beings, as it requires them to have far more beliefs than they can be expected to hold. Positive doxastic duties are typically hypothetical, in the sense that they are conditional on the interests of agents—if Amy is a physics student, she ought to form beliefs about the theory of relativity; if Basil isn't, and has no interest in the topic, he need not. When positive doxastic duties are categorical, in the sense that they are not conditional on the interests of agents, they seem to be grounded in categorical duties to act in certain ways. So, parents ought to know their duties as parents in virtue of having the categorical duty to abide by these duties in their actions. It is not implausible that agents have categorical doxastic duties to hold certain moral beliefs. If agents have such categorical positive doxastic duties, however, it is plausible that they have them in virtue of having the categorical duty to express these beliefs through their actions.

Positive doxastic duties, whether hypothetical or categorical, might conflict with, and sometimes even override, the negative duty all believers have to refrain from holding beliefs of whose truth they have insufficient indication. In the case mentioned above, Basil might have the prudential duty to believe that he has decent chances of recovering from his illness even though he has every indication that his medical case is hopeless. Still, as stressed above, Basil remains subject to his negative doxastic duty as a *prima facie* duty, and his belief remains liable to criticism as such for having inadequate evidential support. That this is so is inexplicable, I claim, on the view that beliefs inherit their alethic aim as a byproduct of the practical aims of agents. That positive doxastic duties typically have a practical grounding, however, might mislead one to endorse this view.

The Transparency of Belief

That the question, ‘Do I believe that p ?’ normally is, and ought to be, transparent to the question, ‘Is it the case that p ?’ in the sense that the former normally is, and ought to be, answered by way of answering the latter, is further evidence that belief is necessarily governed by the norm of truth. In *Authority and Estrangement*, Richard Moran argues that the transparency of first-person, present-tense attitude reports, noticed by Roy Edgley and Gareth Evans, is both a presupposition of, and a normative requirement on, rational agency. Given the foregoing discussion, this is just what should be expected in the case of agents who are capable of reflecting on their own attitudes. Moran writes:

[P]art of what it is to be a rational agent is to be able to subject one’s attitudes to review in a way that makes a difference to what one’s attitude is. One is an agent with respect to one’s attitudes insofar as one orients oneself toward the question of one’s beliefs by reflecting on what’s true, or orients oneself toward the question of one’s desires by reflecting on what’s worthwhile or diverting or satisfying. (Moran 2001: 64)

And further:

[A]s I conceive of myself as a rational agent, my awareness of my belief is awareness of my commitment to its truth, a commitment to something that transcends any description of my psychological state. And the expression of this commitment lies in the fact that my reports on my belief are obliged to conform to the condition of transparency: that I can report on my *belief* about X by considering (nothing but) X itself. (*Ibid*: 84)

Moran distinguishes two stances an agent may take toward her own doxastic attitudes, namely, the theoretical stance and the deliberative stance. When asking, ‘Do I believe that p ?’ as a theoretical question, the agent’s stance toward her own belief is just like that of any outside observer. Proper evidence for her inquiry is evidence that she believes that p and may include her own behavior, and the answer, ‘Yes, but it is not the case that p ’ does not involve her in any outright incoherence, since the fact that she believes that p and the fact that p are distinct. When asking, ‘Do I believe that p ?’ as a deliberative question, by contrast, the agent’s stance toward

her own belief is unlike that of any outside observer. The deliberative question is equivalent to the question, ‘Shall I believe that p ?’ and does not aim merely to yield an accurate attitude report; it aims to yield belief—i.e., an attitude that conforms to what Moran calls ‘the condition of transparency’ in the quote above.²³ In the deliberative stance, the agent’s inquiry is not properly focused on evidence that she believes that p . It is properly focused on evidence that p , that is, on alethic reasons for believing that p .

Moran argues that the capacity to take a deliberative stance toward one’s own judgment-sensitive attitudes is a prerequisite of rational agency. Agents are responsible for their attitudes inasmuch as they have the capacity to avow them, where an avowal is a statement of one’s attitude that conforms to the condition of transparency (*Ibid*: 101). Avowability also sets a normative requirement on rational agency. Inasmuch as agents have the capacity to avow their judgment-sensitive attitudes, they are committed to satisfying the condition of transparency in their judgment-sensitive attitudes. In the case of belief, the commitment is to have the belief that p only if one’s report, ‘I believe that p ’ would be a proper answer to the deliberative question, ‘Do I believe that p ?’

The thesis that belief is necessarily governed by the norm of truth explains why the deliberative question, ‘Do I believe that p ?’ normally is, and ought to be, answered by way of answering the question, ‘Is it the case that p ?’ Before I can present the details of this explanation, I need formulate an account of the manner in which the norm of truth governs belief. Accordingly, I will postpone this issue until the end of Chapter 4, when I will address it in the section entitled “The Deliberative Stance”, and leave matters at the intuitive level for now. I should note also that Moran’s general account of rationality offers a deeper explanation of the thesis for

²³ Also: “A statement of one’s belief about X is said to obey the Transparency Condition when the statement is made by consideration of the facts about X itself, and not by either an ‘inward glance’ or by observation of one’s own behavior” (*Ibid*: 101).

which I have been arguing. The broader order of explanation goes as follows: The capacity to take a deliberative stance toward one's own judgment-sensitive attitudes is a prerequisite of rational agency. The belief that p is sensitive to the judgment that p . That the belief that p is sensitive to the judgment that p explains why belief is necessarily governed by the norm of truth. That belief is necessarily governed by the norm of truth explains why the deliberative question, 'Do I believe that p ?' is transparent to the question, 'Is it the case that p ?'

Moore's paradox: Moran's account of the transparency of belief in first-person deliberation explains why assertions of the forms, 'I believe that p , but p is false' and ' p , but I don't believe that p ,' strike us as paradoxical. Their third-person counterparts, i.e., 'She believes that p , but p is false' and ' p , but she doesn't believe that p ,' are seen as entirely unproblematic. Moore's Paradox presents us with a problem, since each assertion schema, both in the first-person and the third-person formulations, consists of two conjuncts that have distinct subject matters and that can concomitantly be true. Whatever the problem is with the first-person schemas, it isn't formal contradiction (presumably I, too, can have false beliefs and fail to believe some truths).

Moran notes that we view as paradoxical not only assertions, but also thoughts that conform to the first-person schemas. For this reason, he rejects the prevalent pragmatic explanation of Moore's Paradox, according to which the problem in the first-person cases results from a self-defeat of the purpose of assertion (*Ibid*: 84). Instead, Moran proposes that the reason why assertions (and thoughts) of the forms, 'I believe that p , but p is false' and ' p , but I don't believe that p ,' strike us as paradoxical is that they express a normative contradiction within the deliberative stance.²⁴ Inasmuch as I am capable of reflecting on my beliefs, I have a commitment to

²⁴ Although Moran does not use the expression, 'normative contradiction,' it seems to capture his line of thought.

satisfying the condition of transparency with regard to my beliefs. That is, I am committed to having the belief that p only if my report, 'I believe that p ' would be a proper answer to the deliberative question, 'Do I believe that p ?' But I could not properly answer the deliberative question with 'I believe that p ' while also holding that p is false, since my holding that p is false would commit me to answering the deliberative question with, 'I don't believe that p .' Likewise, I could not properly answer the deliberative question with 'I don't believe that p ' while holding that p , since my holding that p would commit me to answering the question with, 'I believe that p .'

Moran's account of the properly first-person stance as the stance in which one can avow one's judgment-sensitive attitudes, as well as his attractive answer to Moore's Paradox, tie together the strands I have been pursuing above. That beliefs typically are, and always ought rationally be—at least *prima facie*—responsive to (the appearance of) truth clearly has a close relation to the condition of transparency Moran discusses. To require that the beliefs of an agent capable of reflecting on her beliefs be avowable is to require that they be responsive only to what is properly taken by her lights as indicative of their truth, that is, to alethic reasons. That beliefs are subject to, and typically comply with, such a requirement is highly plausible. This requirement, however, does not make sense if belief is not necessarily governed by the norm of truth. If belief is not necessarily governed by the norm of truth, why shouldn't the agent who seeks to answer the deliberative question, 'Do I believe that p ?' let consideration of, say, the use or appeal of believing that p guide her inquiry? After all, as I noted above, believing that p despite overwhelming evidence that not- p can at times be of far greater value to the agent than refraining from believing that p or believing that not- p .

Belief and Believer Evaluability

I argued earlier that the claim that belief aims at truth constitutively should be understood (in part) as casting belief as subject to evaluation relative to truth along the dimensions of success and performance independently of any other aims and desires the believer might have, and as casting the believer as subject to corresponding evaluations and attendant praise or blame. Although I doubt that anyone would seriously dispute these claims, I would feel remiss if I did not revisit them at this stage.

Clearly, truth constitutes the dimension of success for belief. It is quite clear also that belief can independently be evaluated along a dimension of performance, which has received various characterizations and various names—epistemic rationality, entitlement, warrant, justification.²⁵ I hope the foregoing discussion has made it equally clear that belief can be evaluated along these two dimensions independently of any other aim and desire the believer might have. Finally, the claim that believers are responsible for their beliefs is fairly uncontroversial. It might be that evaluations of believers and their beliefs along the dimension of performance might occasionally yield different verdicts. I considered earlier the case in which Basil has a prudential duty, by which he abides, to believe that he has decent chances of recovering from an illness even though he has every indication that his medical case is hopeless. I claimed, plausibly I think, that Basil remains subject to his doxastic duty to refrain from holding beliefs of whose truth he has insufficient indication as a *prima facie* duty, and that his belief remains liable to criticism as such for having inadequate evidential support. In this case, Basil's belief is false and epistemically irrational. However, it is not quite clear how Basil should be evaluated for his epistemic performance. Basil certainly does not seem blameworthy for having this belief. Given that epistemic considerations are properly overridden in this case, Basil seems irreproachable in

²⁵ These might be distinct, though related, dimensions of evaluation.

terms of overall rationality. And since epistemic considerations are properly overridden in this case, Basil does not clearly seem to deserve reproach in terms of epistemic rationality either. This, however, is consistent with the thesis that belief aims at truth constitutively, which claims only that both believer and belief are liable to evaluation along the dimensions of success and performance. It does not claim that evaluations of believer and belief must yield coinciding verdicts in every case. This concludes my *prima facie* case for the thesis that belief aims at truth constitutively.

Theoretical Refinements

Should the Thesis be Understood Literally?

The online *Oxford English Dictionary* lists the following senses of the verb, ‘to aim’ (obsolete and clearly irrelevant senses omitted):

5. To calculate one’s course with a view to arriving (at a point); to direct one’s course, to make it one’s object to attain.
Hence *fig.* To have it as an object, to endeavour earnestly.
- b. *intr.* Also with infinitive: to intend; to attempt.
6. a. *intr.* To calculate or estimate the direction of anything about to be launched (at an object); to deliver a blow, or discharge a missile (at anything) with design or endeavour to strike.
Hence *fig.* To try to hit, gain, or bring into one’s power; to have designs upon, to seek to obtain.
7. *trans.* To direct (a missile, or blow); especially, to direct it with the eye before its discharge; to point or level a gun, etc. (at).
Hence *fig.* To direct any act or proceeding against. Also *transf.*
8. *absol.* In both the preceding senses: To take aim; to form designs.

In non-figurative uses, the active voice of the verb, ‘to aim,’ attributes an intention to the referent of its grammatical subject. In figurative uses such as “Our translation, aiming at the sense, rendereth it, etc.” (*OED* quotation, listed under 6a), the referent of the grammatical subject of ‘to aim’ is represented as governed by an intention. In such cases, the intention is not attributed directly to referent of the grammatical subject of ‘to aim,’ but indirectly to the agent whose

intention governs the referent of this grammatical subject.²⁶ Aiming, then, seems to be subject to the following condition:

Necessarily, for any x , if x is aiming, then x is either an agent, or purposive, where the relevant sense of ‘purposive’ entails that (i) x is an event, and (ii) x is intentional under some description.²⁷

Since beliefs are obviously not agents, the thesis that belief aims at truth constitutively entails that belief is purposive. Yet given that belief is not an event but a state, and given that belief does not seem to be intentional, the thesis seems thereby to commit two category mistakes. Let us consider these two issues in reverse order.

Concern that belief is not intentional

The thesis that belief aims at truth constitutively entails that belief is necessarily intentional. That is, it entails that, as a matter of conceptual necessity, every belief is generated (and presumably maintained) by an alethic intention. Minimally, this intention would be a doxastic analogue of what John Searle calls ‘intentions-in-action.’²⁸ Using Searle’s notation (Searle 2001: 44-45), the content of this alleged alethic intention for the belief that p would be represented as follows:

²⁶ *Mutatis mutandis* for plural uses.

²⁷ In “Our translation, aiming at the sense...” (quoted above), ‘translation’ can be interpreted as referring either to an action or to a document. On the latter interpretation, ‘aiming’ seems to be used in a transferred sense, where the aim properly ascribable to the action it governs is transferred to the document that this action produces.

²⁸ Searle distinguishes intentions-in-action from prior intentions. While prior intentions have a propositional content, the content of an intention-in-action directly involves the event-type that the intention is to cause. Searle offers the following example to illustrate the distinction:

[S]uppose that in a meeting I want to vote for the motion by raising my right arm. I thus form the prior intention that I raise my arm. The intentional content of the prior intention to raise my arm can be represented as follows:

p.i. (that I raise my arm and that this p.i. causes that I raise my arm). ...

The prior intention has to be distinguished from what I call the *intention-in-action*. The intention-in-action is the intention I have while I am actually performing an action. Thus in this case, when the moment to vote comes, and the chair says “All those in favor raise your arm,” I will act on my prior intention and thus have an intention in action whose conditions of satisfaction are that that very intention-in-action should cause the bodily movements of my arm going up. We can represent that as follows:

i.a. (my arm goes up and this a.i. causes that my arm goes up). ...

Of course, not all actions are premeditated. Many of the things I do, I do quite spontaneously. In such a case I have an intention-in-action but no prior intention (*Ibid*; italics in the original).

(1) a.i. (I believe that p only if p and that this a.i. causes that I believe that p only if p)

This theoretical commitment flies in the faces of introspection and of a fairly wide philosophical consensus that believing is not voluntary.²⁹

Introspection offers no evidence that there exist any such alethic intentions—in fact the immediate etiology of my beliefs does not typically seem to include any intentions to believe at all. I may at times have an intention to form a belief concerning a given subject matter—say, whether p . Note, however, that my intention in this case would have a content of a different form than (1) above, namely,

(2) p.i. (that I form a belief as to whether p and that this p.i. causes that I form a belief as to whether p)

My intention to form a belief as to whether p might cause me to engage in a variety of activities, such as gathering evidence, reflecting on it, etc. Should I eventually form a belief as to whether p , say, the belief that p , my intention to form a belief as to whether p and resultant activities would belong to the distal etiology of my belief that p . Yet my belief that p itself does not seem to be caused by an additional intention to believe that p only if p . Once I am faced with the relevant evidence, my belief that p seems simply to occur as a response to overwhelming evidence that p (or to a preponderance of evidence that p)—no additional intention seems involved, and none seems needed.

Williams argues that, as a matter of conceptual necessity, belief cannot be voluntary. Interestingly, Williams' main reason for endorsing this position is that belief aims at truth. He writes:

One reason [I cannot acquire a belief at will] is connected with the characteristic of beliefs that they aim at truth. If I could acquire a belief at will, I could acquire it whether it was true or not ... With regard to no belief could I know—or, if this is to be done in full con-

²⁹ This view is not unchallenged; see e.g., Funkhouser 2003, quoted above.

sciousness, even suspect—that I had acquired it at will. But if I can acquire beliefs at will, I must know I am able to do this; and could I know that I was capable of this feat, if with regard to every feat of this kind I had performed I necessarily had to believe that it had not taken place? (Williams 1973: 148)

The relevant sense of ‘voluntary,’ implicit in Williams’ argument and explicit in Jonathan Bennett’s critique of it, is ‘immediately responsive to practical reasons’ (Bennett 1990: 90).

While finding doxastic involuntarism intuitively compelling, Bennett produces a counterexample to William’s claim that one could not know that she has the capacity to acquire beliefs at will and regard the states she acquired by exercising this capacity as beliefs. So long as she does not know of any particular state that she acquired it by exercising this capacity, Bennett argues, it is perfectly possible for her to know that she has this capacity and to regard the states she acquired by exercising this capacity as beliefs. Although Bennett had originally set out to establish the truth of doxastic involuntarism in “Why Is Belief Involuntary?” the essay is explicitly offered as an instructive chronicle of his failure.

Like Williams, Pamela Hieronymi believes that doxastic involuntarism is a consequence of the constitutive alethic aim of belief and offers a convincing argument for Williams’ conclusion. Hieronymi holds that belief is partly “commitment-constituted”; that is, she holds that “you believe p just in case you take p to be true in a way that leaves you answerable to certain questions and criticisms—namely, those that come with believing p ” (Hieronymi 2006: 50). She invokes a distinction between “content-related” reasons for beliefs (i.e., reasons that count in favor of believing that p by showing that p is true) and “attitude-related” reasons for belief (i.e., reasons that count in favor of believing that p by showing that the belief that p is a good attitude to have). Since content-related reasons for belief support the commitment constitutive of belief, she calls them “constitutive” reasons for belief. Any non-constitutive reasons for belief she calls

“extrinsic.” On this nomenclature, the question that interested Williams and Bennett is, ‘Can one believe immediately (i.e., without acting on oneself) for extrinsic reasons?’

Hieronymi distinguishes two types of control we might exercise over our beliefs, namely, “managerial control” and “evaluative control.” Managerial control over belief is analogous to the control we have over ordinary objects as we manipulate them so that they accord with our purposes. Evaluative control over belief, by contrast, is exercised by evaluating what is true.

She notes that, unlike managerial control, evaluative control is immediate:

Even in cases of basic action, such as raising your right hand, there remains a conceptual difference between deciding to act (intending) and acting ... In contrast, there are no possible cases in which you answer positively the question whether p but are prevented from believing p . In answering the question positively, one has already, therein, believed. The immediacy of evaluative control is thus not temporal or causal but rather a consequence of the constitutive relation between the commitment to p as true and the belief. Since we believe p insofar as we are committed to p as true, when we take p to be true in the relevant way, our beliefs therein change. (*Ibid*: 53-54)

Hieronymi denies on the following grounds that one can believe that p for extrinsic reasons by an exercise of evaluative control:

Since one evaluatively controls a belief that p by answering the question of whether p , and since extrinsic reasons are those reasons that are *not* taken to bear on the question of whether p , one cannot evaluatively control a belief that p in response to extrinsic reasons. (*Ibid*: 59-60)

Since any other type of control over one’s belief must be mediated by an intention, Hieronymi concludes that one cannot believe immediately for extrinsic reasons.

Hieronymi holds that, like belief, intention is partly commitment-constituted and, accordingly, finds that belief and intention are analogous in the following respects:

[N]otice that we can draw the same two distinctions in reasons for intending that we have drawn in reasons for believing. We can distinguish, first, between content-related reasons (which bear on the question of whether to ϕ) and attitude-related reasons (which bear on the question of whether it would be good, useful, desirable, appropriate, or important in some way to intend to ϕ). As with belief, the content-related reasons for an intention support the commitment that constitutes it. We can call these the constitutive reasons for in-

tending. By taking them to settle the question of whether to ϕ , you will, therein, intend to ϕ . One evaluatively controls an intention to ϕ by answering for oneself (however implicitly) the question of whether to ϕ . (*Ibid*: 56-57)

Hieronimi's plausible characterization of belief and intention suggests an argument showing that no alethic intention of the sort represented in (1) above could generate the belief that p in the absence the belief that p . Accordingly, any account that appeals to such alethic intentions to explain the alethic character of belief is viciously circular. On most accounts, a constitutive reason capable of settling the evaluative question of whether to ϕ , i.e., a total constitutive reason³⁰ to ϕ , consists of a desire to ψ (where 'desire' is understood broadly as a pro-attitude toward ψ ing) coupled with the belief that ϕ ing is the best means of satisfying ψ . For the intention-in-action alleged to generate/sustain the belief that p , i.e., the intention to believe that p only if p , a total constitutive reason to believe that p only if p would consist of some desire coupled with the belief that believing that p only if p is the best means of satisfying this desire. Now, what could be the content of this desire? There is only one candidate that accords with the spirit of the thesis that belief aims at truth, namely, to have a true belief as to whether p . Accordingly, the means-end belief involved is the belief that believing that p only if p is the best means of satisfying the desire to have a true belief as to whether p .

Although this result might seem innocuous at first blush, the illusion of harmlessness is dispelled once one realizes that the alleged intention-in-action to believe that p only if p is motivationally inert on its own. It can move one to believe that p when p only once the question as to whether p is settled. But to settle the question as to whether p , it seems, is already to believe that p (as Hieronymy plausibly stresses in the first indented quote above). Put another way, the problem is that the desire to have a true belief as to whether p , coupled with the means-

³⁰ I am borrowing the term, 'total reason,' from Searle (2001: 115ff). 'Reason' is used in an internal sense here, as that which not only rationalizes, but also causally explains, intention and action.

end belief that believing that p only if p is the best means of satisfying this desire, cannot move one to believe that p when p unless it is conjoined with the belief that p (assuming the standard account of intention/action explanation). One might answer that motivational efficacy for the intention to believe that p only if p does not require the belief that p , but merely, e.g., a perception as of p . This response does seem available to the proponent of alethic intentions to believe. However, it is liable to two objections. First, since the alleged intention is to believe that p only if p , not to believe that p only if it appears to one as if p , motivational efficacy for the alleged intention requires that, e.g., the perception as of p be taken as veridical. More generally, it requires that the question as to whether p be settled. That is, it requires belief that p (though the belief that p need not be reflected or even conscious). If believing that p requires the intention-in-action to believe that p only if p , then, there arises a vicious psycho-causal circle.^{31,32} Accordingly, believing that p can require no such intentions. Yet the first objection might be thought to beg the question against the proponent of alethic intentions to believe—why shouldn't the intention to believe that p only if p be motivated by the desire to have a true belief as to whether p together with a perception as of p ? The second objection answers this question: because this requires rejecting the standard account of intention/action explanation, which is highly plausible on independent grounds.

³¹ The view that belief is necessarily caused/sustained by an alethic intention also threatens to generate an infinite regress on the standard account of intention/action explanation, since the means-end belief underwriting the alethic intention would require a further alethic intention.

³² Total constitutive reasons for some intentions-in-action might not consist of a desire together with a means-end belief, but of a desire, which might be non-propositional, together with a non-propositional recognition of how best to satisfy this desire in one's circumstances (where the recognition is made possible by a non-propositional know-how that the agent possesses). Such seems to be the case when we engage in behavior that is clearly intentional, such as a riding bicycle, while not possessing the conceptual apparatus that would be required to formulate the relevant means-end beliefs even implicitly. Conceiving alethic intentions-in-action along this model is unlikely to provide an avenue of escape to the alethic-intentions theorist. Indeed, given the theoretical work such alethic intentions are called to play, the only desire that could plausibly be involved in a total constitutive reason for the intention-in-action to believe that p only if p is the desire to have a true belief as to whether p . Even if the formation of the intention-in-action to believe that p only if p involved an exercise of a belief-forming know-how, the resulting intention-in-action would be powerless to motivate the belief that p when p unless it was accompanied by a recognition that p . But a recognition that p , it seems, is already a belief that p . I should note, however, that no regress threatens on this model of alethic intentions-in-actions.

For these reasons, the thesis that belief aims at truth should not be understood literally as entailing that belief is necessarily purposive. Instead, I propose to interpret the thesis as claiming that belief is necessarily governed by truth as a regulating principle (i.e., necessarily governed by the norm of truth). Further, borrowing Hieronymi's expression, I propose to interpret the thesis as claiming that belief is commitment-constituted relative to the truth of its content just as action is commitment-constituted relative to the intention that governs it. That is, as I argued earlier, to describe an action in terms of the intention under which it was undertaken *ipso facto* casts this action and the agent who undertook it as evaluable along the dimensions of success and performance relative to that intention, at least *prima facie*.³³ Likewise, to describe something as a belief that *p ipso facto* casts this belief and the person who has this belief as evaluable along the dimensions of success and performance relative to *p*. However, although the structures of belief and action are analogous in these respects, they are radically different in the following respect: Intention has a world-to-mind direction of fit. Accordingly, the commitment constituted by intention for action is ultimately grounded in the mind (i.e., in the agent's desires, together with her beliefs as to how to satisfy them). Hence, constitutive reasons for intention (and, so, for action) are desires, coupled with corresponding means-end beliefs. By contrast, the principle governing belief has a mind-to-world direction of fit. As such, then, the principle governing belief is not an intention, and the commitment it constitutes for belief is not grounded in the mind, but in the world. Hence, constitutive reasons for belief do not involve desires, but are reasons for thinking that the content of the belief is true.

In *The Possibility of Practical Reason*, J. David Velleman proposes that the claim that belief aims at truth should be explicated as follows:

³³ This, of course, does not entail that an action that carries out its intention successfully and well is thereby rational.

Belief is the attitude of accepting a proposition with the aim of thereby accepting a truth ... An acceptance has the aim of being the acceptance of a truth when it is regulated, either by the subject's intentions or by some other mechanisms in ways designed to ensure that it is true. (Velleman 2000: 252-54)

Velleman's thesis strikes me as unhelpful at best and implausible at worst. I have argued already that any attempt to explicate the alethic characteristics of belief in terms of an alethic intention leads to vicious circularity. Further, I am suspicious of Velleman's use of the notions of accepting a proposition and of design, as they are no clearer than the notion they are meant to elucidate, namely, the notion that belief (or believing) aims at truth constitutively, and as they seem to have conceptual commitments that belief (or believing) doesn't have.

Accepting a proposition seems to require higher cognitive powers than believing. That infants and some animals are capable of forming rudimentary beliefs is *prima facie* plausible. But it is far less plausible that they are capable of accepting propositions. Unlike believing, it seems, accepting a proposition is subject to voluntary control. For this reason, unlike believing, accepting a proposition seems to require the capacity to reflect on the content of one's prospective acceptance. Of course, Velleman might be using 'acceptance' as a term of art that is devoid of the commitments the term has in common parlance. If so, until a clear definition is offered, the term is rather unhelpful. If Velleman has the common-sense notion of acceptance in mind, his account is implausible.

The notion of design, too, is troublesome for the following reason (although the disagreement, here, might not be especially deep). Suppose I unreflectively form an ordinary perceptual belief in the ordinary way. Now suppose there is a world that is now an exact duplicate of this world but that spontaneously sprang into being a few minutes ago. In this case, the mechanisms that regulate the acquisition of my duplicate's perceptual belief were clearly not designed to ensure that it is true. Yet my duplicate's state is nevertheless plausibly a state of belief. Of

course, since Velleman merely offers the disjunction as a sufficient condition for an acceptance to have the aim of being an acceptance of a truth, this case does not provide a counterexample to his account. Yet this case puts pressure on Velleman's account by casting doubt on the claim that, in the actual world, my acceptance has the aim of being an acceptance of a truth in virtue of being regulated by mechanisms *designed* to ensure that my acceptance is an acceptance of a truth.

Unlike Velleman, then, I do not identify believing with accepting a proposition, or the attitude of belief with the attitude of having accepted a proposition, with the aim of thereby accepting a truth. Instead, I propose to understand believing as a type of mental event individuated by its function in generating/sustaining belief, and belief as a type of mental state individuated by its function in sanctioning and preserving information as such in conceptualized form³⁴ (I will call these functions the 'proper cognitive functions' of believing and belief, respectively).^{35, 36} Further, I propose that believing is governed by the norm of truth in virtue of its proper function and in virtue of the proper function of belief. These proposals will be fleshed out in the next two chapters. Finally, I confess to having no idea how to answer the question, 'What are the necessary and sufficient conditions for an event or a state to have such functions?' But it is not entirely clear to me that much would be gained for my purposes here if I were able to offer such necessary and sufficient conditions.

Concern that belief is not an event, but a state

I stated above that the thesis that belief aims at truth constitutively should be understood as claiming that that belief is necessarily governed by truth as a regulating principle. Yet it clearly

³⁴ What is meant by 'conceptualized form' will be clarified in the following chapter.

³⁵ Believings/beliefs with particular contents, as types and as tokens, involve further individuating conditions.

³⁶ A dispositional belief would be one that the believer would be disposed to believe immediately, were the issue to arise.

makes very little sense to say that a state embodies a principle that governs it, since a state, as such, is not amenable to change. The specification of belief as a state, then, seems incompatible with the claim that belief aims at truth. How should we respond to this problem?

In my view, we should say that what aims at truth constitutively (in the non-literal sense proposed above) is not belief itself, but believing, i.e., the event of forming or maintaining a belief. On this interpretation, belief itself cannot, as a matter of conceptual necessity, aim at anything or be governed by anything. When we say that belief aims at truth constitutively, we speak loosely, as we do when we say, pointing at a dart that just hit the bull's eye, 'Nice shot!' The dart protruding from the board is the outcome of a nice shot, not a nice shot. Likewise, belief, as such, is not governed by the norm of truth; it is the outcome of an event governed by the norm of truth.

The thesis requires a further clarification. I argued earlier that the evaluability of belief along the dimensions of performance and success is grounded in truth as the commitment-constituting principle governing belief. Now that I have concluded that this principle governs not belief, but believing, I have committed myself to holding that evaluability along the dimensions of both performance and success properly belongs only to believing. This might seem to contradict the rather obvious truism that it is belief as a representational state, and not believing as an event, that is evaluable as true or false. The contradiction, however, is merely apparent. An evaluation of a teleological event along the dimension of success evaluates the outcome of this event. By contrast, an evaluation of a teleological event along the dimension of performance evaluates the event itself. Accordingly, on the view I have been proposing, evaluation as true/false properly belongs to belief as a representational state, while evaluation as doxasti-

cally/epistemically rational/irrational, justified/unjustified, warranted/unwarranted, etc., properly belongs to believing as an event. This is a happy result, as it is independently plausible.

Believing and Guessing

In “Does Belief Have an Aim?” Owens argues, correctly I think, that guessing (understood as a mental act) aims at truth constitutively, and challenges the proponents of the truth-aim hypothesis of belief to explain why the rationality of guessing differs from the rationality of believing. Since the explanatory burden Owens proffers belongs properly to the proponents of a literal interpretation of the hypothesis, who posit that belief is purposive and purport to explain the rationality of belief in terms of a constitutive alethic intention, his challenge does not affect my position. Nevertheless, a brief discussion of the differences between guessing and believing should be useful at this stage.

Owens thinks that although truth sets the standard of correctness for guessing, the rationality of guessing does not differ from the rationality of any other type of act. While truth is the constitutive aim of guessing, it operates in conjunction with the agent’s other aims and purposes in rationalizing her guesses. So, on Owens’ view,

guessing that p is reasonable when aiming at the truth by means of a guess that p would maximise expected utility. That requirement integrates the purpose constitutive of guessing with the subject’s wider purposes in a familiar and well understood fashion.
(Owens 2003: 292)

Although an agent cannot be said to guess that p unless she intends to guess that p only if p , a reason to guess that p can be evidence that p , but it need not: the practical consideration that, e.g., time is running out can also be a reason to guess that p , provided that the agent’s evidence favors p over not- p .

The account I have been outlining is not merely compatible with the differences between guessing and believing Owens stresses; it predicts and explains them. The norm of truth governs

both believing and guessing. However, unlike believing, guessing is governed by an intention, which must include an alethic component. Because guessing is governed by an intention, i.e., the intention to guess, it is evaluable along the dimensions of success and performance relative to this intention (at least *prima facie*). A mental act intended as a guess might fail in virtue of not being a guess at all. If it succeeds in being a guess, it is then also evaluable as a true or false guess. Further, a mental act intended as a guess can be evaluated along the dimension of performance relative to the agent's practical interests (in the manner in which actions are typically evaluated along this dimension) and relative to the agent's alethic interest (in the manner in which beliefs are typically evaluated along this dimension). That the norm of truth governs guessing, but not believing, via an intention that must include an alethic component, explains why reasons for guessing, unlike reasons for believing, can include both practical and alethic considerations.

Reflection on another distinctive feature of guessing should prove illuminating for my account of the manner in which the norm of truth governs belief. Consider the following situation: I see that there is now a glass on the table to the right of my computer. Although guessing constitutively aims at truth, I cannot guess that there is now a glass on the table to the right of my computer, regardless of how useful such a guess might be to me—I could not guess it if my life depended on it. If I am to form an attitude about this proposition, I must believe it. Why is this? An answer to this question, I think, requires an account of how guessing and believing fit in the wider context of our cognitive lives. Accordingly, I will now offer a very brief sketch of such an account. Since defending such an account would exceed the scope of this project by far, I will not argue for my claims, but I hope that they have enough *prima facie* plausibility to be worth presenting nevertheless.

Believing and guessing are cognitive events. That is, they are mental events whose function is to generate/sustain cognitive states. Cognitive states are mental states whose function is to carry information. Cognitive events are individuated in part in terms of the cognitive states they generate/sustain. Cognitive states are individuated in part in terms of the manner in which they carry their informational content and in terms of their informational content. Our common-sense psychological concepts seem to distinguish different types of cognitive state in terms the manner in which states of these types carry information. So, for example, we seem to distinguish belief from perception as types of cognitive state inasmuch as belief carries information in fully conceptualized form, and carries it *as* information, while perception does not.³⁷ All cognitive events are governed by the norm of truth, inasmuch as their function is to generate/sustain states whose function is to carry information ('information' is understood factively here). The manner in which the norm of truth governs cognitive events of a given type, however, depends on the manner in which the states they generate/sustain carry their information. Further, the manner in which the norm of truth governs a cognitive event of a given type depends, in addition, on the cognitive content of the cognitive state it generates. Even imagining seems governed by the norm of truth. Truth does not set the standard of correctness for the products of my imaginings—after all, the informational content of the states my imaginings generate is not posited *as* information. Nevertheless, my imaginings seem subject, e.g., to a requirement of conceptual consistency—try as I might, I cannot imagine that a triangle has four sides—and my subjection to this requirement seems to be dictated by the informational content of the states my imaginings produce.

³⁷ I suspect that these concepts express natural kinds.

Now, cognitive events are not isolated events. My believings and guessings are events in a wider cognitive process, which encompasses my cognitive life. The function of this process is to provide me with information. Accordingly, this process, too, is governed by the norm of truth. I speculate that the reason why I cannot guess that p in the face of overwhelming evidence that p is that doing so is proscribed by the norm of truth, to which I am subject as a cognizer. If I form an attitude as to p in the face of overwhelming evidence that p , my attitude must sanction its content as information. This is accomplished by believing, but not by guessing or conjecturing. In more general terms, the norm of truth seems to govern my cognitive life both “prospectively” and “retrospectively.” The norm of truth governs my cognitive life retrospectively inasmuch as it determines the commitments that I do or would undertake by undertaking a given cognitive attitude. So, if I believe that p , it must be the case that p and I must have suitable evidence that p . The norm of truth governs my cognitive life prospectively inasmuch as it determines what cognitive attitude it is appropriate for me to take in response to my evidence (should I have an interest in having a cognitive attitude concerning this evidence). So, if I have overwhelming evidence that p and an interest in having a cognitive attitude concerning this evidence, I cannot merely guess or conjecture that p ; I must believe that p .

The Norm of Truth and the Concept of Belief

The bar against guessing in the face of overwhelming evidence that p seems to be metaphysical rather than conceptual (if such a distinction can legitimately be drawn). That is, the problem isn't that that I cannot conceive of myself as guessing in the face of overwhelming evidence that p . Rather, the problem is that I cannot guess in the face of overwhelming evidence that p . In “How Truth Governs Belief,” Nishi Shah argues that belief is subject to the norm of

truth as a matter of conceptual necessity.³⁸ More precisely, he claims that the norm of truth is properly understood as governing deployments of the concept of belief, rather than belief itself:

What I suggest is that by framing his deliberation as answering to the question whether to believe that p , a disposition to be moved by considerations that he regards as relevant to the truth of p and a disposition blocking considerations that he regards as irrelevant to the truth of p are activated. That is, part of possessing the concept of belief involves being disposed in this way when one applies the concept to frame one's reasoning. But, on this view, when one doesn't exercise the concept of belief, as in cases of non-deliberative belief-formation, there is no guarantee that one's cognitive activity will be regulated by these dispositions. (Shah 2003: 467)

Shah's main motivation for endorsing this view is to avoid a dilemma that he thinks faces any teleological account of the normative authority of truth over belief. According to Shah, any such account is caught on the horns of the following dilemma:

On one horn, the teleologist must allow the disposition that constitutes aiming at truth to be so weak as to allow paradigm cases in which beliefs are caused by such non-evidential processes as wishful-thinking, in which case he cannot capture the exclusive role of evidence in one particular type of belief-forming process, reasoning. On the other horn, in order to account for the exclusive role of evidence in reasoning about what to believe, the teleologist must strengthen the disposition that constitutes aiming at truth so that it excludes the influence of non-truth-regarding considerations from such reasoning. However, by strengthening the truth-aimed disposition, the teleologist cannot accommodate the cases of wishful thinking, in which non-evidential factors clearly exercise influence over belief. This dilemma arises because teleological accounts try to reduce the essentially normative conceptual relation that truth has to belief to a descriptive, causal relation. (*Ibid*: 461)

Although the account of the manner in which the norm of truth governs belief I have been proposing is non-purposive, it is teleological. Accordingly, I need to address Shah's objection. Before I address it, I should note that Shah's own position is implausible for the following two reasons: First, a belief that is improperly responsive to evidence is not immune to epistemic criticism by virtue of not having been conceived as a belief, or by virtue of not having been

³⁸ Shah also argues that the phenomenon of transparency discussed by Moran occurs only within first-person doxastic deliberations involving an application of the concept of belief. This seems mistaken. Transparency seems to be a feature of first-person doxastic deliberation, yet such deliberation need not involve a deployment of the concept of belief. In Chapter 4, however, I will argue that the capacity to take the deliberative stance toward one's beliefs requires the concept of belief. I will also argue that evaluability as doxastically/epistemically rational/irrational requires the capacity to take the deliberative stance toward one's beliefs.

formed in the context of a suitably framed deliberation. This indicates that belief itself, rather than belief as formed/revised within a deliberative context involving an exercise of the concept of belief, is governed by the norm of truth. Second, it is quite clear that, *pace* Shah, we are not saved from the threat of irrationality by a mere reflexive deliberative deployment of the concept of belief. Accordingly, Shah's own position is caught on the horns of the very dilemma he outlines, albeit at the level of first-person doxastic deliberation deploying the concept of belief.

It should be quite clear that the view I have been proposing has no difficulty explaining the possibility of wishful believing. As Hieronymi stresses, we cannot believe that p by an exercise of evaluative control for reasons that fail to show that p is true—and this is so regardless of whether our evaluation involves a deployment of the concept of belief. In other words, we cannot believe immediately (i.e., without acting on ourselves) for extrinsic reasons. But we are perfectly capable of believing mediately for extrinsic reasons. On this view, wishful believing, unlike believing, is purposive. This is not to say that the intention involved in wishful believing needs to be conscious. It can, however, be conscious, and it can even consciously govern a process of doxastic deliberation involving an exercise of the concept of belief, *pace* Shah.

I do agree with Shah in one important respect, namely, that the normative relation that truth bears to believing and belief is essentially conceptual. This, however, is not because the concept of belief, rather than belief, is properly understood as a relatum in the relation that underwrites doxastic norms. Rather, it is because the normative relation that truth bears to believing and belief is underwritten by the cognitive content of belief, which is wholly constituted by concepts in given modes of combination (as I will argue in the next chapter).

CHAPTER 3 THE COGNITIVE CONTENT OF BELIEF

In the previous chapter, I proposed that belief is a cognitive state whose proper function is to sanction and preserve information as such in conceptualized form and argued that the normative relation that truth bears to belief is properly understood as underwritten by the proper functions of believing, understood as an event, and belief. In this chapter, I argue that the cognitive content of belief (i.e., the content that belief, as such, contributes to the cognitive processes of which it is a part) is wholly constituted by concepts as combined in certain ways, where concepts are conceived as abstract objects of a certain sort.

Any good account of the cognitive content of belief must accommodate the following facts: (i) belief states are concrete particulars, and (ii) two distinct belief states can have the same cognitive content. (ii) seems to require one to hold that the cognitive content of belief is not constituted by concrete particulars. Yet (i) seems to require one to hold that the cognitive content of belief is amenable to realization in concrete particulars. There is, of course, little mystery as to how (i) and (ii) can jointly be true—assertion provides a straightforward model for understanding how these seemingly incompatible requirements can be met. Like beliefs, assertions are contentful concrete particulars that can share the same content despite being numerically distinct. They can do so by virtue of tokening suitable types of contentful particulars, namely, expressions. Likewise, beliefs are contentful concrete particulars that can share the same cognitive content despite being numerically distinct by virtue of tokening suitable types of contentful particulars, namely, concepts.

Concepts as Types of Thought Content

The model of assertion suggests that concepts are abstract objects that can be tokened as concrete, psychological objects, i.e., that concepts are types of thought content. Note that the

view that concepts are such objects does not commit one to any specific view concerning how concepts are tokened in physical systems. In particular, it does not commit one to holding that the tokening of conceptual thought in physical systems must be realized by way of mental representation tokens that have a combinatorial syntax and semantics, so that, in any physical system capable of conceptual thought, complex concept tokens are wholly constituted by combinations of simple concept tokens, as the Language of Thought Hypothesis would have it. Yet reflection on the model of assertion does suggest a thesis at the level of the cognitive content of conceptual thought that bears a strong resemblance to the Language of Thought Hypothesis (minus its causal claims).

Complex Concepts and Simple Concepts

Suppose (*per impossibile* perhaps), that some creature not subject to our limitations is a competent user of a purely semantic language consisting of an infinite number of primitives, each homonymous to a distinct English numeral¹ and each synonymous to either an English primitive expression or an English complex expression, and that her language has the same expressive power as English. Since it is doubtful that such a language could be learned, let us suppose that her capacity to use this language is innate.

Such a language would in principle be translatable into English—not by a Davidsonian radical interpreter, to be sure, but with the help of a suitable manual. Suppose that this creature utters a sentence homonymous to the English, ‘thirty-seven,’ and, in her language, synonymous to the English, ‘The cat is on the mat.’ Armed with my translation manual, I can recover the information encoded in her utterance, namely, that the referent of ‘the cat’ bears the relation expressed by ‘on’ to the referent of ‘the mat’ in the corresponding English sentence (assuming

¹ Since the supposition that each primitive in my fictional language is homonymous to a distinct English numeral is made for ease of exposition only and does not affect any of my conclusions, I am taking the liberty to ignore the syntactic complexity of most English numerals.

that her utterance is true). Furthermore, this information is structured, in the sense that it is information about an entity of one type bearing a spatial relation to an entity of another type, and it certainly does seem to be built systematically out of simple elements whose contribution is invariant relative to the complexes in which they figure. Since this information is recoverable for me using nothing but my translation manual and my understanding of the corresponding English sentence, it is contained in her utterance. The moral of this story so far is that even if we can imagine a purely semantic language whose expressive power equals that of English, we would not thereby have imagined a language that is limited to carrying an unstructured content—the thesis that linguistic content is subject to a compositionality and systematicity constraint does not entail that the vehicle which carries this content is subject to this constraint also.²

Being a competent user of a language plausibly requires the capacity not only to employ sentences of that language under appropriate circumstances, but also to understand the content of these sentences. What such a capacity requires is controversial. Minimally, however, the supposition that our imagined speaker is a competent user of her language plausibly entails that she, too, has the capacity to recover the information encoded in her utterance. Given that I am a competent speaker of English, knowledge of the meaning and truth of her utterance alone would license me to infer directly that a number of other English sentences, each of which carries part of the information encoded in her original utterance, must be true also, namely, ‘there is a cat,’ ‘there is a mat,’ ‘something is on something else,’ ‘something is on the mat,’ ‘the cat is on something’ (among others perhaps). For our imagined speaker to count as having the capacity to recover the information encoded in her utterance, it seems, she too would have to be capable of

² For context-sensitive sentences, my translation manual would give the corresponding English sentence; understanding of the corresponding English sentence in context would suffice for me to understand the creature’s sentence in context.

inferring the corresponding sentences in her language. This is a capacity she can have, e.g., simply by virtue of embodying a suitable system of rules of inference.

Further, she would have to possess the relevant concepts—i.e., the *complex* concepts expressed by her original sentence, and the *complex* concepts expressed by the sentences she can infer on its basis. This plausibly requires that she have the capacity to token these complex concepts. Yet just as we can imagine a language that represents the relevant information via linguistic primitives carrying complex, structured information, so too can we imagine an analogous internal system of representation. So, nothing we have said so far commits us to requiring that she token these complex concepts in such a way that the structure of her complex concept tokens be isomorphic to the structure of the complex concepts of which they are tokens. What this reflection suggests, then, is that the conclusion we reached above concerning linguistic content generalizes to cognitive content; that is, the thesis that cognitive content is subject to a compositionality and systematicity constraint does not entail that the vehicle which carries this content is subject to this constraint also. Concerning the cognitive content of thought, I do endorse the thesis that the cognitive content of any complex concept token is a function of the cognitive content of the simple concepts that, so combined, constitute the tokened complex concept. Yet since this thesis does not entail that the vehicle that carries this cognitive content is itself subject to a compositionality/systematicity constraint, it is compatible with the denial of the thesis that complex concept tokens are fully constituted by combinations of simple concept tokens.

Informational Content, Linguistic Content and Cognitive Content

I am now in a position to offer a more precise characterization of what I mean by ‘cognitive content’ as it applies to belief. Once again, linguistic communication provides a useful

reflective starting point. As H. Paul Grice famously pointed out, much of what we learn in our conversations is not mentioned in them. So, if I ask you, ‘would you like a cookie?’ and you answer, ‘I’m on a diet,’ although you have not said whether you would or wouldn’t like one, you have said enough for me to gather that you would rather pass. Another way of putting the point is that assertions typically convey far more (putative) information than they contain in virtue of their meaning. Even silence can convey information—if I ask you, ‘What do you have to say in your defense?’ and you clearly refrain from answering my question despite having heard and understood it, I may take this to indicate that there is nothing that you have to say, or are willing to say, in your defense, and I may very well do so rightly. In such a case, although it seems odd to say that your silence, as such, contains information, it nevertheless seems appropriate to say that your silence conveys information.

In *Knowledge and the Flow of Information*, Dretske proposes a theoretical framework that enables one to explain how, e.g., silence can manage to convey information and how assertions can convey more information than they contain in virtue of their meaning. On Dretske’s account, information carrying does not require meaningfulness—any contingent event or state of affairs can carry information. Dretske defines informational content as follows (Dretske 1981: 65):

A signal r carries the information that s is F = The conditional probability of s ’s being F , given r (and k), is 1 (but, given k alone, less than 1),

[where k stands for what the receiver already knows (if anything) about the possibilities that exist at the source³ and r stands for an event, condition or state of affairs the existence (occurrence) of which depends on s ’s being F .]⁴

³ Dretske says rather too little about what k is supposed to include. I understand it to include every piece of knowledge that is required for the relevant information to be available to the receiver. So, in order for a tree-trunk’s having seven rings to carry for me the information that the tree to which it belongs is seven years old, I would require knowledge of the relevant nomic correlations. If I did not possess this knowledge, the information would not be available to me.

Provided that one grants the perhaps odd-sounding claim that, for any particular, o , and any property, G , that o instances, o 's being G (or that o is G) depends on o 's being G ,⁵ Dretske's definition entails that, for any property, H , that r instances, r carries the information that r is H (given the relevant knowledge).

The claim is far more innocuous than it might seem. So, consider the following example: I witness Basil being assaulted by bears. Given the knowledge that I do possess and that enables me to discriminate Basil from other individuals, bears from other animals, bears from individuals wearing bear suits, genuine attacks from staged mock-attacks, etc., the event carries, for me, the information that Basil is assaulted by bears (for someone else, it might merely carry the information that, e.g., some individual is assaulted by large, furry creatures). Now, assertions are events and, as such, can carry information in Dretske's sense. Suppose that, upon witnessing Basil being assaulted by bears, you exclaim, 'Basil is being assaulted by wolverines!' Since your assertion is false, it does not carry the information that Basil is being assaulted by wolverines. Nevertheless, your assertion may carry plenty of information for me, given the relevant knowledge—e.g., that the event surprises, or frightens, or pleases you, that you are about to take flight, or gawk, or rush to Basil's rescue, and much else besides. Within this informational profusion, there is one type of information that is of special interest for my purposes here. So, given that I

⁴ As stated, Dretske's definition might strike one as rather puzzling—probability relations, it seems, hold between event types rather than events. Since Dretske makes it clear that he intends the relevant probability to hold between event types—he writes:

In saying that the conditional probability (given r) of s 's being F is 1, I mean to say that there is a nomic (lawful) regularity between these event types, a regularity which *nominally precludes* r 's occurrence when s is not F . (*Ibid*: 65, n. 1)

—we may take his definition as short-hand for a more complex definition including a probability statement holding between event types and a statement to the effect that the occurrence of r is an event of a type mentioned in the probability statement.

⁵ The oddness of the claim may be due to nothing more than its triviality. So, if some event, A , requires the occurrence of some event, B , for its occurrence, it seems appropriate to say that the occurrence of A depends on the occurrence of B . By parity of use, since the occurrence of A trivially requires the occurrence of A , it seems appropriate (though utterly uninformative) to say that the occurrence of A depends on the occurrence of A . Of course, if the dependence relation at hand were, say, causal, it could not be born reflexively by A . Yet what holds for causal dependence relations need not hold for all dependence relations.

know that your assertion is made in English and given that I am a competent speaker of English, your utterance carries for me the information that it means that Basil is being assaulted by wolverines and that it is true just in case Basil is being assaulted by wolverines. If this reflection is correct, then, on Dretske's account, any meaningful utterance carries information concerning its meaning and truth conditions, given the relevant knowledge—where the relevant knowledge, on Dretske's account, is knowledge of the information-carrying role in a given language of the type of utterance of which the utterance is a token (by parity of reasoning, the same can be said of any meaningful expression token, *mutatis mutandis*).⁶ This is a happy consequence; if Dretske's account failed to yield this result, it would leave the possibility of linguistic communication—especially when it involves false sentences—rather mysterious. On this view, then, any meaningful utterance (and expression) carries information about its linguistic content, where 'linguistic content' may be defined as follows:

Let *s* be a particular utterance in language *L*, and let *k* be the knowledge that any competent speaker of language *L* has just by virtue of being a competent speaker.

Utterance *s* carries the linguistic content that *p* just in case, from *k* alone and the supposition that *s* is true, it is immediately *a priori* derivable that *p*.⁷

This approach is theoretically fruitful. So, by adding to *k* above knowledge of the conversational context in which *s* is uttered and knowledge that the speaker of *s* is conversationally cooperative (in the manner specified by Grice), one may derive the pragmatically implied content of *s* from

⁶ I am adopting Dretske's phraseology here. According to Dretske,

The symbol token inherits its meaning from the symbol type of which it is a token; and the symbol type has an information-carrying role independent of the success (if any) of its tokens in carrying this information. (*Ibid*: 192)

On this view, knowledge of the meaning of a symbol is knowledge of the information-carrying role of the symbol type of which it is a token. Although I find this manner of speaking congenial, I do not think that anything of great consequence to my project hinges on endorsing information-theoretic semantics—with suitable modifications, I think, any account of meaning and linguistic competence would do.

⁷ The qualifiers 'alone' and 'immediately' mean to address the uniqueness problem which *a priori* derivability generates (since an infinity of propositions are *a priori* derivable from any sentence).

its linguistic content. Along similar lines, it is plausible that one may derive the content pragmatically conveyed by silence in certain conversational contexts.

I submit that the cognitive content of belief, i.e., the content that belief, as such, contributes to the cognitive processes of which it is a part, is properly understood along similar lines. Like utterances and expression tokens, belief states are contingent particulars whose function is to carry information. Like utterance and expression tokens, belief states perform this function by way of instantiating representational properties. What representational properties a given belief state instances can be characterized in terms of the knowledge required for the information that this belief has the representational properties it has to be available to one. What this information is about is the cognitive content of belief. Using the above definition of linguistic content as a model, I propose that the cognitive content of belief will include at least the following (I will argue later that there is more to the cognitive content of belief than what the following captures):

Let b be a particular belief state such that being in b requires tokening concepts $c_1 - c_n$, and let k be the knowledge that anyone who possesses $c_1 - c_n$ has just by virtue of possessing those concepts.

Belief b carries the cognitive content that p if from k alone and the supposition that b is true, it is immediately *a priori* derivable that p .⁸

I should stress that I take the cognitive content of a belief to be distinct from its referential content. The distinction between referential content and cognitive content means to capture the intuitive distinction between what a belief represents and how it represents whatever it represents. This distinction applies to both utterances and beliefs. So, suppose I uttered sincerely (prior to Ronald Regan's death), 'Ronald is a clown.' Depending on whom I had in mind as I made this utterance, my utterance can be reported truly by the sentence, 'Elka said that the corporate mascot of the fast food giant McDonald is a clown,' or by the sentence, 'Elka said that

⁸ As it goes for linguistic content, so it goes for cognitive content (see n. 7, this chapter).

the 40th president of the United States is a clown,’ or by any other indirect discourse sentence that attributes a declarative utterance to me and whose subordinate clause correctly discloses the referent of my original assertion and the property that I attributed to that referent. Suppose I had the president in mind. If so, my utterance can be reported truly by the sentence, ‘Elka said that the 40th president of the United States is a clown.’ Yet that the 40th president of the United States is a clown is not immediately derivable *a priori* given knowledge of English alone and the supposition that my utterance is true. The (putative) information encoded in my utterance in virtue of its meaning, then, is merely that the individual to whom the name, ‘Ronald,’ refers instances the property of being a clown (the context in which I made this utterance might supply additional information as to whom I referred to by so using the expression, ‘Ronald’).⁹ Nor is the (putative) information that the 40th president of the United States is a clown encoded in the belief my utterance expressed—I could have had this belief while failing to have the concept **president**. Of both beliefs and statements, then, the following seems true: Two beliefs (statements) have the same truth conditions if they have the same referential content; however, two beliefs (statements) may differ in cognitive (linguistic) content even if they have the same referential content. So, if content is involved in the individuation conditions of beliefs (and statements), it would have to be cognitive (linguistic) content. The account I am outlining here means to take a stance on the nature of the cognitive content of belief only; it means to remain neutral concerning the nature of the referential content of belief (though I do hold that the cognitive content of a belief determines its referential content—whatever it is). The cognitive content of belief, I claim, is wholly constituted by the concepts the belief tokens.

⁹ Obviously, the cognitive content of a belief that would correctly be expressed by the utterance, ‘Ronald is a clown,’ need not include any (putative) information concerning Ronald’s name.

Directly Referring Concepts, Predicative Concepts and Logical Concepts

The claim that the cognitive content of a belief is wholly constituted by the concepts it tokens, together with (i) the claim that the cognitive content of any complex concept token is a function of the cognitive content and mode of combination of the simple concepts that, so combined, constitute the tokened complex concept, (ii) the view that any content that can be carried linguistically by an utterance can also be carried cognitively by a belief, and (iii) the view that linguistic contents can have directly referring components, commits me to the view that there are directly referring concepts.

Directly referring concepts, however, do not sit well with certain phrases that are commonly employed in discussions of concepts. Things are said to fall under concepts. A thing falls under a concept by virtue of instancing the property that the concept picks out. In such cases, it seems, the relation between concept and world is analogous to the relation between predicate and world. What a predicate contributes to the linguistic content of a sentence is the (putative) information that the property expressed by the predicate is instanced. Similarly, it might be thought that what a concept contributes to the cognitive content of a thought is the (putative) information that the corresponding property is instanced. If there are directly referring concepts, however, the relationship between such concepts and the things to which these concepts apply is markedly different and is not amenable to explication along a predicative model.

Consider the belief that I expressed by uttering the sentence, ‘Ronald is a clown.’ This belief is about a particular individual, namely, Ronald. Let us call my concept of this individual ‘**Ronald**.’ Even if the Ronald I have in mind has the property of being named ‘Ronald,’ it does not seem that the content that the concept, **Ronald**, contributes to the cognitive content of my belief is that the property of being named ‘Ronald’ is instanced. I may very well know two

individuals named ‘Ronald,’ believe that each is so named, and believe that one is a clown and that the other isn’t, without there being any incompatibility in cognitive content among these beliefs. Further, although I would need to believe that Ronald is named ‘Ronald’ in order to express my belief as I did, I can certainly believe that Ronald is a clown without having any belief concerning Ronald’s name. Nor does the concept **Ronald** seem to contribute (putative) information concerning the instancing of some other complex property, such as being the 40th president of the United States or being the actor who played the part of Drake McHugh in *Kings Row*—for any such property, it seems, I may come to believe that Ronald fails to instance it without thereby coming to believe something whose cognitive content is incompatible with the cognitive content of my belief that Ronald is a clown. I may, of course, rely on some such property to recognize Ronald. Nevertheless, it does not seem that the (putative) information that **Ronald** contributes to the cognitive content of my belief that Ronald is a clown concerns the instancing of any such property. Rather, what the concept, **Ronald**, seems to contribute to the cognitive content of my belief is a referent, namely, Ronald. I can similarly be shown to be committed to the existence of logical concepts, whose contribution to the cognitive content of thought can obviously not be understood along a predicative model, but must instead be seen as analogous to the contribution made by logical terms to the linguistic content of utterances.

To mark this difference in contribution to the cognitive content of thought, I introduce a distinction between directly referring concepts, predicative concepts and logical concepts. As just noted, the cognitive role of a logical concept is properly understood as analogous to the linguistic role of the corresponding logical term. The cognitive role of a predicative concept is to contribute information concerning the instancing of some property, which the concept can be said to predicate. An object falls under a predicative concept in virtue of instancing the property

that the concept predicates. Finally, the cognitive role of a directly referring concept is to contribute a referent for predicative thought (or for thought *tout court*, if I can properly be said to think, e.g., of Ronald, without attributing any distinct property to him).

Synchronic Cognitive Content vs. Individuating Cognitive Content

Now, it may very well be that we typically token beliefs about particulars by way of tokening complex referring concepts whose cognitive content is constituted by a simple directly referring concept and some predicative concept (either a simple predicative concepts or a complex predicative concept constituted by simple predicative concepts and logical concepts), using the predicative component to keep track of the referent(s) of our thought. Nevertheless, it is plausible that the cognitive content of our beliefs about particulars does involve simple directly referring concepts, and that these play a crucial role in our cognitive economy.

Indeed, though the predicative components we use to keep track of the referents of such beliefs often change over time, the core components of their cognitive content may remain unchanged (that they can so remain seems to be, for creatures like us, a prerequisite for cognitive success). So, suppose that some individual wearing a bear suit (in a room full of ordinarily dressed people) makes a clever comment. Since his wearing a bear suit in such a setting is rather striking, I may rely on this feature to keep track of him and come to have a belief that carries cognitively the content carried linguistically by the sentence, ‘The individual wearing a bear suit has made a clever comment.’ Should I see him remove his costume to reveal Spiderman pajamas, I may come to keep track of him as the individual wearing Spiderman pajamas. In such a case, if all goes as it should, I would also quite naturally adjust the cognitive content of my belief by updating its tracking concepts to preserve the core components of its cognitive content,

namely, that **[he]** has made a clever comment (where ‘**he**’ stands for a simple directly referring concept whose referent is the individual who made the clever comment).

Although this example doubtlessly fails to do justice to the intricacy of the predicative components typically involved in our complex referring concepts, the basic picture it suggests seems correct. The cognitive content of our beliefs about particulars does seem to involve simple directly referring concepts. Further, it is plausible that simple directly referring concepts are seldom tokened as such (i.e., as components of cognitive contents whose logico-semantic form can be represented as, e.g., ‘**a is F,**’ where ‘**a**’ stands for a simple directly referring concept and ‘**F**’ stands for a predicative concept), but are typically tokened instead as constituents of complex referring concepts involving a predicative element whose cognitive role seems to be to track the referent of the concept (i.e., as constituents of cognitive contents whose logico-semantic form can be represented as, e.g., ‘**a[the X] is F,**’ where ‘**X**’ stands for a predicative concept, ‘**a[the X]**’ stands for a complex referring concept, and ‘**[the X]**’ represents **X** as predicating *X* uniquely of *a*). Nevertheless, the cognitive content of beliefs involving complex referring concepts seems to be individuated (over time at least) in terms of the simple directly referring component, rather than the predicative component, of their constituent complex referring concept(s). So, in the example just considered, it does not seem that I relinquished one belief and acquired another once the individual who made a clever comment removed his bear suit; rather, what is essential to the identity of my belief over time seems to have remained constant despite my using different predicative concepts to keep track of the object of my belief. Reflection on this humdrum (if somewhat idealized) case suggests the following theses:

Let there be two cognitive contents whose constituents are:

- (i) **a[the X] is F**
- (ii) **a[the Y] is F**

where ‘**a**’ stands for a simple directly referring concept,
and ‘**F**,’ ‘**X**’ and ‘**Y**’ stand for distinct predicative concepts;

let (i) be the cognitive content of some belief state, b_1 , at some time, t_i ,

let (ii) be the cognitive content of some other belief state, b_2 , at some time, t_j ,

Then:

- (d) b_1 and b_2 have the same individuating cognitive content, namely, **a is F**, where ‘individuating cognitive content’ stands for the cognitive content that individuates a belief over time;
- (e) if some believer, S , has b_1 at t_i , then, at t_i , the canonical way for S to verify that a is F is by way of verifying that the X is F (the same is true, *mutatis mutandis*, for b_2);
- (f) having b_1 and b_2 need not place S in a position to judge that the $X =$ the Y .

Henceforth, I will use ‘synchronic cognitive content’ to refer to the cognitive content that a belief contributes at a given time to the cognitive processes of which it is a part, and ‘individuating cognitive content’ to refer to that part of its cognitive content that individuates it over time.

Note, however, that I am not suggesting that each belief has two distinct contents. Rather, I am proposing that, in the case of our beliefs about particulars at least, each concept that is included in the content that a belief contributes at a time to the cognitive processes of which it is a part need not be included in the cognitive content that individuates that belief over time.

The individuating and synchronic cognitive contents of belief may be defined as follows:

Individuating Cognitive Content:

Let b be a particular belief state such that being in b requires tokening concepts $c_1 - c_n$, and let k be the knowledge that anyone who possesses $c_1 - c_n$ has just by virtue of possessing those concepts.

Belief b is individuated by the cognitive content that p just in case from k alone and the supposition that b is true, it is immediately *a priori* derivable that p .

Any two particular belief states, b_1 and b_2 , that meet the above condition count as being of the same individuating type (i.e., as being the same belief, where ‘belief’ is used to refer to a type of state rather than to a particular state) by virtue of having the same individuating cognitive content.

Synchronic Cognitive Content:

Let b be a particular belief state individuated by the content that p , and let S have b at t .

If the individuating cognitive content of b can be represented as **... a is F ...**, then the synchronic cognitive content of b at t can be represented as **... a [the X] is F ...** just in case the canonical way for S to verify at t that a is F is for S to verify that the X is F .

Whenever two particular belief states, b_1 and b_2 , meet the above condition, b_1 and b_2 have the same synchronic cognitive content (i.e., b_1 and b_2 make the same type of contribution, while they are undergone, to the cognitive processes of which they are a part).

Introspection corroborates the theses outlined above, i.e., theses (a), (b) and (c). Our beliefs concerning particulars seem to be about particulars, rather than about particulars as picked out in specific ways. Although the ways in which we pick out particulars may change significantly over time—so significantly, indeed, that none of the properties in light of which we identify an object at one time need be included among the properties in light of which we identify it at a later time—our beliefs about them may nevertheless remain constant. Further, although our cognitive make-up seems to be such that, at any given time, we typically identify particulars in light of a specific set of properties (we may even be subject to a nomological compulsion to do so when the particulars involved are concrete particulars with which we are perceptually acquainted) and, so, think of them at that time as falling under the conjunctive predicative concept that corresponds to this set of properties, the role that such predicative concepts play in our cognitive lives is ancillary (albeit important). Although they are useful (in some cases perhaps even indispensable) in identifying the particular object of our belief at a time (or period of time) and seem to be part of the synchronic cognitive content of that belief at that time (or period of time), they are not part its individuating cognitive content.

Predicative concepts, then, seem to make different contributions to the cognitive content of our beliefs about particulars depending on whether their cognitive role is to identify the object of our belief (let us call this role ‘the identifying role’; henceforth, I will use ‘identifying concept’

to mean ‘predicative concept tokened in an identifying role’) or to attribute a property to that object (let us call this role ‘the attributive role’; henceforth, I will use ‘attributive concept’ to mean ‘predicative concept tokened in an attributive role’). When, relative to a given belief, a predicative concept plays an identifying role, it is not part of the individuating cognitive content of that belief, though it is part of its synchronic cognitive content while playing that role. By contrast, when a predicative concept plays an attributive role relative to a given belief, it is part both of the individuating cognitive content of that belief and of any synchronic cognitive content that belief might have at any given time. These considerations lend support to Thesis (a) above.

Although identifying concepts are not part of the individuating cognitive content of our beliefs about particulars, they are nevertheless crucial to our cognitive lives. Indeed, at any given moment, our cognitive access to particular objects of belief is typically by way of identifying concepts—at any given moment, we are typically, as it were, behind the veil of the identifying concepts we token at that moment.¹⁰ For this reason, the only way we typically have of answering questions as to what to believe about a given object of belief is by considering the object *as we currently identify it*. For this reason also, in cases in which we hold several beliefs about a single object, either simultaneously or over time, and in which each belief identifies this object in a distinct way, we may fail to realize (and to be in a position to realize) that they are beliefs about a single object. This explains an unfortunate feature of our cognitive lives. Although we seem cognitively incapable of consciously accepting beliefs that identify their object identically

¹⁰ The qualification ‘typically’ means to allow for the existence of beliefs whose contents include simple directly referring concepts tokened independently of any identifying concepts. The only beliefs that plausibly answer this description are first-person beliefs. Further, it is plausible that, when a simple directly referring concept is tokened by its referent in combination with an identifying concept through which the referent identifies itself, this simple directly referring concept is not tokened first-personally. On the account I am proposing, then, properly *de se* beliefs are beliefs that involve a simple directly referring concept tokened by its referent independently of any identifying concepts. I will return to this issue in Chapter 5.

and attribute contradictory properties to it—e.g., to beliefs whose synchronic cognitive contents have the form ‘**a[the X] is F and a[the X] is not-F**,’ at least in cases when the complexity of the attributive concepts involved is not too great—we can (and, sadly, often do) fall prey to beliefs whose individuating cognitive content is contradictory but that identify their object in distinct ways—e.g., to beliefs whose synchronic cognitive contents have the form ‘**a[the X] is F and a[the Y] is not-F**.’

Consideration of the phenomenology of belief, then, lends support to the proposed distinction between synchronic and individuating cognitive content, and to the attendant account of the contribution that identifying and attributive concepts make to these contents (i.e., to theses (a)-(c) above). Since these theses could not be formulated without positing the existence of simple directly referring concepts, the plausibility of these theses speaks in favor of the view that there are such concepts and that they are required components of both the synchronic and the individuating cognitive contents of our beliefs about particulars (any belief that does not include such a component is simply not a belief about a particular).

Beliefs About Non-Existent Particulars

Yet this view faces a challenge, which mirrors the challenge faced by semantic theories that posit the existence of directly referring expressions. We can, it seems, have beliefs about non-existent particulars. Little Basil can believe that Rudolph is a flying reindeer that has a shiny red nose, and crazy Jane can, in a fit of paranoid delusion, believe that the man in black who has been following her all day is a secret agent, and she can do so even if no one has been following her at all. What is a proponent of the view that simple directly referring concepts are required components of the cognitive content of our beliefs about particulars to say about the (putative) content of (putative) beliefs about non-existent particulars?

There seem to be two theses that a proponent of directly referring concepts might endorse in response to this challenge. According to the first thesis, there can only exist a directly referring concept if the corresponding object exists also. On this view, we cannot have beliefs *about* non-existent particulars. We can, however, be under the illusion of having such a belief when we token an identifying concept and utilize it in thought as if it were associated with a simple directly referring concept. So, e.g., one may be under the illusion that the synchronic cognitive content of one's belief has the form, '**a[the X] is F,**' and that its individuating cognitive content has the form, '**a is F,**' when the actual form of its synchronic and individuating cognitive content is '**the X is F.**' On this view, then, such beliefs are not about particulars at all (though one may be under the illusion that they are). If one construes the satisfaction conditions of such beliefs to require merely that there be a unique object that falls under the identifying concept corresponding to '**X**' and the attributive concept corresponding to '**F**' in the schema above, such beliefs are truth-evaluable (when there is no such object, they are false).

On this view, the form of the existential commitment generated by beliefs whose cognitive content involves a simple directly referring component is distinct from that generated by beliefs whose cognitive content involves no such component. In the former case, believers are committed to the existence of the referent of any of the simple directly referring concepts their beliefs token (in other words, the commitment involves a referential component). When the synchronic cognitive content of such a belief includes identifying concepts, believers undertake the additional commitment that the referent of their simple directly referring concepts fall under the relevant identifying concepts (at the relevant time). In the latter case, by contrast, the commitment involves no referential component (since no such component is involved in the synchronic

cognitive content of the relevant belief). So, the commitment is merely to the existence of some unique object falling under the relevant identifying concepts (at the relevant time).

According to the second thesis, there can exist a directly referring concept even if there exists no corresponding object. On this view, simple directly referring concepts are concepts that operate in thought just as if they contributed a referent, and do so irrespective of the existence of their putative referent. According to the second option, the cognitive content of beliefs about non-existent particulars can have the same logico-semantic form as the cognitive content of beliefs about existent particulars. The difference between beliefs about existent and non-existent particulars is that the former, but not the latter, are truth-evaluable. Just as the logico-semantic form of the cognitive content of beliefs about particulars is independent of the existence of their putative object, so is the form of the commitment that such beliefs generate. By virtue of having beliefs about particulars (be they existent or non-existent), believers are committed to the existence of the referent of any of the simple directly referring concepts their beliefs token. When the synchronic cognitive content of such a belief includes identifying concepts, believers undertake the additional commitment that the referent of their simple directly referring concepts fall under the relevant identifying concepts (at the relevant time).

The virtue of the second approach is that it offers a unified account of the cognitive contents of beliefs whose cognitive contents seem, introspectively, to be identical. Its drawback is that it posits the existence of concepts which, though capable of being possessed, are semantically defective in such a way that beliefs involving them are not truth evaluable. Such a view would be especially uncongenial to proponents of the view that concepts are constituents of propositions (if propositions are taken to be truth-evaluable entities). The first approach avoids the difficulties faced by the second approach. Yet the first approach conflicts with introspection,

as it assigns distinct synchronic cognitive contents to beliefs whose cognitive contents seem, introspectively, to be identical. So, the first approach seems committed to denying that believers have privileged access to the cognitive content of their own beliefs. Since settling this issue is not crucial to the remainder of my project, I need not choose between these approaches. That simple directly referring concepts are required to account for the identity of beliefs about particulars over time, I take it, is reason enough to posit their existence, despite these theoretical difficulties.

Beliefs about Misidentified Particulars

A proponent of the view I have been advocating must answer another question, namely, what are the truth conditions of beliefs whose synchronic cognitive content includes a complex directly referring concept in cases when the referent of their component simple directly referring concept does not fall under their component identifying concept? In other words, what are the truth conditions of beliefs whose synchronic cognitive content includes, e.g., a concept of the form, '**a[the x]**,' when *a* either does not fall, or is not the unique object falling, under *X*? One might doubt that such a case is possible. Indeed, for a belief to be about a given particular, it seems, that belief must somehow succeed in identifying that particular. On this view, a belief cannot succeed in referring to an object while misidentifying it—beliefs either succeed in identifying their object (by identifying it correctly), or fail to refer altogether. I will show in Chapter 5 that such beliefs are possible, as they may arise in the context of reasoning.

Accounting for the possibility that a belief may misidentify its object requires one to hold that the truth-value of a belief is determined by its individuating cognitive content. Let us suppose that the synchronic cognitive content of the belief in question has the form, '**a[the x] is F.**' On the view that individuating cognitive content determines truth-value, this belief is

true just in case a falls under F . Nevertheless, a proponent of this answer may hold that a belief that includes a concept of the form ‘ $\mathbf{a[the\ X]}$,’ though true, is doxastically defective when a either does not fall, or is not the unique object falling, under X in that it fails to place the believer in a suitable position to verify that her belief is true—given that, *per* thesis (b) above, the canonical way for S to verify that a is F is by way of verifying that the X is F , when a does not fall, or is not the unique object falling, under X , S is not in a position to verify that a is F .

The alternative view, according to which the truth-value of a belief is determined by its synchronic cognitive content, is implausible, since it is committed to the untoward result that the truth-value of a belief may change even if there is no corresponding change in the particular that this belief represents (this is so regardless of whether one takes beliefs that include a concept of the form ‘ $\mathbf{a[the\ X]}$ ’ to be false, or non-truth evaluable, when a does not fall, or is not the unique object falling, under X). So, suppose that a is F . Suppose, further, that a is always the Y and never the X . Now, suppose I have a belief individuated by the content, $\mathbf{a\ is\ F}$, and which has the synchronic cognitive content, $\mathbf{a[the\ X]\ is\ F}$, at t_1 and the synchronic cognitive content, $\mathbf{a[the\ Y]\ is\ F}$, at t_2 . On the view that the truth-value of a belief is determined by its synchronic cognitive content, my belief would have been false (or non-truth evaluable) at t_1 , then true at t_2 , even though the particular represented by my belief underwent no relevant change (and even though my belief has had the same individuating cognitive content throughout).

Concepts as Individuated by their Possession Conditions

These issues set aside, the general thesis that the foregoing reflection suggests is that there are different types of concept, which differ by virtue of the type of contribution they make to the synchronic/individuating cognitive content of thought. The distinction among directly referring concepts, predicative concepts and logical concepts seems to provide an exhaustive taxonomy of

these types (concerning predicative concepts, a further distinction must be made between their contribution when tokened in an identifying role and their contribution when tokened in an attributive role). But this says little about the nature of concepts. Concerning the nature of concepts generally, I endorse the general outlines of the account Peacocke proposes in *A Study of Concepts*, according to which concepts are fully individuated by their possession conditions, which specify the conditions under which a thinker possessing the relevant concept finds transitions to and from propositional attitudes containing this concept primitively compelling. Note that the type of contribution that a given concept makes to the cognitive content of thought is specified by the possession conditions of that concept. On Peacocke's view, as on the view I presented above, concepts are either simple or complex, and the possession conditions of complex concepts are a function of the possession conditions of their simple component concepts and their modes of combination. On the view I am endorsing, then, according to which the content of belief is fully constituted by concepts in given modes of combination, the synchronic and individuating cognitive contents of belief are complex concepts; so, the synchronic and individuating cognitive contents of belief are also fully individuated by their possession conditions. So, my endorsement of Peacocke's account of concepts seems to commit me to holding that one can only have some belief, *b*, if one fully satisfies the possession conditions of its simple component concepts.¹¹

Before I proceed to discussing Peacocke's account of concepts, I should pause and take stock of what I have accomplished so far. I stated at the outset that I intended to articulate and defend the view that belief is a psychological state individuated by its aim, function and cognitive content. I have argued that the cognitive content of belief is wholly constituted by concepts.

¹¹ Peacocke invokes a distinction between belief possession conditions and belief attribution conditions to deny this commitment; I discuss this distinction below.

Reflection, however, revealed the need to distinguish between the cognitive content a belief has at a time and the cognitive content that individuates this belief over time. In the case of our beliefs about particulars, individuating cognitive content may be a proper part of synchronic cognitive content (where both types of content are conceived as sets of simple concepts, rather than complex concepts), while in the case of our general beliefs, individuating and synchronic cognitive contents presumably always coincide. If so, then belief is a psychological state individuated by its aim, function and individuating cognitive content. If this reflection is correct, then it may be useful to distinguish between belief and the cognitive states that realize belief. On the proposed nomenclature, the cognitive states that realize belief differ from belief in that their individuating cognitive content always coincides with their synchronic cognitive content. If so, then the relation between our general beliefs and the cognitive states that realize them is identity, while the relation between our beliefs about particulars and the cognitive states that realize them is, presumably, supervenience. This is so because, if two beliefs about particulars have distinct individuating cognitive contents, then the cognitive states that realize these beliefs must also have distinct individuating cognitive contents, while two cognitive states that have distinct individuating cognitive contents may nevertheless realize the same belief.

Let us now turn to Peacocke's account of concepts, according to which concepts are fully individuated by their possession conditions, which specify the conditions under which a thinker who possesses the relevant concept finds transitions to and from propositional attitudes containing this concept primitively compelling. This account must answer the following questions:

- (i) What does it take for a set of possession conditions to be the possession conditions of a concept?
- (ii) Does there correspond a single concept, or a multiplicity of concepts, to a vague expression, and is(are) the relevant concept(s) also vague?

Requirements on Concept Possession Conditions

Peacocke argues that every concept must have a determination theory, which he characterizes as follows:

The determination theory for a given concept (together with the world in empirical cases) assigns semantic values in such a way that the belief-forming practices mentioned in the concept's possession condition are correct. That is, in the case of belief formation, the practices result in true beliefs, and in the case of principles of inference, they result in truth-preserving inferences, when semantic values are assigned in accordance with the determination theory. (Peacocke 1992: 19)

Peacocke talks of concepts generally as referring, and of the requirement that a concept have a determination theory as the requirement that its possession condition determine consistent reference conditions—on his nomenclature, the semantic value of a concept is the referent of that concept (*Ibid*: 16-17). Since I do not think that predicative or logical concepts can properly be said to refer to anything (predicative concepts predicate properties, but do not refer to them, while logical concepts operate on concepts), I prefer to talk of deployment conditions. On this nomenclature, the determination theory of a concept spells out consistent deployment conditions for that concept. Consistent deployment conditions for referential concepts spell their reference conditions; consistent deployment conditions for predicative and logical concepts spell their satisfaction conditions. Peacocke's requirement on concepthood, as I understand it, is that the possession conditions of any concept must determine a consistent set of deployment conditions.

Further, Peacocke identifies possessing a concept with knowing what it is for something to be its semantic value (he calls this feature of his account 'The Identification' (*Ibid*: 23)). For Peacocke, grasping a concept simply is knowing what it is for something to be its semantic value,¹² where the relevant knowledge either does not involve thinking of the concept under a

¹² Since I do not think that all concepts refer, I will understand Peacocke's Identification as involving knowing what things must be like for the concept to apply, rather than knowing what it is for something to be the semantic value of that concept (i.e., on Peacocke's nomenclature, the referent of that concept).

mode of presentation at all, or involves thinking of it under a mode of presentation available only to those who already possess the concept. Peacocke claims that his determination theory requirement underwrites his Identification. Since he states that his Identification applies to arbitrarily complex concepts, including complete propositional concepts (*Ibid*: 22), the question arises whether his determination theory requirement, too, applies to arbitrarily complex concepts, including complete propositional concepts. I will argue later that Peacocke's determination theory requirement needs to be restricted to simple concepts.

Peacocke offers **tonk** as an instance of a putative concept that has no determination theory. Arthur Prior defines 'tonk' by the following two rules:

- (1) From p , it can be inferred that p tonk q
- (2) From p tonk q , it can be inferred that p (*Ibid*: 21)

If **tonk** were a concept, it would be that concept **C** to possess which a thinker must find transitions to and from propositional attitudes that are instances of the forms, (1) and (2), primitively compelling, and must do so because they are of these forms. Peacocke's view is that **tonk** is a spurious concept because (1) and (2) fail to determine a consistent set of satisfaction conditions. In the case in which p is true and q is false, (1) requires p tonk q to be true, while (2) requires p tonk q to be false. Accordingly, **tonk** has no determination theory, as the belief-forming practices mentioned in its possession condition fail to result in truth-preserving inferences.

Yet the expression, 'tonk,' enters philosophical discussions—Peacocke's included. If both Peacocke's view that **tonk** is not a genuine concept and my account of the cognitive content of belief and its relation to linguistic content are correct, what are we to make of his statement that **tonk** is not a genuine concept? Are we to hold that this is a statement that can neither be believed nor disbelieved—more precisely, that this statement carries a linguistic content that

cannot be carried cognitively by any belief, *contra* the claim I endorsed earlier that any content that can be carried linguistically by an utterance can also be carried cognitively by a belief—and that the statement, ‘Peacocke believes that **tonk** is not a genuine concept’ is either false or non-truth-evaluable? Peacocke is not without resources for answering this question. So, he could answer that in the statement, ‘**tonk** is not a genuine concept,’ ‘**tonk**’ expresses a set of possession conditions rather than a concept. To this set there corresponds a complex concept, which can suitably be combined with the concept, **concept**, and so constitute the cognitive content of the belief expressible by the sentence, ‘**tonk** is not a genuine concept.’ Alternatively, Peacocke could elect to deal with this issue meta-linguistically. So, he could claim that, properly cast, his view is that the expression, ‘tonk’ does not express a genuine concept. On either answer, then, there is a sentence (either ‘**tonk** is not a genuine concept,’ where ‘**tonk**’ expresses a set of possession conditions, or ‘‘tonk’ does not express a genuine concept’) that is amenable to belief, and belief attributions involving this sentence are truth-evaluable.

These two (compatible) lines of answer are plausible. Yet the problem just raised with regard to ‘tonk’ generalizes to expressions concerning which analogous answers are less plausible. So, Peacocke’s account seems committed to holding that no genuine concept corresponds to any expression for which semantic paradoxes can be generated. So, e.g., Peacocke seems committed to the view that ‘truth’ does not express a genuine concept. Despite the central role that truth is called to play in my project, this is bullet I have to bite, though without needing therefore to abandon the claim that truth (and **truth**) plays a crucial role in belief evaluations. The alternative—abandoning Peacocke’s requirement of a determination theory—is far too implausible, at least as it applies to simple concepts, as it is not clear what a putative concept with inconsistent deployment conditions could be a concept of. Accordingly, I introduce the technical term,

‘notion,’ for sets of possession conditions that do not meet Peacocke’s requirement for concepthood, according to which something is a concept only if its possession conditions determine a consistent set of deployment conditions; a notion, then, is individuated by a set of “possession”¹³ conditions for which no determination theory can be formulated. Some notions, such as **tonk**, wear their failure to satisfy Peacocke’s requirement for concepthood on their face.

Others, such as **truth**, do not, and discovering that they are not genuine concepts may require a great deal of careful investigation. Being a notion, I think, is no bar to playing a significant role—indeed, an indispensable role, in the case of **truth**—in the cognitive lives of their “possessioners.”¹⁴

As I noted earlier, I have committed myself to the view that the synchronic and individuating cognitive contents of belief are complex concepts, which are fully individuated by their possession conditions. I have also committed myself to the view that the cognitive content of any complex concept is a function of the cognitive content of its simple constituent concepts and their mode of combination. This accords with Peacocke’s view, according to which the possession conditions of any complex concept are reducible to the possession conditions of its simple constituent concepts and their modes of combination (the former can be derived from the latter via a procedure he calls ‘multiplying out’ (*Ibid*: 111-112)). Now, since I have endorsed Pea-

¹³ I am using scare-quotes here because it is not clear that there is anything that is amenable to being possessed in such a case.

¹⁴ It is not clear that a paradox-free substitute such as **truth-in-English** could play the role that **truth** plays in our cognitive lives. When I wonder whether some belief of mine is true, I do not wonder whether it would be true-in-English should I express it in English, true-in-French should I express it in French, etc.; when I set out to determine whether it is true, I do not consider language.

Note that although the liar paradox generates difficulties at the conceptual level insofar as we can have beliefs about the truth of sentences/utterances, it is not clear that the liar paradox can be run directly on non-linguistic beliefs (the sentence, ‘this belief is true’ grossly underdescribes the relevant self-referential belief). **Truth** as applied to non-linguistic beliefs might be paradox-free because the “grammar” of belief distinguishes between levels of beliefs (i.e., object-level beliefs and successive levels of meta-beliefs). Since the possession conditions of **truth** permit us to apply it to liar sentences, **truth** is a notion. Nevertheless, **truth** might be unobjectionable as applied to non-linguistic beliefs. A conceptual fix might involve circumscribing **truth** to non-linguistic items and employing language-relative analogues, e.g., **truth-in-English** for linguistic items.

Peacocke's requirement that every concept have a determination theory, I have committed myself to the view that no genuine belief can involve a notion. This is because any inconsistency in the deployment conditions of a constituent concept would result in an inconsistency in the deployment conditions of any complex concept that comprised it. In endorsing Peacocke's requirement, then, I have also committed myself to the more general view that there can be no genuine belief that deploys inconsistent concepts, or that deploys concepts inconsistently. So, I have committed myself to the view that there can be no genuine belief with the contents, e.g., **my account of belief is true, man is condemned to be free**,¹⁵ or **I now have no belief**. I propose to use the term, 'belief-like attitude' for purported beliefs whose content fails to meet Peacocke's requirement of a determination theory, interpreted as a requirement on deployment conditions.

However, there is a significant difference between, on the one hand, the belief-like attitude that my account of belief is true and, on the other hand, the belief-like attitudes that man is condemned to be free or that I now have no belief. All three attitudes purport to represent some content as true. Given that the first attitude involves a notion, there is a sense in which it simply fails to represent; although the component concept, **my account of belief**, succeeds in referring to my account of belief, the notion, **true**, fails to attribute any property to my account of belief (by virtue of failing predicate any property at all). Accordingly, this belief-like attitude is not truth-evaluable, since its content does not fully specify what the world would need to be like for the attitude to be true. By contrast, given that the attitudes that man is condemned to be free and that I now have no belief do not employ notions, they do fully specify what the world

¹⁵ I take it that 'condemned' here is used to mean something along the lines of 'forced,' 'compelled,' 'obligated,' or 'required,' all of which I take to be incompatible with being free (even on a hypothetical interpretation of 'being free'). One who does not think that the content, **man is condemned to be free**, involves incompatible concepts may substitute the less tendentious though somewhat more boring, **this triangle is square**.

would need to be like for the attitude to be true. Accordingly, they are truth-evaluable—given the concepts involved and their modes of combinations, they are necessarily false. I take this to indicate that Peacocke’s requirement that all concepts have a determination theory, interpreted as the requirement that the possession conditions of any concept, simple or complex, must determine a consistent set of deployment conditions, is too strong, and should be restricted to simple concepts.

Note that my restriction of Peacocke’s requirement of a determination theory, interpreted as a requirement on deployment conditions, accords well with his identification of grasping a concept with knowing what it is for something to be its semantic value. For the belief contents, **man is condemned to be free** and **I now have no belief**, anyone who grasps these complex concepts knows just what it is for something to be their semantic value. With some reflection, anyone who grasps them might even know that the belief that man is condemned to be free is necessarily false, and that the belief that I now have no belief is necessarily false while entertained reflexively. Accordingly, I propose to call purported beliefs involving notions ‘belief-like attitudes.’ Purported beliefs that are incoherent by virtue of deploying inconsistent concepts, or by virtue of deploying concepts inconsistently, are genuine beliefs, albeit necessarily false.

Vagueness

Let us now turn to the second question I introduced earlier, namely:

- (ii) Does there correspond a single concept, or a multiplicity of concepts, to a vague expression, and is(are) the relevant concept(s) also vague?

A different problem arises for concepts expressed by vague expressions (which plausibly constitute a significant portion of the vocabularies of natural languages). Concerning vague expressions, one may ask whether, on Peacocke’s account, there corresponds to them a single concept,

or a multiplicity of concepts, and whether the relevant concepts are themselves vague. Peacocke holds that possessing a concept requires knowledge of what it is for something to be its semantic value, and may require tacit knowledge of its possession conditions, which specify the conditions under which a thinker possessing the relevant concept finds transitions to and from propositional attitudes containing this concept primitively compelling (*Ibid*: 149-150). Accordingly, he seems committed to the view that whatever concepts correspond to vague expressions must themselves be vague, on pains of being committed to denying that competent speakers of natural languages possess the concepts that correspond to the vague expressions they employ and, so, to denying that their statements involving vague expressions express genuine beliefs—vague expressions are thought to be vague precisely because competent speakers are unable, even implicitly, to draw precise boundaries for their application. Since competent speakers also differ with regard to when they are uncertain as to whether a given vague expression applies or fails to apply, the view that a single, vague concept corresponds to each vague expression is implausible. Accordingly, Peacocke seems to be committed to the view that there corresponds a multiplicity of overlapping concepts to each vague expression. This yields the seemingly unpalatable result that two competent speakers saying, of the same object at the same time, ‘this is red’ can thereby express beliefs with distinct (though overlapping) individuating cognitive contents. Note, however, that the foregoing argument assumes a principle I endorsed earlier, namely, that any content that can be carried linguistically by an utterance can also be carried cognitively by a belief. I take this principle to be eminently plausible, but Peacocke might deny it. The commitment I have just foisted on Peacocke, then, is conditional on his endorsement on this principle.¹⁶

¹⁶ The foregoing should not be understood as expressing any epistemicist tendencies toward vagueness. In my view, there are both vague concepts and vague expressions (though I find it terrifically hard to imagine that there are any such things as vague properties).

Peacocke seems otherwise committed to this result. According to Peacocke, the possession conditions for the concept, **red**, are:

The concept **red** is that concept **C** to possess which a thinker must meet these conditions:

- (1) He must be disposed to believe a content that consists of a singular perceptual-demonstrative mode of presentation *m* in predicational combination with **C** when the perceptual experience that makes *m* available presents its object in a red' region of the subject's visual field [where the prime indicates a sensational property] and does so in conditions he takes to be normal, and when in addition he takes his perceptual mechanisms to be working properly. The thinker must also be disposed to form the belief for the reason that the object is so presented.
- (2) The thinker must be disposed to believe a content consisting of any singular mode of presentation *k* not meeting all the conditions on *m* in (1) when he takes its object to have the primary quality ground (if any) of the disposition of objects to cause experiences of the sort mentioned in (1).¹⁷ (*Ibid*: 7-8)

Yet Peacocke holds that someone who satisfies condition (1) but not condition (2) would nevertheless possess a distinct, though equally genuine, concept (**red***). This view is plausible.

There is evidence that children do not make the appearance/reality distinction until they are about three years old. Yet they certainly seem capable of recognizing colors and of forming color beliefs when presented with colored objects. Yet this view, together with the account of the individuating cognitive content of beliefs I have been outlining, commits me to holding that a child of two and a typical adult, both of whom say, of the same object at the same time, 'this is red,' may thereby express beliefs with distinct individuating cognitive contents (i.e., distinct beliefs). In this case, the overlap between the individuating cognitive contents of their respective beliefs would be far less than in the case of competent speakers who differ with regard to when they are uncertain as to whether a given vague expression applies or fails to apply. This case, however, poses a lesser threat to our ordinary practices of belief attribution and evaluation since

¹⁷ For the sake of consistency, I have added the font I have been using to mark concepts to Peacocke's text.

there is good reason to think that a child who fails to distinguish between appearance and reality is not a competent user of the expression, ‘red.’

This second bullet, too, I will bite—two competent speakers saying, of the same object at the same time, ‘this is red’ can thereby express beliefs with distinct individuating cognitive contents. Nevertheless, I think that the overlap between the individuating cognitive content of the beliefs that competent language-speakers would express using synonymous sentences is typically so significant that this result does not threaten our ordinary practices of belief attribution and evaluation or render the success of these practices inexplicable.

The Sorites Paradox: Most vague predicates are soritical. That is, their application conditions are allegedly shown to lead to paradoxical results by arguments of the following form:¹⁸

Let P be some predicate

Let F be the property, is in the extension of ‘ P ’

Let $S = \langle a_1, \dots, a_n \rangle$, where, for any a_i , a_i is indiscriminable with respect to F from a_{i+1}

Then:

(P1) Fa_1

(P2) For any a_i , if Fa_i then Fa_{i+1}

(C) So, Fa_n

(where (P1) seems clearly true and (C) seems clearly false).

Where $S = \langle$ a man with no hair on his head, ... , a man with 10^6 hairs on his head \rangle (one man per possible number of hair) and P is the predicate, ‘is bald,’ the argument, if sound, shows that a man with 10^6 hairs on his head is in the extension of ‘is bald.’ By semantic descent, we get the result that a man with 10^6 hairs on his head is bald. This is paradoxical, since a man with 10^6 hairs on his head is quite clearly hirsute (the average number of hair on a human head is about

¹⁸ Typically, *sorites* arguments are formulated in the material mode.

10⁵). Since corresponding arguments can be run at the conceptual level, they would show, if sound, that on Peacocke's view putative soritical concepts are not concepts at all, but notions.

Peacocke is not committed to this result. R. Mark Sainsbury (1990: 259-260), and Kirk Ludwig and Greg Ray (2002: 432-434; n. 20) convincingly argue that the persuasive force of the second step in *sorites* arguments (P2 in the argument schema above) originates in a failure to distinguish between what a rule does not command and what it forbids (Sainsbury) or permits (Ludwig and Ray). On Sainsbury's view, although the rules governing the use of vague predicates do not command us to withhold a vague predicate in a case that differs only marginally from a case to which one has previously applied that predicate, these rules do not forbid us to do so either—they permit us to draw the line that is needed to block the second step in *sorites* arguments (albeit without fixing where the line is to be drawn).¹⁹ According to Ludwig and Ray, the practices governing the use of vague predicates simply do not offer any guidance as to how these predicates are to be used in borderline cases. They write:

Our use of vague predicates is tolerant with respect to small variations along relevant dimensions. The key point we urge is that this should not be misunderstood as a license on the part of our practices to apply the term to any case that differs only incrementally along a relevant dimension of variation if it applies to the given case. This would make the practice incoherent, rather than simply incomplete. It is failure to distinguish the practice giving no guidance from its providing a license to move from one case to an incrementally different one which is the source of much of the confusion about the sorites arguments. (Ludwig and Ray 2002: n. 20)

Note that the two views are consistent. Assuming that the practices governing the use of vague predicates give us no guidance as to how to use these predicates in borderline cases, it might be permissible for us to refrain from using a vague predicate in some borderline case even though we used it previously in a marginally different case (thereby drawing a line by default). Nevertheless, it might not be permissible for us to apply this predicate to all marginally different cases

¹⁹ Presumably, Sainsbury's point applies only to borderline cases. Note that Sainsbury formulates his view in terms of vague concepts. For the sake of expository consistency, I reformulated it in terms of vague predicates.

along this dimension of variation, as doing so would require a positive license to disregard this dimension of variation as irrelevant.²⁰

Ludwig and Ray argue that vague predicates are incomplete. They define semantic incompleteness as a failure of semantic completeness, which they define as follows:

We will say that predicates in a language *L* are *semantically complete* provided that, relative to *L*, (i) a range of application is determined, and (ii) for anything which may fall in the range of application it is determined whether the predicate either applies or fails-to-apply to it. We will say that such predicates have an *extension* in *L*. Other predicates in *L* we call *semantically incomplete*, and we say that they *lack an extension* in *L*. (*Ibid*: 423-424)

Ludwig and Ray propose that our use of vague terms typically involves a pretense that they are precise and, so, semantically complete. *Sorites* arguments are often presented in the material mode, which tempts us to engage in our usual pretense and endorse their first premise. We then maintain this pretense in the face of their second premise and endorse it also, thus committing ourselves to their untoward conclusion. Once we abandon the pretense and realize that, given that the predicate has no extension, neither of the argument's premises is truth-evaluable, the paradox dissolves.

Although Sainsbury might balk at the lesson Ludwig and Ray draw from the incompleteness of vague predicates, namely, that no sentence containing a vague predicate is truth-evaluable,²¹ he agrees that vague predicates (and vague concepts) lack extensions altogether, since extensions, as sets, must have sharp boundaries, while vague predicates (and vague concepts) draw no boundaries. He proposes that vague concepts (and, presumably, vague predicates) do classify, but they do not classify by dividing the items in their range of application into

²⁰ A positive license might be an explicit license to apply some predicate (e.g., 'is a mammal') to the cases at hand by virtue of their instancing some property (i.e., their being animals possessing mammary glands), which might implicitly permit us to disregard incremental variations along various dimensions (e.g., the number of hairs on their head) as irrelevant. A positive license to disregard some dimension of variation need not be explicit.

²¹ Sainsbury proposes that the way forward in describing the logic of vagueness involves taking vague objects as basic (1990: 262), presumably to function as "truth-makers" of sorts for sentences containing vague predicates.

two sets, the set of items to which the concept (or predicate) applies and the set of items to which the concept fails-to-apply. He writes:

Boundaryless concepts tend to come in systems of contraries: opposed pairs like child/adult, hot/cold, weak/strong, true/false, and the more complex systems exemplified by our colour terms. This is a natural upshot of boundarylessness, as we can see by reflecting on what is involved in grasping a concept.

Such a grasp, it must be agreed on all sides, involves knowing how something would have to be for the concept to apply to it, and how something would have to be for the concept not to apply ... Hence it is very natural to see the division between what a concept includes and what it excludes in terms of a boundary. Certainly, perception of a boundary would be enough; but the proponent of boundarylessness will insist that it is not the only way.

On the alternative picture, what a concept excludes is graspable in a positive way, mediated by other contrary concepts. A grasp of red attains grasp of what is not red at a derivative level, via a grasp of yellow, green, blue and so on. A system of such concepts is grasped as a whole, as can be seen in the way paradigms are used in learning. There are paradigms of red, but nothing is non-derivatively classifiable as a paradigm of not-red. Any paradigm of another colour will serve as a paradigm of how not to be red, but only in virtue of its positive classification as another colour. (Sainsbury 1990: 258)

On this view, vague concepts belong to what Peacocke calls local holisms, i.e., sets of concepts whose possession conditions cannot be given independently of the possession conditions of the others (they are, then, sets of concepts that cannot be possessed independently of the others).²²

On the picture Sainsbury proposes, the items in the range of application of, e.g., **child** are classified as **not-child** not by virtue of falling outside a boundary of applicability determined by the deployment conditions of **child**, but by virtue of falling under **adult**. Both concepts permit the drawing of a boundary of applicability between them, but neither mandates the exact location of this boundary, or that such a boundary be drawn.

²² This view is attractive for some vague concepts, but not for others. So, I have beliefs about the Everest. On the view I have been proposing, such beliefs would involve a directly referring concept (call it **Everest**), together with some identifying concepts. On other views, such beliefs might involve the predicative concept, **is Everest**. Whether directly referring or predicative, this concept is plausibly vague, but it does not clearly belong to a system of contrary concepts (although **is a mountain** might).

Setting aside for now the issue of the truth-evaluability of beliefs involving vague concepts, let us ask whether such (putative) concepts can satisfy Peacocke's requirements for concepthood. The answers offered by Sainsbury and by Ludwig and Ray to the *sorites* paradox show that the possession conditions of vague concepts need not be thought to generate inconsistent deployment conditions (even if we take ourselves to possess "classical" logical concepts).²³ Nevertheless, such concepts seem to have incomplete determination theories, in the sense that their deployment conditions fail to specify, for some objects that belong in their range of application, whether the concept applies or fails-to-apply to them. Note that this is so even on Sainsbury's view, interpreted as proposing that vague concepts form local holisms, and according to which the vagueness of the concepts involved in such local holisms is no bar to a complete classification of the objects that belong in their range. On this view, although the deployment conditions of these concepts permit one to classify even borderline cases, they do not mandate how borderline cases are to be classified.

The incompleteness of the determination theories of vague concepts does not sit comfortably with Peacocke's identification of possessing a concept with knowing what it is for something to be its semantic value.²⁴ In virtue of possessing a vague concept, we know what some things

²³ Ludwig and Ray think that our practice of using vague expressions should not be taken to indicate that there is a corresponding vagueness in the concepts and beliefs we express using vague expressions. Given the account of belief content I have been outlining, I cannot agree with them. Their claim also strikes me as highly counterintuitive. When I state, 'Amy bought a green car,' thereby expressing a belief based on Basil's testimony rather than first-hand perceptual experience, isn't the concept **green** involved in my belief just as vague as the expression I use to express it? After all, Amy's car could be of any of the shades to which I normally apply my concept, **green**, including some shades to which I might apply it with some hesitancy—blueish green (or would that be greenish blue?), mustard green (or would that be mustard yellow?). That Basil might count something as blueish green that I would tend to count as greenish blue does not worry me. The expression, 'green' that I use to express my belief, with full sincerity, and the concept, **green**, that I thereby express, both seem equally tolerant of such borderline differences. Upon seeing Amy's car and finding it to be greenish blue by my lights, I might revise my belief and deem her car blue. But it is not clear that I would judge my former belief, or the belief Basil expressed in his testimony, as strictly speaking incorrect—it was a borderline case.

²⁴ This is so especially if one endorses a conceptual analogue of Ludwig and Ray's linguistic thesis, according to which vague expressions neither apply, nor fail-to-apply, to anything. If 'knowing what it is for something to be the semantic value of some concept' is taken to mean 'what something must be like for that concept to apply to it,'

must be like in order for the concept to apply to them, and such knowledge certainly seems required for one to count as possessing the concept. But what seems equally characteristic of (and required for) possessing a vague concept is an *uncertainty* as to whether to apply the concept in borderline cases. As Sainsbury stresses, our vague concepts allow us to resolve our uncertainty and apply/withhold vague concepts to/from borderline cases. In such cases, however, the resolution is not underwritten by knowledge. Rather, it is underwritten by something akin to a decision. One who always classified objects as red just as I do but conceived the continuum of red shades as sharply delimited on both sides and, so, was never unsure (on conceptual grounds) as to whether to count a given shade as red, would not seem to be deploying the same concept as I do, but a precise, and distinct, analogue of my concept. Peacocke's identification of possessing a concept with knowing what it is for something to be its semantic value, then, might need to be revised to allow for the distinctive freedom that the deployment conditions of vague concepts allow in borderline cases.²⁵

Let us now return to the issue of the truth-evaluability of beliefs involving vague concepts. I have been interpreting Peacocke's requirement that every concept have a determination theory as the requirement that the possession conditions of every concept determine consistent deployment conditions for that concept. Peacocke formulates this requirement in different terms, namely:

The determination theory for a given concept (together with the world in empirical cases) assigns semantic values in such a way that the belief-forming practices mentioned in the concept's possession condition are correct. That is, in the case of belief formation, the practices result in true beliefs, and in the case of principles of inference, they result in

nothing is amenable to the relevant knowledge in the case of vague concepts. Such a view, together with Peacocke's identification, entails that (putative) vague concepts cannot be grasped (and, so, can presumably not be possessed).

²⁵ I will not attempt the revision here, as I do not take the issue just discussed to indicate that vague concepts are unlikely to be *bona fide* concepts. Rather, the issue indicates that a more careful account of the relation between the possession conditions of a concept and grasp of this concept is needed.

truth-preserving inferences, when semantic values are assigned in accordance with the determination theory. (Peacocke 1992: 19)

I have argued that *sorites* arguments (reformulated to apply to concepts) need not be taken to show that vague concepts have inconsistent deployment conditions. Peacocke's requirement, however, asks for more, and vague concepts fail to meet it. Consider the case I outlined earlier (n. 23, this chapter). I formed, unobjectionably it seems, the belief that Amy's car is green. Should I see another car of the same shade as Amy's, I would believe that it is blue. Something seems amiss in my belief formation, and it seems traceable to the very freedom one has in deploying **green** in borderline cases—the belief-forming practices mentioned in the concept's possession conditions may not result in true beliefs when the concept is applied to borderline cases (this is so especially if one thinks that vague concepts applied to borderline cases, or to any case whatsoever, cannot yield truth-evaluable beliefs). Should we therefore conclude that (putative) vague concepts are not genuine concepts?

I think that this conclusion would amount to throwing our conceptual baby out with its theoretical bathwater. Arguably, most of our concepts are vague. A requirement that entails that most of our putative concepts are not genuine concepts is probably too strong. Peacocke is right, I think, to require that the belief-forming practices mentioned in a concept's possession condition be correct, in addition to requiring that its possession conditions determine consistent deployment conditions. His requirement is too strong, however, because he interprets correctness in terms of truth for all concepts.

I proposed near the end of last chapter that the proper function of belief is to sanction and preserve information as such in conceptualized form. Beliefs perform this function by way of deploying concepts in given modes of combination. Beliefs deploying vague concepts do sanction and preserve information, albeit not precisely. Ludwig and Ray stress that sentences

containing vague predicates can be used to convey information about the world. They can do so because, as competent speakers of the relevant language,

we have at our disposal whatever we know of what there is to know about the partially determined meaning of a vague predicate. This may amount to such things as (i) knowing what are supposed to be determining factors for the application of the predicate (quantity of hair on the head, for example), and (ii) some ranges of such determining factors which are intended cases for the predicate. (Ludwig and Ray 2002: 442)

Something similar can be said of beliefs whose content includes vague concepts. My belief that Homer is bald carries for me the (putative) information that the number of hairs on Homer's head falls within the deployment range of my concept, **bald**, and this (putative) information is available to me by virtue of my possessing this concept—by virtue of possessing a concept, I have implicit knowledge of its possession conditions, and the possession conditions for **bald** specify the determining factors for its deployment, i.e., the number of hair on one's head, together with a range of its deployment.

My belief does not carry any information concerning the exact number of hairs on Homer's head, but for my purposes it need not. The attempt to form a true belief concerning the exact number of hairs on Homer's head would require much effort, would have a high probability of resulting in a false belief despite my best efforts, and the (putative) information such a belief would carry would most likely be useless to me. Beliefs deploying vague concepts are often economical in terms of the doxastic effort required to acquire them, less prone to error than their precise analogues, and perfectly suited to our purposes. It would seem misguided, then, to count all beliefs deploying vague concepts as incorrect and somehow failed by virtue of deploying vague concepts. If conceptual precision is required for the success of a belief in performing its proper function to be measurable in terms of truth and falsity, then truth and falsity are not an appropriate measure of the success of all beliefs.

Accordingly, I propose to reformulate Peacocke's requirement as follows:

The determination theory for a given concept (together with the world in empirical cases) assigns semantic values in such a way that the belief-forming practices mentioned in the concept's possession condition are correct. That is, in the case of belief formation, the practices result in information-carrying beliefs, and in the case of principles of inference, they result in information-preserving inferences, when semantic values are assigned in accordance with the determination theory.

With regard to precise concepts, the revised requirement is co-extensive with Peacocke's requirement. Not so with regard to vague concepts, as it allows that vague concepts can be involved in the formation of correct beliefs. The revised requirement is also co-extensive with the requirement that concept possession conditions determine consistent deployment conditions. **Tonk**, for instance, fails it, and does so by virtue of having inconsistent deployment conditions. The revised requirement can be understood as explaining why concepts possession would need to determine consistent deployment conditions—failure to do so would render the putative concept incapable of contributing content to information-carrying states. I take these considerations to indicate that the proposed reformulation is appropriate.

However useful, vague concepts are epistemically dangerous and should be handled with care. Just as we tend to operate under the pretense that vague expressions are precise, we tend to operate under the pretense that vague concepts are precise. Although most of our concepts are vague, we tend to assume that our beliefs offer us a clear window onto the world and sometimes come to epistemic and practical grief as a result. In the case I considered earlier, (n. 23, this chapter) had I come upon Amy's car not knowing that it was Amy's car, I would most likely not have recognized it as Amy's car—I would have believed that Amy's car was green but thought that the car I beheld was blue. In most circumstances, my failure to recognize Amy's car on the

basis of my belief would be innocuous; in some circumstances, however, it could be disastrous.²⁶

In matters of import, keeping in mind that most of our concepts are vague can be as useful, I think, as vague concepts are in mundane circumstances.

Belief Possession Conditions vs. Belief Attribution Conditions

Peacocke endorses the view that the possession conditions for a concept may differ from its attribution conditions. His view is that one may truly attribute a concept to someone who fails to satisfy its possession conditions fully because attribution conditions are weaker than possession conditions (Peacocke 1992: 29). If the view that the individuating cognitive content of belief is fully constituted by concepts is correct, the same can be said of beliefs. Although this view seems implausible, there are theoretical reasons for endorsing it. So, consider the two concepts mentioned earlier, **red** and **red***. Suppose someone finds herself in a circumstance of the type described in the first possession condition for the concept **red**, and sincerely utters, ‘this is red.’ Does the belief she thereby expresses deploy the concept, **red**, or the concept, **red***?

The reasons that, e.g., Willard V. O. Quine and Davidson adduce in support of meaning holism and Saul Kripke’s views concerning rule following speak in favor of the view that there is no determinate answer to this question. Yet introspection seems to favor the contrary view. If, upon sincerely uttering the sentence, ‘this is red,’ under circumstances of the type described in the first possession condition for the concept **red**, I was asked whether I meant to express a belief deploying the concept, **red**, or the concept, **red***, the answer would seem readily available to me. Indeed, the view that there is no determinate answer to this question seems preposterous when the belief at issue is my own—the answer, it seems, is known to me immediately;

²⁶ This case is especially flagrant because it involves a vague concept deployed in a borderline case. But vague concepts can lead us astray even when deployed in clear cases, when we treat beliefs involving them as if they were precise while our purposes require precise beliefs.

since I am the author of my belief, my answer to the question (or my disposition so to answer the question) determines the matter and ensures its own truth. The introspective stance on belief content, then, seems to favor the view that a belief attribution is correct just in case the believer satisfies the possession conditions of the complex concept attributed to her and tokens it at the relevant time.

The definition of the individuating cognitive content of belief I offered earlier—namely:

Let b be a particular belief state such that being in b requires tokening concepts c_1-c_n , and let k be the knowledge that anyone who possesses c_1-c_n has just by virtue of possessing those concepts.

Belief b is individuated by the content that p just in case from k alone and the supposition that b is true, it is immediately *a priori* derivable that p .

Any two particular belief states, b_1 and b_2 , that meet the above condition count as being of the same individuating type (i.e., as being the same belief, where ‘belief’ is used to refer to a type of state rather than to a particular state) by virtue of having the same individuating cognitive content.

—lends credence to the side favored by the introspective stance on belief content. On the account I have outlined, having b entails that one satisfies the possession conditions for c_1-c_n . To one who has b and knows that she has b , b carries the information that p (where b carries the information that p by virtue of being constituted by c_1-c_n in a given mode of combination). To one who does not possess c_1-c_n , or does not know that it is b that is being tokened (as a belief attributor might), b might fail to carry this information altogether. I think that the reasons that Quine and Davidson adduce in support of meaning holism need not lead one to endorse the view that the individuating cognitive content of belief is subject to radical indeterminacy, since their evidence can be accounted for by the thesis that environmental and behavioral evidence (including linguistic evidence) radically underdetermines belief attribution.

The problem raised by Kripke is trickier, since, from the attributor’s stance at least, it is hard to say what could count as a suitable ground for giving a definite answer to the question,

e.g., ‘Does little Amy believe that the ball she beholds is red, or does she believe that it is red*?’

Introspection does suggest an answer to this puzzle. The natural reaction to this puzzle when it is raised with regard to one’s own belief seems to be, ‘I determine the content of my belief!’

Introspection, that is, seems to advocate the view that what determines the individuating cognitive content of belief is something akin to a cognitive intention of the believer. Yet this claim seems to entail that every belief is accompanied by a propositional attitude concerning the cognitive content of that belief. If so, given the view endorsed here that the individuating cognitive content of propositional attitudes is constituted by concepts and their mode of combination, this claim seems to entail that believing requires one to have the concept of a concept, concepts of whatever concepts constitute the cognitive content of one’s belief and perhaps also a concept of the relevant mode of concept combination. This result is wildly implausible, since young children and cognitively sophisticated animals seem to have beliefs but do not seem capable of having any such concepts.

Searle’s account of the form of the intentional content of intentions-in-action suggests a way to avoid this untoward consequence (Searle 2001: 44-5).²⁷ On Searle’s account, the intentional content of an intention-in-action does not involve the proposition whose truth the intention is to bring about causally; rather, it directly involves the event (or type of event) that the intention is to cause. This account is plausible, since it is doubtful that we possess the concepts that would be required to have propositional intentions-in-action to perform most of the humdrum motions we carry out, intentionally it seems, as we go about our daily lives. If cognitive intentions are involved in securing a determinate cognitive content for belief, it seems, their intentional content would have to be akin to the intentional content of intentions-in-action.

²⁷ For a brief explanation of what Searle means intentions-in-action (as opposed to prior intentions) by see n. 28, Chapter 2.

However, it is unclear that such cognitive intentions are distinct mental states that include, as do intentions-in-action according to Searle, a self-referential causal element designating the sought effect of one's intention (for cognitive intentions, the sought effect would be concept deployment). Introspection, for one, yields nothing to suggest that believing involves two distinct and causally related stages in cognitive processing, one to secure the cognitive content of belief and the other to secure belief itself. Rather, securing a synchronic/individuating cognitive content seems to be achieved by believing itself—believing that the cat is on the mat, it seems, simply is to believe that a specific object [**a**] falling under the concept **cat** bears the relation conceptualized as **on** to a specific object [**b**] falling under the concept **mat**. Introspection, however, does not yield clear evidence that there are intentions-in-action of the sort described by Searle. The reasoning that leads one to endorse their existence seems rather to be that, since the bodily movements we perform while executing those actions which we intend to perform are introspected as intentional even when not consciously intended, their etiology must involve intentions. Likewise, if it turned out that for cognitive states to have a determinate content, their etiology would have to involve cognitive intentions, one might argue that, since cognitive states are introspected as having a determinate content, their etiology does involve such intentions.²⁸

My purpose here is not to argue for the existence of cognitive intentions, nor to take a stance concerning their nature and relation to cognitive states. My purpose is rather to shift the burden of proof—until conclusive evidence to the effect that cognitive states have no determinate content is offered, the view favored by the introspective stance on belief, namely, that

²⁸ One might worry that the proposal that the individuating content of belief is secured by a cognitive intention-in-action is liable to a regress objection. Intentions are judgment-sensitive attitudes. That is, they are at least capable of being motivated by reasons. The worry is that such reasons would themselves need to have a cognitive content, which would need to be secured by a further cognitive intention. As I noted in n. 32, Chapter 2, this need not be the case. A total reason for an intention-in-action might consist of a non-propositional desire, together with a non-propositional recognition of how best to satisfy this desire in one's circumstances (where the recognition is made possible by a non-propositional know-how that the agent possesses). Concept-possession provides just the required know-how, and one's various interests might provide the required desire. No further intentions are required.

beliefs have determinate contents, is properly taken as one's theoretical starting point. Reflection on the concept of belief provides further reason to side with introspection.

The Concept of Belief

I said above that behavioral evidence radically underdetermines belief attribution. This has been taken to indicate that the individuating cognitive content of belief itself is subject to radical indeterminacy. This conclusion, I think, follows only if one takes **belief** to be an essentially third-personal concept. There are good reasons to think that **belief** is an essentially first-personal concept. According to Peacocke, the possession conditions of **belief** are:

A relational concept **R** is the concept of belief only if

- (F) the thinker finds the first-person content that he stands in **R** to **p** primitively compelling whenever he has the conscious belief that *p*, and he finds it compelling because he has that conscious belief; and
- (T) in judging a thought of the third person form **aRp**, the thinker thereby incurs a commitment to *a*'s being in a state that has the same content-dependent role in making *a* intelligible as the role of his own state of standing in **R** to **p** in making him intelligible, were he to be in that state. (Peacocke 1992: 163-64)

As formulated, (T) strikes me as incomplete, since it does not specify a transition to a thought content containing **R** and, so, does not determine a third-person deployment condition for **belief**.²⁹ (T), then, should be reformulated along the following lines:

- (T*) the thinker is disposed³⁰ to believe the third-person content of the form, **aRp**, when he takes *a* to ..., and (T).

²⁹ Although Peacocke presents (F) and (T) as merely necessary for the possession of **belief**, he offers them as full-fledged possession conditions.

³⁰ Peacocke seems to wish to reserve the phrase, 'primitively compelling' for valid transitions (e.g., the transitions mentioned in the possession conditions of logical concepts and in (F) above). Although he does not explicitly restrict the use of this phrase to such transitions, his own uses seem to abide by such a restriction. This seems to introduce an unnecessary complication in his lexicon. However, it seems appropriate not to use this phrase in (T*). More generally, for experiential concepts (where 'experience' is understood broadly enough to include introspective awareness), it seems appropriate to formulate the clause that mentions experience in terms of a transition that the thinker finds primitively compelling, and to formulate the other clause(s) in terms of transitions that the thinker is disposed to make when s/he takes the primary ground for the experience to obtain. The clause formulated in terms of a transition that the thinker finds primitively compelling would then be primary, while the other(s) would be derivative on the primary clause. By virtue of using such primary and derivative clauses, such a formulation would secure that these possession conditions are the possession conditions of a single concept; by virtue of formulating

The task of completing (T*) is rather tall, however, since (i) behavioral evidence radically underdetermines belief attribution, (ii) (T*) includes a variable (**p**) ranging over attributed belief contents (i.e., complex concepts), and (iii) according to Peacocke's determination requirement, the belief-forming practice mentioned (T*) must be correct. Given (i), (ii) and (iii), the ellipsis marks in (T*) cannot simply be replaced by 'be in a state that has the same content-dependent role in making *a* intelligible as the role of his own state of standing in **R** to **p** in making him intelligible, were he to be in that state.'

I am not sure how to reconcile (i), (ii) and (iii). Fortunately, I need not do so for my purposes here. It is enough to note that, given (F), (T) will have to be included in any suitable version of (T*), either explicitly or by virtue of being entailed by it, since this is required to ensure that (F) and (T*) are the possession conditions of a single concept, which conceptualizes a single type of state. In other words, this is required to ensure that a state conceptualized as **belief**_(T*) when attributed to another on the basis of behavioral evidence is the same type of state as one conceptualized as **belief**_(F) when attributed directly to oneself. Suppose some thinker satisfies (T[#]), where (T[#]) is some possession condition specifying when a thinker is disposed to deploy the relational concept **R**[#] third-personally. Suppose that deploying **R**[#] in accordance with (T[#]) does not incur the commitment mentioned in (T). Suppose also that this thinker satisfies (F). Regardless of whether (T[#]) succeeds in capturing some features of **belief** as deployed third-personally, it does not seem that the relation between **R**[#] and **R** would be identity. In other words, it does not seem that, by virtue of satisfying (F) and (T[#]), this thinker would be satisfying the possession conditions of a single concept. Note that a putative concept

the primary clause in terms of a transition that the thinker finds primitively compelling, and the derivative clause(s) in terms of what the thinker is disposed to believe, such a formulation would capture the phenomenal difference involved in deploying such concepts in experiential and non-experiential circumstances; by virtue of mentioning experience in the primary clause, such a formulation would secure that these possession conditions are the possession conditions of an experiential concept.

individuated by (F) and (T[#]) would fail Peacocke's determination requirement. The deployment conditions for a concept are not relative to the thinker who happens to deploy this concept. Accordingly, for the possession conditions mentioned in (F) and (T[#]) to determine consistent deployment conditions, they must be such that, for any two thinkers, *x* and *y*, and any content, *p*, *x* is in a position appropriately to deploy **R**[#] with respect to *y* and *p* just in case *y* is in a position appropriately to deploy **R** with respect to *y* and *p*. (F) and a suitable version of (T*) that either included, or entailed (T) would satisfy this requirement. (F) and (T[#]) obviously fail it.

The foregoing reflection merely shows that the first-person and the third-person clauses of **belief** must be specified as possession conditions of a relational concept (relating an individual to a content) that the thinker is disposed to deploy only when some ground for deployment obtains, i.e., only when the corresponding relation is instanced. The first-person and the third-person clauses might differ in terms of the thinker's cognitive access to this ground, i.e., in terms of how the instancing of the corresponding relation is available to the thinker's awareness—presumably, introspective awareness in the first-person clause and behavioral evidence in the third-person clause—but they cannot differ by virtue of offering distinct relations as grounds for deployment. Yet this does not show that the possession conditions of belief must designate the first-person cognitive access to the instancing of the relevant property as primary (and, so that the first person clause must be primary, and the third-person clause must be derivative, in the sense specified in n. 30, this chapter). That they must be so specified is overwhelmingly plausible, however, given the well-established asymmetries in first-person and third-person authority over present-tense ascriptions of conscious belief. The foregoing reflection, then, shows that it is overwhelmingly plausible that belief is an essentially first-personal concept.

The phrase, “the thinker thereby incurs a commitment to *a*’s being in a state that has the same content-dependent role in making *a* intelligible as the role of his own state of standing in **R** to *p*” does not warrant Peacocke’s view that belief attribution conditions are weaker than belief possession conditions. The thinker’s own belief states are individuated in part by their individuating cognitive content, which is constituted by concepts in given modes of combination. In the case of conscious belief, this content is immediately accessible to the believer, at least on Peacocke’s view (Peacocke 1992: 158-159). To attribute a state that has the “same *content-dependent* role in making *a* intelligible,” then, is to attribute a state with a definite content, constituted by concepts in given modes of combination, whose possession conditions *a* satisfies and that *a* tokens, and it is to attribute a state to whose content the attributor would have immediate access, were it his/her own state and were it conscious. So much seems to be dictated by the concept of belief.

The appearance that belief attribution conditions are weaker than belief possession conditions is generated by the very difference in cognitive access that the first-person and third-person clauses of **belief** would have to capture—the synchronic and individuating cognitive contents of conscious belief are immediately accessible to the believer, but not to the attributor. But regardless of whether the state conceptualized as a belief is one’s own or another’s, it is a state with a definite conceptual content to which the believer has immediate access (if the belief is conscious). In this respect, **belief** is analogous to **red**, in the sense that both concepts have one clause that specifies the conditions under which one who possesses the concept is disposed to deploy it in the light of experience (where ‘experience’ is understood broadly enough to include introspective awareness), while the other(s) specifies(specify) the conditions under which one who possesses the concept is disposed to deploy it while not having the experience

mentioned in the first clause but taking things to be such that she would have the relevant experience if she were appropriately placed. Such clauses, then, differ not in terms of the property (or range of properties for **red**) or relation that they specify as ground for deploying the relevant concept, but in terms of the thinker's cognitive access to it.

CHAPTER 4 RATIONAL BELIEF¹

‘Rational’ is used here in the broad sense of ‘manifesting rationality as a capacity,’ rather than the narrow sense of ‘fulfilling some relevant set of norms of rationality.’ These two senses are related as follows: only beings that are rational in the broad sense can have states that are either rational or irrational in the narrow sense. Likewise, only states that are rational in the broad sense can be rational or irrational in the narrow sense. This chapter focuses on belief as rational in the broad sense. More specifically, it seeks to answer the question, what is required for a belief to be rational in the broad sense, i.e., to be subject to evaluation as rational or irrational? The question is motivated by the view, implicit in some of the claims I made earlier, that some beings whose states are not typically taken to be evaluable as rational or irrational—to wit, very young children and some cognitively sophisticated animals—are nevertheless capable of having beliefs.² If so, then having the capacity to form a rationally evaluable belief requires more than the mere capacity to form a belief. In what follows, I argue that a belief is evaluable as rational or irrational only if the believer has the capacity to take a deliberative stance toward it. Taking a deliberative stance toward a belief, I argue, minimally requires the capacity to entertain a question and possession of the first-person concept and the concept of belief, in addition to the concepts that constitute the individuating and current synchronic cognitive contents of that belief.

¹ I argued in Chapter 2 that, strictly speaking, evaluations as true or false properly belong to beliefs as representational states, while evaluations as epistemically rational or irrational (and justified or unjustified, warranted or unwarranted, etc.) properly belong to believings as events. This is because evaluations of teleological events along the dimension of success evaluate the outcome of these events, while evaluations of teleological events along the dimension of performance evaluate these events themselves. In this chapter, I am adopting the widespread use of ‘rational’ as a predicate of states. However, I should stress that, given my view that ‘rational’ expresses a teleological property, I am using ‘rational’ in a transferred sense whenever I use it as a predicate of a state.

² This view has been denied, by Davidson for example, on the ground that belief requires rationality (understood in the broad sense specified above). One might be theoretically driven to the view that cognitively sophisticated, but not fully rational, beings cannot have beliefs. Such a restrictive view of belief should not, however, be endorsed as a theoretical starting point.

Rational Evaluability and Control

Intuitively, an individual, or an event or state she undergoes, is properly evaluated as irrational (in the narrow sense) inasmuch as she or it fails to satisfy some norm of rationality that she or it ought to satisfy, and as rational (in the narrow sense) inasmuch as she or it satisfies every norm of rationality that she or it ought to satisfy, given the relevant circumstances. When the evaluation is concerned with a specific type of rationality, e.g., doxastic rationality, the subset of relevant norms is correspondingly circumscribed. An individual, or an event or state she undergoes, that is evaluated as irrational is *ipso facto* evaluated as *prima facie* rationally blameworthy, and an individual, or an event or state she undergoes, that is evaluated as rational is *ipso facto* evaluated as *prima facie* rationally praiseworthy. When the evaluation is concerned with a specific type of rationality, liability to *prima facie* rational praise or blame is relative to those respects concerning which one is evaluated.

Yet it seems unfair to assign normative praise or blame for matters over which one has no control. Liability to normative praise and blame seems to require a dual form of control. First, one cannot rightly be praised or blamed for something for which she is not responsible. If one is praised or blamed for the occurrence of an event, then she must be (at least partly) responsible for causing that event, or she must have been in a position to prevent it. If one is praised or blamed for a property being instanced, then she must be (at least partly) responsible for making it the case that this property is instanced, or she must have been in a position to make it the case that this property is not instanced. So, I can be blamed for reasoning poorly within the confines of my intellectual limitations. But unless I can be shown to be responsible for my intellectual limitations (I might have done too little to overcome them), I cannot properly be blamed for

them. Second, the liability to praise and blame that attends normative³ evaluability seems to require an additional form of control, namely, the capacity to abide by the relevant norm.

Note that abiding by a norm is not tantamount to acting in accordance with it. Consider the following cases. Little Amy knows that she must be home every evening by 6 PM, or else she doesn't get any dessert. She has a watch, knows how to read it, but she doesn't usually bother—she doesn't really care about being home on time, and normally cares more about playing outside than getting dessert. Tonight, however, the aroma of the peach cobbler that bakes in the oven is so deliciously strong that it interrupts her play in the garden and draws her into the kitchen—as luck would have it, by 6 PM. Although Amy gets both dessert and praise, I tend to think that she does not deserve the latter—she did not abide by her rule, but merely acted in accordance with it. Now, consider the following case: Amy knows that peach cobbler—her favorite—is on the menu tonight. For once, she keeps an eye on her watch and makes sure that she gets home by 6 PM. In this case, it seems, Amy deserves not only dessert but also praise—she did abide by her rule (although her motivation might have been somewhat mercenary). Now imagine that Amy's curfew is exceptionally brought forward to 5:30 PM, but no one tells Amy. Once again, highly motivated by the prospect of peach cobbler, Amy makes sure that she is home by 6 PM. In this case, Amy certainly doesn't seem to deserve blame for unknowingly breaking her unusual curfew (and seems to deserve some cobbler, too).

Reflection on these cases indicates that abiding by a rule, in the manner required for normative praise, requires that this rule be recognized as relevant and applied, rather than merely complied with. Conversely, failure to abide by a rule, in the manner required for normative blame, requires that this rule be recognized as relevant and disregarded. Intuitively, then,

³ I am unsure whether all evaluations that properly issue in praise or blame are normative. I may blame someone for not following a singular order to act in a specific way at a given time. In such a case, it is not clear that my order constitutes a norm, which his action fails to meet—norms are general, while my order isn't.

liability to normative praise and blame requires responsibility and the capacity to abide by the norm relative to which one is evaluated. Minimally, the capacity to abide by the norm relative to which one is evaluated requires the capacities to recognize this norm as governing the matter concerning which one is evaluated, to determine whether or not to obey this norm, and either to obey, or not to obey, it as a result of this determination.

In *Reason without Freedom* (2000), Owens argues that the intuitive requirements just outlined, albeit correct with regard to practical normativity, are not pertinent to doxastic normativity. Owens calls the view according to which normativity requires responsibility, and responsibility requires control, the “juridical theory of responsibility” (henceforth JTR). He argues that proponents of JTR are not interested in control in the sense of subjection to the will as such, but rather in the sense of subjection to the will as controlled by higher-order normative judgments, i.e., in reflective control:

We are interested in subjection to the will only in so far as subjection to the will is a form of control, is something which (on the juridical theory) might underwrite the agent’s responsibility for his or her own agency. And the agent exercises control by means of higher order normative judgements about what ought to happen. These judgements are arrived at precisely by applying the relevant norms, so norms do constrain what can happen at will in the sense of that phrase which connotes control and therefore (on the juridical theory) responsibility. (Owens 2000: 80)

This agrees with the intuitions outlined above. On this intuitive view, the liability to normative praise and blame that attends normative evaluability requires not only the rather thin responsibility of authorship (according to which one might be the author of an event or state that satisfies or fails a norm either positively, by bringing this event or state into being, or negatively, by failing to prevent its occurrence). In addition, such liability requires the capacity to determine whether or not to obey the relevant norm, and as a result either obey it, or not. In other words, it requires the capacity to exercise reflective control in light of the relevant norm.

Owen offers the following definition of reflective control as it applies to states:

A state is under reflective control when it can be motivated by higher order judgements about the probative force of the reasons for it. (*Ibid*: 4)

Note that this definition seems to set an overly strong requirement for reflective control, as it is plausible that a state that is regulated, but not motivated, by such higher-order judgments is nevertheless under reflective control. I will return to this issue later, when I will show that Owens' argument against JTR as it applies to doxastic normativity hinges on the strength of this requirement. I will argue that given a proper understanding of reflective control, JTR sets a highly plausible general requirement on normativity. But before I subject Owens' position to critical assessment, let us examine his arguments for JTR as it applies to practical normativity and against JTR as it applies to doxastic normativity.

Owens isolates the following principles as his argumentative target:

Reflective Motivation: if R is a *prima facie* reason to believe that p , reflection on R provides the rational subject with a motive to believe that p ,

[which is implied by the conjunction of:]

Rational Motivation: If a rational subject is aware of something R which provides them with a reason for belief, it also provides them with a motive for belief, [and]

Reflective Rationalisation: If R provides this subject with a reason for belief, the judgement that they have reason R also provides them with a reason for belief. (*Ibid*: 20)

Owens argues that Reflective Motivation fails, and traces this failure to a failure of Reflective Rationalization.

Owens' case for endorsing JTR with regard to practical normativity and for rejecting it with regard to doxastic normativity begins with an asymmetry he notices between the authorities that higher-order normative judgments have over actions, on one hand, and over beliefs, on the other. A judgment about what I ought to do, he rightly claims, has intrinsic authority over my actions (*Ibid*: 104-108). This is not to say that my judgments about what I ought to do cannot

themselves be irrational. I can be confused about what I want, or want most, and I can make poor judgments about what I ought to do to get what I want, thereby manifesting different forms of irrationality. But if I judge that I ought to do something and fail to act in accordance with this judgment, I, and my action, thereby manifest a further form of irrationality, *akrasia*, regardless of whether my judgment was rational or irrational to start with. Owens writes:

My views about what I ought to do have a right to control my actions which is independent of the cogency of those views: my practical judgements have an intrinsic authority over what I do, not merely an authority which derives from the beliefs and desires which should motivate such action. Of course, the practical judgement must make some sense in terms of the agent's perceived needs and interests, otherwise it would be hard to see it as a judgement of the agent at all. But the judgement's authority, its ability to determine the rationality of both intention and action, is not a product of its cogency. (*Ibid*: 104)

By contrast, a judgment as to what I ought to believe has no intrinsic authority over my belief. A judgment as to what I ought to believe has authority over my belief only inasmuch as it is properly motivated. In cases of conflict between belief and judgment as to what to believe, the judgment as to what to believe does not render the conflicting belief irrational. The conflict is evidence of irrationality, but whether this irrationality affects the higher-order normative judgment or the conflicting belief is undetermined. Should the conflict be discovered, the question as to how it should be resolved—either by revising the higher-order normative judgment, or by revising the conflicting belief—is open, and is properly settled only by conclusive grounds to believe (*Ibid*: 108-110).

According to Owens, what underwrites the difference in the authorities that higher-order normative judgments have over actions, on one hand, and over beliefs, on the other, is a difference in motivational power. (Henceforth, I will reserve the terms, 'practical judgment' for higher-order normative judgments pertaining to action, i.e., judgments of the form, I ought to ϕ , and 'doxastic judgment' for higher-order normative judgments pertaining to belief, i.e., judg-

ments of the form, I ought to believe that p .) Unlike practical judgment, Owens claims, doxastic judgment is motivationally inert. Rational motivation for the belief that p requires a conclusive ground to think p true. Our evidence, however, is typically inconclusive (it always is for perceptual beliefs). What determine whether we count our evidence as a conclusive ground for belief are pragmatic factors such as the importance of the issue, the time and resources available, etc. Since such considerations are necessary for the justification of most of our beliefs but reflection on them cannot rationally motivate belief or provide a reason for belief, they must operate beneath the level of reflection. Further, since our evidence for most of our beliefs is inconclusive independently of such factors, reflection on our evidence independently of such factors cannot rationally motivate belief or provide a reason for belief either. Accordingly, Owens concludes that belief is rationally motivated solely by a non-reflective awareness of considerations in favor of thinking p true against a background of non-reflective pragmatic considerations. Doxastic judgments are “an idle wheel in our motivational economy” (Owens 2000: 18). Reflective Motivation fails by way of a failure of Reflective Rationalization. By contrast, neither action nor intention requires conclusive grounds. While rational belief that p requires taking all apparent evidence that not- p as misleading, rational intending to ϕ is compatible with having reasons not to ϕ ; all that is required is that ϕ ing be, on balance, more desirable than not- ϕ ing. Further, since pragmatic factors such as the importance of the issue, the time and resources available, etc., can properly be taken as reasons for action and intention, Owens views the practical realm as subject to reflective control—by reflecting on our reasons we can control what we do, but we cannot control what we believe.

Owens’ argument against Reflective Motivation is liable to the following objection:

Owens is quite right that pragmatic factors determine whether we count our evidence as consti-

tuting a conclusive (or sufficient) ground for belief. But these pragmatic factors are not properly counted as reasons for our beliefs—we would not include them in an answer to the question, ‘Why do you believe that...?’ Accordingly, they are beyond the pale of Owens’ Reflective Motivation, Rational Motivation and Reflective Rationalization, which set requirements on *prima facie* reasons for belief. That these pragmatic factors are not reasons for belief, however, should not lead one to conclude that they must be motivationally inert at the level of doxastic reflection. First, a moment’s introspection will reveal that doxastic reflection is often motivated by pragmatic considerations—especially considerations concerning the significance of the issue. Second, one can endorse Owens’ highly plausible view concerning the role that pragmatic considerations play in belief formation but avoid his conclusion concerning the idleness of doxastic reflection by endorsing the equally plausible view that pragmatic considerations play the same role in doxastic reflection and, so, in the formation of reflected belief. Third, although pragmatic considerations are not reasons for belief, they nevertheless seem subject to evaluation in doxastic reflection (where ‘doxastic reflection’ is used broadly to mean ‘reflection as to what to believe’). So, how important it is for me to get things right, whether I should devote more time and resources to seeking evidence, etc. are all issues on which I may form a considered opinion, ultimately revising my pre-reflective belief as a result. In fact, the capacity to subject such pragmatic considerations to scrutiny seems to afford a significant avenue through which rational beings govern their cognitive economy.

Owens’ view has a price. On his view, practical judgment is based on inconclusive reasons. Consequently, unlike doxastic judgment, practical judgment isn’t idle, but has intrinsic authority over intention and action. According to Owens, however, rational belief requires conclusive grounds. Further, he holds that being in a position to believe something requires

being in a position to claim to know it. So, by his own admission, Owens must conclude that practical judgment, although truth-evaluable, is not a form of belief and is not amenable to knowledge (Owens 2000: 111-112). This is a rather steep price for the view that belief is not amenable to reflective control, which seems to counter ordinary introspective experience as it is.

Reflective Control: Motivating vs. Regulating

Owens is open to a more general objection, namely, that he sets too strong a requirement on reflective control. Consider the following situation: I am sitting in my car, stopped at an intersection where a policeman is directing traffic. Hidden from my view, but not from the policeman's, an elderly pedestrian is slowly crossing the wide avenue. The policeman signals that I should go, and so I do. I round the corner and, suddenly discovering the pedestrian in my path, I nudge him slightly before I manage to stop. In this case, I can be said to have "motive" control over the motions of my car, in the sense that I control whether my car starts or stops, and where it goes. The policeman, by contrast, merely has "regulating" control over the motions of my car, and over traffic generally at that intersection, in the sense that he controls whether cars are allowed to proceed in a given direction. Nevertheless, in this case, it is the policeman, not I, who is responsible for the nudging (assuming that I did not round the corner too fast and that there was nothing I could have done to stop my car sooner once I became aware of the pedestrian). In line with this case, I propose a general distinction between motivating and regulating control.

The analogy between this case and deliberate action is imperfect in many respects. For one thing, the policeman and drivers are all agents who act on reasons, either reflectively or unreflectively—in particular, the drivers at the intersection presumably all have reasons for intending to drive in given directions. For another, unless the policeman does a terrible job, all drivers will eventually get a chance to proceed in their chosen direction, and whether a particular driver is

allowed to proceed depends on the relational properties that this driver bears to the other drivers—i.e., on the position of this driver’s car in one queue relative to the positions of the other cars in the other queues at the intersection—and on the principles that the policeman follows in letting cars go, which should be impartial to drivers. However imperfect, this example nevertheless sheds light on the distinction between motivating and regulating control. Much like a desire, I provide the impetus for the motion of my car in a given direction and, in this case, I am blind to what lies in my path. Much like practical judgment, the policeman can see what lies in my path and has the power to stop me. This case underscores, I think, that regulating control can be quite sufficient (while motivating control can be insufficient) for responsibility.

Owens himself seems to think that regulating control is sufficient for responsibility. With regard to practical rationality, the controlling power he ultimately assigns to practical judgment is regulating rather than motivating:

An important point of clarification: I am not saying that my judgement that I ought to ϕ in itself provides me with a positive reason to ϕ in just the way my first-order beliefs and desires provide me with reasons to ϕ ... If this judgement did provide a reason for ϕ -ing, it would be sensible to ask whether this was an overriding reason to ϕ , or whether it might be outweighed by my first-order reasons for not ϕ -ing, in which case it might (on balance) be rational for me not to ϕ , even though I judge that I ought to ϕ . But this question is bizarre: we have obviously misconstrued the intrinsic authority of practical judgement.

Instead, we must treat my practical judgement as having a power of veto over my actions. My action cannot be rational unless it conforms to my practical judgement, no matter how strong the considerations in the action’s favour ... What fully rational action requires is congruence between the first-order reasons (of which I am non-reflectively aware) and my higher order practical judgement (where such a judgement is made). (*Ibid*: 106)

Owens seems to be guilty of a bait and switch. He argues that doxastic normativity is not subject to reflective control because it fails Reflective Motivation, i.e., the principle according to which any *prima facie* reason, R , provides the believer with a reason for belief only if reflection on R provides the rational subject with a motive for belief. Doxastic normativity fails Reflective Motivation, he argues, because it fails Reflective Rationalization, i.e., the principle according to

which any *prima facie* reason, *R*, provides the believer with a reason for belief only if the judgment that she has *R* also provides her with a reason for belief. *Ergo*, he concludes, doxastic reflection is *motivationally* idle. However, although he claims that practical normativity is subject to reflective control, he nevertheless grants that practical normativity fails the practical analogue of Reflective Rationalization. Given Owens' own argumentative strategy, parity of reasoning seems to commit him the view that practical normativity fails the practical analogue of Reflective Motivation and, so, that practical reflection is *motivationally* idle also.

Yet while Owens holds belief accountable to the definition of reflective control I quoted earlier—namely:

A state is under reflective control when it can be motivated by higher order judgements about the probative force of the reasons for it. (*Ibid*: 4)

—he clearly holds the states and events involved in practical rationality (i.e., intentions and actions) accountable to a weaker definition of reflective control, namely:

A state or event is under reflective control when it can be regulated by higher-order judgments about the probative force of the reasons for it.

The reason why Owens does not seem to think that failure to satisfy the practical analogue of Reflective Rationalization should be taken to indicate that practical reflection is idle is that higher-order normative judgment has intrinsic authority over intention and action. Although Owens ostensibly focuses his case against JTR as it applies to doxastic normativity on Reflective Motivation, authority rather than motivational power seems to be the heart of the issue. In order to answer Owens's claim that doxastic reflection is idle, then, I need to show that the disparity in authority between higher-order normative judgments in the practical and doxastic realms can be explained consistently with doxastic reflective control, understood as regulating.

Regulating Control: Practical vs. Doxastic

According to the highly plausible account of practical rationality Owens proposes, reflective control over action is exercised via reflected practical judgments of the form, I ought to ϕ , which issue in causally efficacious intentions to ϕ . As I noted above, however, the judgment that I ought to ϕ should not be understood as providing a reason (or, presumably, a motive) for intending to ϕ or for ϕ ing, over and above my reasons for judging that I ought to ϕ . Judging that one ought to ϕ , intending to ϕ and ϕ ing are three distinct and presumably causally linked events. Each stage in this process is liable to failure—agents can fail to respond to their reasons appropriately in their judgments, their intentions can fail to accord with their judgments, and their actions can fail to accord with their judgments and their prior intentions. Each failure is attended by a corresponding failure of rationality. Although the etiology of some actions fails to include any practical judgment or prior intention, all actions are subject to rational evaluation by virtue of being subject to reflective control. An action is rational just in case it either is/was, or could be/have been, undertaken by the agent as a result of a failure-free process of reflective control.

Practical judgment and intention each play crucial, albeit distinct, roles in this process. As I noted earlier, Owens rightly argues that practical judgment has intrinsic authority over action, in the sense that no action can be rational unless it conforms to the agent's practical judgment, regardless of the rational status of that practical judgment and of the strength of the reasons favoring that action. This is because the role of practical judgment is to exert regulating control over action (via intention). Even when irrational, practical judgment retains its regulating function and, so, its authority over intention and action. The role of intention, by contrast, is to “settle the motivational question”:

[T]o intend to enter the hospital at [a] time is to be in a state which, so long as I am in it, guarantees that I will (at least try to) go. Failure to (try to) ϕ when one intended to ϕ and now believes that the time for ϕ -ing has arrived means that one must have abandoned the intention to ϕ ; at some point, one ceased to be set on ϕ -ing ... The will commits the agent in a way that his desires and judgements do not: the will is an executive phenomenon. (*Ibid*: 78-79)

Clearly, the case of belief is different. Belief that p is not typically—or ever, if doxastic involuntarism is true (and I believe that it is)—directly generated by the judgment that one ought to believe that p and resultant intention to believe that p . Although believers can form such a doxastic judgment and intention, and can seek to implement them by undertaking actions aimed at causing themselves to believe that p , they cannot typically⁴ thereby secure a rational belief that p . Further, Owens' case to the effect that doxastic judgments do not have intrinsic authority over belief strikes me as unimpeachable. Accordingly, if we have reflective control over our beliefs, the process through which this control is exercised will have to differ significantly from its practical counterpart.

Most of our actions, and most of our beliefs, are unreflective. They are simply responses to reasons—either for action, or for belief—of which we are unreflectively aware and which we process unreflectively against a background of non-reflective pragmatic considerations (to borrow Owens' phraseology). Reflection is often motivated by pragmatic considerations, especially a consideration of the significance of the matter, and initiated by the questions, 'What should I do/believe?' or 'Should I ϕ /believe that p ?' although these questions need not be enter-

⁴ The qualifier, 'typically' is meant to allow for my uncertainty with regard to the following type of case: Suppose I have a false and seemingly irrational belief, say, that my friend is a governmental spy who means me harm, which I believe to be both false and irrational but can't seem to relinquish. I seek the help of a hypnotist, thanks to whose services I revise my belief so that it becomes both true and no longer irrational in its former way. If my revised belief is grounded in evidence at all, however, it is so grounded in an unusual manner, mediated by hypnosis. Should my friend become a governmental spy who means me harm, I might be incapable of revising my belief accordingly. For this reason, it is not clear to me that my revised belief should be counted as rational. Nevertheless, I am doing my doxastic best under the circumstances, and hardly seem to deserve blame for my imperfect belief. So it is not clear to me that my revised belief should be counted as irrational either. Perhaps, given that my original belief and subsequent revised beliefs are not amenable to revision as a result of reflection alone, they should not be counted as liable to rational evaluation at all, but as pathological.

tained consciously. More generally, what is characteristic of reflection is that it involves an inquiring frame of mind. Reflection is a conscious, purposive activity that is aimed at the rational resolution of one's uncertainty as to how to act or whether to act in a given manner, or as to what to believe or whether to believe a given proposition, and which is concluded either once one's uncertainty is resolved, or once one's preoccupation with one's uncertainty vanishes.

In the practical realm, the question, 'What should I do [in some set of circumstances]?' is properly answered by considering one's relevant desires, values, etc. (i.e., the various psychological states that are usually subsumed under the term, 'pro-attitude'), by ranking them by order of importance or preference if necessary, then by determining what course of action has the best chance of satisfying or optimizing these (ranked) pro-attitudes, given the circumstances. ' ϕ ' in the judgment schema, 'I ought to ϕ ' would stand for the course of action deemed to have the best chance of satisfying or optimizing one's (ranked) pro-attitudes at the close of an idealized process of practical reflection. The less open-ended question, 'Should I ϕ ?' is properly answered along similar lines. The role of intention, whose characteristic form is 'I will ϕ ,' is to ensure that one's practical judgment is translated into action.

In the doxastic realm, by contrast, the question, 'What should I believe [concerning some topic]?' is properly answered by considering evidence concerning the topic itself. Note, however, that in this respect, doxastic reflection and practical reflection are exactly alike. Reflection is concerned with one's reasons. In the practical realm, one's reasons are one's pro-attitudes and one's means-end beliefs concerning how to satisfy or optimize these pro-attitudes. In the doxastic realm, one's reasons are one's evidence. What is distinctive of reflection in the practical realm is that the agent does not merely respond to her reasons, but attends to them—she might rank her pro-attitudes, she seeks to discover the best way to satisfy or optimize them, she might

determine that the satisfaction of some of her pro-attitudes will have to be forgone for the sake of satisfying others, etc. Even if practical judgment cannot properly be said to constitute a new reason or motivation to act, practical reflection certainly can make a vast difference to what one does. Likewise, what is distinctive of reflection in the doxastic realm is that the believer does not merely respond to her evidence, but attends to it. In a process of doxastic reflection, a believer might subject evidence that would unreflectively have moved her to belief to scrutiny; she might query her grounds for accepting it as evidence and discount some of it as misleading or irrelevant, she might come to view distinct evidence as relevant, and ultimately determine that she needs to seek additional evidence prior to forming a belief on the matter. Reflection on her evidence might also lead her to maintain her unreflective belief, or to endorse a belief that she would also have endorsed unreflectively. Nevertheless, as the previous case illustrates, doxastic reflection can make a difference as to what, or whether, one believes. Even when doxastic reflection issues in a belief that was, or would have been endorsed unreflectively, it makes a difference to how one believes—doxastic reflection issues in reflected belief.

The conclusion of practical reflection is a practical judgment. When the judgment counsels acting rather than refraining from acting and when no breach of rationality ensues, the practical content of this judgment provides the practical content of an intention to act (i.e., of a commitment of the will to undertaking a (type of) action). By contrast, the conclusion of doxastic reflection need not be a doxastic judgment. Rather, the conclusion of doxastic reflection is typically belief or disbelief, which are types of doxastic commitment, or abstention from belief and disbelief. Doxastic reflection can, but need not, issue in a doxastic judgment, which can conflict with the corresponding first-order belief. Thus, the first asymmetry between practical and doxastic reflection that requires explaining is that the former must conclude in a second-

order normative judgment, while the latter need not. As we shall see, the obvious explanation of this asymmetry will also provide an explanation—consistent with doxastic regulating control—of the asymmetry in authority between practical and doxastic judgment that misleads Owens to conclude that doxastic reflection is idle.

I said earlier that reflection is a conscious, purposive activity aimed at the rational resolution of one's uncertainty as to what to do or believe. Intuitively, rationally to resolve one's uncertainty as to what to do is to be in a position rationally to commit to a course of action, i.e., to form an intention to act in accordance with one's reflected practical judgment. Rationally to resolve one's uncertainty as to what to believe is to be in a position rationally to form a doxastic commitment, i.e., to form a reflected belief. Further, *ceteris paribus*, to fail to intend in accordance with one's reflected practical judgment is a failure of rationality, i.e., *akrasia*. Likewise, *ceteris paribus*, to fail to believe while being in a position to believe at the close of a process of reflection is a failure of rationality. Generally, then, rational reflection aims to generate a rational commitment. It does so by focusing one's attention on one's reasons and by processing them consciously in accordance with the norms of rationality (in the manner adumbrated earlier). Uncertainty, however, is a doxastic issue. Rationally to resolve one's uncertainty on any matter is to be in a position rationally to form a doxastic commitment on that matter, i.e., rationally to form a belief. As I just noted, *ceteris paribus*, to fail to believe while being in a position to believe at the close of a process of reflection is a failure of rationality. Rational reflection on any matter, then, aims to generate a doxastic commitment, i.e., a belief.

In the practical realm, the doxastic commitment that seals the rational resolution of one's uncertainty as to what to do, i.e., the belief that ϕ ing is best suited to satisfy or optimize one's relevant pro-attitudes in the circumstances one faces (henceforth, the best means), and the

practical commitment that practical reflection is ultimately meant to generate, i.e., the intention to ϕ , are distinct. Given that practical reflection is concerned with resolving one's uncertainty as to what to do, the reasons with which it is primarily concerned are one's relevant pro-attitudes, which have a world-to-mind direction of fit. It is concerned derivatively with determining the means to satisfy or optimize these pro-attitudes in the circumstances one faces, i.e., with means-end beliefs, which have a mind-to-world direction of fit. One's uncertainty as to what to do is rationally resolved once one discovers the best means. The belief that ϕ ing is the best means expresses one's doxastic commitment to the result of one's reflection. But to be rationally committed to the view that some course of action is superior to all alternatives is not yet to be committed to the view that pursuing this course of action is rationally required. For this reason, practical reflection requires practical judgment. ϕ , the means believed to be the best at the close of the reflective process, provides the practical content of the judgment, I ought to ϕ , which has a mind-to-word direction of fit and expresses one's doxastic commitment to the view that ϕ ing is rationally required. ϕ , in turn, provides the practical content of the intention, I will ϕ , which has a world-to-mind direction of fit and expresses one's practical commitment to implementing ϕ in action.

The role of practical reflection, then, is to take us from pro-attitudes, which have a world-to-mind direction of fit, via a means-end belief and a practical judgment, which have a mind-to-world direction of fit and express doxastic commitments, to the practical commitment embodied in intention, which has a world-to-mind direction of fit. Practical reflection requires practical judgment by virtue of requiring a bridge between the doxastic commitment expressed by the belief that ϕ is the best means, which evaluates ϕ but does not counsel that ϕ be undertaken, and the intention to ϕ . Practical judgment has intrinsic authority over intention and action because it

expresses the mind's reflected view concerning what action ought to be undertaken so that the world accords with one's pro-attitudes (as closely as possible). Failure to commit in one's intention to the course of action recommended by such a judgment is irrational because one would thereby be violating one's doxastic commitments (to this course of action being the best and to it being rationally required). This is so even if these doxastic commitments were undertaken irrationally, i.e., if they were undertaken in the course of a process of practical reflection that was rationally flawed.

In the doxastic realm, by contrast, the commitment that seals the rational resolution of one's uncertainty as to what to believe and the commitment that doxastic reflection is ultimately meant to generate are one and the same. Accordingly, doxastic reflection does not require a higher-order normative judgment to function as a bridge between distinct types of commitment. Given that doxastic reflection is concerned with resolving one's uncertainty as to what to believe, the reasons with which it is concerned are one's reasons to believe. These are reasons to think true that, accordingly, have a mind-to-world direction of fit. In the course of doxastic reflection, reasons for one's reasons might be sought, putative reasons might be dismissed, alternative reasons might be discovered, etc. Pragmatic considerations, such as the significance of the issue, the time and resources available, etc. might operate in the background of one's reflection. *Pace* Owens, however, they can be subjected to reflective evaluation. So, e.g., in the process of one's reflection, one might come to see the issue as doxastically fundamental and, consequently, reevaluate one's view of its significance and determine to devote more time and resources to reflecting on the matter. Whatever course doxastic reflection takes, its role is to take one from reasons to believe, which have a mind-to-world direction of fit, to a doxastic commitment that has a mind-to-world direction of fit (provided that reflection discovers that one's

reasons warrant such a commitment⁵). All-out⁶ doxastic commitment is embodied in belief. But doxastic reflection has also achieved its purpose of rationally resolving one's uncertainty as to what to believe once one is in a position rationally to form a belief, and the corresponding belief expresses one's doxastic commitment to the result of one's reflection. Just as my forming the reflected belief that ϕ ing is the best means does not require in practical reflection the judgment that I ought to believe that ϕ ing is the best means, so my forming the reflected belief that p does not require in doxastic reflection the judgment that I ought to believe that p . In this respect also, doxastic reflection and practical reflection are exactly alike.

Nevertheless, doxastic reflection can issue in a doxastic judgment of the form, I ought to believe that p . However, although such a judgment expresses a doxastic commitment to the belief that p being doxastically required, it does not express a doxastic commitment to p being true. *Ceteris paribus*, reflection that does not issue in the doxastic commitment embodied in a reflected belief (when reflection discovers that one's reasons warrant such a commitment) is thereby incomplete, and thereby manifests irrationality (regardless of whether the reflection issues in a higher-order doxastic judgment). Doxastic reflection that does not issue in reflected belief (when reflection discovers that one's reasons warrant such a commitment) is incomplete and irrational in precisely the same respect in which practical reflection that does not issue in the reflected belief that ϕ is the best means (when reflection discovers that one's reasons warrant such a commitment) is incomplete and irrational—in both cases, the uncertainty that reflection set out to resolve is not properly resolved (intuitively, proper resolution of uncertainty requires

⁵ Practical reflection does not require an analogous qualification. Reflection can issue in the judgment that one ought not to act, followed by the corresponding intention (i.e., the corresponding practical commitment) not to act. By contrast, refraining from believing and disbelieving as a result of reflection does not constitute a doxastic commitment.

⁶ The qualifier 'all-out' means to allow for other types of state, such as conjecture, which embody less-than-complete commitments with a mind-to-world direction of fit.

belief). Typically, judgments of the form, I ought to believe that p , are formed in the face of an inability to believe that p , or in the face of a recalcitrant belief that not- p . Owens is right that, in such a case, the failure to believe that p , or the belief that not- p , is not irrational by virtue of failing to accord with the higher-order judgment. The rational status of not believing that p and of believing that not- p is determined by the relevant reasons. Doxastic judgments of the form, I ought to believe that p have no intrinsic authority over belief, reflected or unreflective, or over failure to believe, and this is because such judgments are not required to generate the rational commitment that doxastic reflection is meant to generate.⁷

If dispensability in reflection is the mark of “an idle wheel in our motivational economy” (Owens 2000: 18), then Owens is right—doxastic judgments are idle wheels in our motivational economy. However, we need to distinguish the claim that doxastic judgments are motivationally idle from the claim that doxastic reflection is idle *tout court*. Owens offers an argument to the effect that doxastic reflection is motivationally impotent. We saw that parity of reasoning commits him to the conclusion that practical reflection is similarly motivationally impotent. We saw that he grants practical reflection regulating power. We determined that doxastic reflection has similar regulating power. We noticed, however, that given the nature of the states and reasons involved, practical and doxastic reflection are radically different processes. In practical reflection, practical judgment plays an indispensable bridging function between the doxastic commitment expressed by the belief that some (type of) action is the best means and the practical commitment expressed by the intention so to act. In doxastic reflection, by contrast, no such

⁷ Failure to judge in accordance with one’s considered practical belief that ϕ is the best means seems *akratic* in the same way that failure to intend in accordance with one’s considered practical judgment is *akratic*. So, if I fail to judge that I ought to ϕ despite my considered belief that ϕ ing is the best means to optimize my desires and values, say because of guilt or some irrational fear, my failure seems to manifest weakness of will just as it would if I had failed to intend to ϕ upon judging that I ought to ϕ because of guilt or some irrational fear. In one case I lack the strength to judge, while in the other I lack the strength to intend. The difference seems immaterial. This indicates that intrinsic authority is a matter of being required to generate reflected rational commitment.

bridge is required. This provided us with an explanation of the asymmetry Owens notices between the authorities that higher-order normative judgments have over actions, on one hand, and over beliefs, on the other. Nothing in this explanation indicates any difference in the regulating power that practical and doxastic reflection exert over the states and events they regulate. Rather, close examination of practical and doxastic reflection reveals that their regulating power is exercised differently by virtue of being exercised over different types of events and states. Regulating reflective control over action is exercised indirectly, via reflected means-end beliefs and practical judgments, which issue in corresponding causally efficacious intentions (although it is exercised directly over reflected means-end beliefs and practical judgments). By contrast, regulating reflective control over reflected belief is exercised directly.

Rationally Evaluable Belief: Minimal Requirements

Owens briefly considers regulating control as securing the control required for doxastic responsibility according to JTR, and rejects it on the following grounds:

Surely rationality does not require continuous reflective intervention; rather it implies the capacity to control what happens by means of reflection. But what exactly does this control, when exercised, involve? Sometimes we exercise control by endorsing (or rejecting) beliefs we have already formed, or decisions we have already taken. Korsgaard calls belief or action 'rational' when reflection would have endorsed it. But such retrospective endorsement can't be the only form reflective control takes. For retrospective judgement to be a form of control at all, it must have some effect on whether we hold the belief (or implement the decision) in question. And if we can influence our beliefs in this way, surely this is because we have a general capacity for epistemic control, one which also enables us to determine our beliefs in prospect: control can't be purely retrospective. (*Ibid*: 11)

The foregoing discussion has shown that the regulating control exercised in doxastic reflection need not be purely retrospective—we can form a reflected belief prospectively, as a result of reflecting on what to believe concerning a given topic, just as we can endorse a reflected belief retrospectively, as a result of reflecting on whether to maintain a belief (in this respect, doxastic reflection functions just like practical reflection—we can form a reflected intention prospec-

tively, as a result of reflecting on what to do in a given circumstance, just as we can form a reflected intention retrospectively, as a result of reflecting on whether to perform some action which we unreflectively intended to perform). Beliefs, I claim, are liable to evaluation as rational or irrational by virtue of being subject to reflective control by the believer, i.e., by virtue of being subject to regulating control by the believer in doxastic reflection in the manner adumbrated earlier (just as actions are liable to evaluation as rational or irrational by virtue of being subject to reflective control by the agent, i.e., by virtue of being subject to regulating control by the agent in practical reflection in the manner adumbrated earlier). A belief formed by a being incapable of doxastic reflection is not rational in the broad sense, i.e., is not evaluable as rational or irrational.

I noted earlier, however, that most of our beliefs (like most of our actions) are unreflective. When we evaluate the rationality of a reflected belief, we evaluate the rationality of the reflective process that generated or maintained it (when we are privy to this process). When we evaluate the rationality of an unreflective belief, or when we evaluate the rationality of a reflected belief whose rational etiology we ignore, we ask whether the believer could have been maintained this belief in the course of a flawless process of rational reflection. So posed, the question is retrospective; but the capacity for doxastic reflection this question presupposes certainly isn't purely retrospective.

Nevertheless, the retrospective question is a good starting-point for determining what the capacity for doxastic reflection minimally requires (and, so, for determining what rational belief in the broad sense minimally requires). Suppose I am deploying my capacity to reflect on my beliefs retrospectively, wondering whether I am right to believe that Basil was assaulted by bears. To do so, of course, I need to possess the concepts that constitute the individuating

cognitive content of my belief that Basil was assaulted by bears, as well as any identifying concepts currently comprised in its synchronic cognitive content, and I need to be capable of tokening this belief. These, however, are conditions on belief possession rather than conditions on doxastic reflection as such. In order to reflect on this belief, I need to be able to represent it as my belief in thought. Accordingly, I need to possess the first-person concept.⁸ Further, I need to be able to represent this belief as a belief in thought. Accordingly, I need to possess the concept of belief. Finally, it seems that I would need to be able to represent this belief as right, correct, or true, in thought. Accordingly, it seems that I would I need to possess the notion of truth. I will provisionally assume that possession of the notion of truth is minimally required for doxastic reflection. I will return to this issue at the end of the next section.

Whether doxastic reflection also minimally requires possession of the concept of a reason (and, so, of the concept of the support relation) is not clear. Reflection does not seem to require one to represent one's reasons as reasons in thought in all cases. Suppose I have recently suffered some episodes of perceptual delusion, which have characteristically been less vivid than genuine perception. As I behold Basil being assaulted by bears, I wonder whether I am right to believe that Basil is assaulted by bears. In such a case, I might simply take my perception as a sufficient reason for my belief, given its vividness, without conceptualizing it as a sufficient reason for my belief (or as a vivid perception). Arguably, then, possessing the concepts of a reason and of the support relation might not belong to the minimal requirements for doxastic reflection. Nevertheless, failure to possess these concepts would severely restrict one's capacity for doxastic reflection. Plausibly, then, possessing the concepts of a reason and of the support relation is minimally required for a full-blown capacity for doxastic reflection, though not for a

⁸ It might be thought that I would need, in addition, to possess the concept of ownership, so as to represent this belief as mine. It will become clear later that the capacity for doxastic reflection does not minimally require possession of the concept of ownership.

minimal capacity for doxastic reflection.⁹ Accordingly, the capacity for doxastic reflection and, so, the capacity to form a rationally evaluable belief, minimally require possession of the first-person concept and the concept of belief and, provisionally, the notion of truth. In addition, the capacity for reflection generally requires the capacity to entertain a question.

I stated earlier that the liability to *prima facie* normative praise and blame that attends the liability to evaluation as rational or irrational requires the capacity to abide by any norm relative to which one is evaluated. This, I argued, minimally requires the capacities to recognize any such norm as governing the matter concerning which one is evaluated, to determine whether or not to obey this norm, and either to obey, or not to obey, it as a result of this determination. To show that that the minimal requirements for doxastic reflection set out above are sufficient to secure the control intuitively required for rational evaluability, I need to show that satisfaction of these minimal requirements suffices to secure the relevant capacities.

By virtue of possessing the concepts that constitute the individuating cognitive content of her belief (as well as any identifying concepts currently comprised in its synchronic cognitive content, when applicable), and by virtue of possessing the concepts minimally required by the capacity for doxastic reflection—i.e., the first-person concept and the concept of belief and, provisionally, the notion of truth—the believer has the capacity to recognize the relevant set of doxastic norms as governing her belief. On Peacocke’s account of concepts, which I endorsed with some modifications in Chapter 3, grasp of a concept entails implicit knowledge of its deployment conditions (or, as Peacocke would put it, of what it is for something to be its semantic value). On the view I have outlined, the individuating and synchronic cognitive contents of a

⁹ When assessing the rational status of a given belief, we make allowances for such conceptual deficiencies. So, we employ less demanding standards when evaluating the beliefs of children than we do when evaluating the beliefs of fully doxastically proficient adults, presumably because children are likely to lack certain concepts that are required for a full-blown capacity for doxastic reflection.

belief are complex concepts. By virtue of possessing the relevant belief, by virtue of being capable of conceptualizing it as her belief, and by virtue of being capable of wondering whether it is true, then, a believer has the capacity to recognize that a given set of doxastic norm governs her belief, in the sense of having the capacity to recognize what it would take for her belief to be true and to recognize her belief as beholden to these truth conditions.

I stated earlier that reflection is a conscious, purposive activity. The capacity for doxastic reflection, then, secures not only the capacity to recognize the norms that govern one's beliefs as governing it, but also the capacities to determine whether or not to obey these norms and either to obey, or not to obey, them as a result. This is not to say that reflected belief requires an intention to believe. I argued earlier that doxastic reflection plays a regulating, rather than a motivating, role in belief formation, and that doxastic reflection does not typically issue in doxastic judgments and corresponding doxastic intentions. When all goes well, doxastic reflection results in rational, reflected belief, formed without any intentional intervention on the part of the believer, at the close of a flawless process of reflecting on her reasons to believe. Nevertheless, doxastic reflection is typically engaged intentionally, and typically involves many intentional steps (in addition to many purely doxastic steps). Sometimes, desire and intention can pervert the reflective process, resulting, e.g., in wishful belief. So, a believer might reflect on her belief with the purpose of comforting it, rather than determining whether she should maintain it. In such a case, although (implicitly) recognizing the norms that govern her belief as governing it, she might actively seek reasons to maintain her belief while refusing to seek, consider or admit any reason to abandon it. In such a case, the believer can be said to have (implicitly) determined not to obey the norms that govern her belief and to fail to obey them as a result—although (implicitly) recognizing what it would take for her belief to be true and (implicitly) recognizing

her belief as beholden to these truth conditions, she has (implicitly) determined that her belief would not abide by these truth conditions, should they fail to be satisfied. Since belief represents its content as true, wishful belief often involves a fair amount of self-deception—the decision to flout the norms that govern her belief isn't typically conscious, and the believer usually tries to convince herself that the world accords with her doxastic desires. Still, the capacity for doxastic reflection secures the (irrational) option to refrain from abiding by the norms governing one's beliefs by making it possible to subject these beliefs to intentional governance—an option seemingly unavailable to beings not blessed with the capacity to form beliefs that are rational in the broad sense.¹⁰

Summing up: I have argued that rational belief, in the broad sense, minimally requires the capacity for doxastic reflection, i.e., possession of the first-person concept and the concept of belief and, provisionally, the notion of truth, as well as the capacity to entertain a question. By virtue of possessing the concepts that constitute the individuating and current synchronic cognitive contents of her belief, by virtue of being capable of conceptualizing it as her own belief, and by virtue of being capable of wondering whether it is true, a believer has the capacity to recognize that a given set of doxastic norms governs her belief, in the sense of having the capacity to recognize what it would take for her belief to be true and to recognize her belief as beholden to these truth conditions. By virtue of being a conscious, purposive activity, doxastic reflection also secures the capacity to refrain from abiding by these norms by making it possible to subject belief to intentional governance. Accordingly, satisfaction of the minimal requirements for the capacity for doxastic reflection secures the control intuitively required for rational evaluability.

¹⁰ Note that the claim isn't that wishful belief can only be formed in a flawed process of reflection. Wishful belief, like any other type of irrational belief, can be formed unreflectively. Still, I suspect that wishful belief requires the same capacity to distance oneself from one's belief, by way of taking a second-order stance on one's beliefs, that reflection requires—on the account of belief content I have been proposing, a being capable merely of first-order belief might be capable of error and illusion, but not of wishful belief.

The Deliberative Stance

In Chapter 2, I briefly discussed the account of transparency Moran proposes as a general requirement for rational agency.¹¹ Moran endorses Edgley's and Evans' view that one's own thinking is transparent to oneself. With regard to belief, Moran characterizes the relation of transparency as follows:

[R]ather than reducibility or indistinguishability, the relation of transparency [Edgley and Evans] are pointing toward concerns a claim about *how* a set of questions is to be answered, what sorts of reasons are to be taken as relevant. The claim, then, is that a first-person present-tense question about one's belief is answered by reference to (or consideration of) the same reasons that would justify an answer to the corresponding question about the world. (Moran 2001: 62; italics in the original)

That is, the question, 'Do I believe that *p*? (and the question, 'Should I believe that *p*?) normally is, and should be, answered by considering (or as if answering) the question, 'Is it the case that *p*?' Moran notes that the question, 'Do I believe that *p*?' cannot, at times, be answered in this manner and must instead be answered by answering (or as if answering) the corresponding, third-person question, i.e., in light of behavioral evidence, testimony, etc. Yet he stresses that such cases are pathological. Transparency, then, is a relation one bears to one's own, non-pathological beliefs (and other judgment-sensitive attitudes¹²).

To distinguish between the former and the latter way of answering (and of posing) the question, Moran introduces the terms, 'deliberative' and 'theoretical':

In characterizing two sorts of questions one may direct toward one's own state of mind, the term 'deliberative' is best seen at this point in contrast to 'theoretical,' the primary point being to mark the difference between that inquiry which terminates in a true description of

¹¹ Some of the outline of Moran's view that follows more or less repeats what was said in Chapter 2.

¹² The statement that belief is a judgment-sensitive attitude does not contradict my earlier agreement with Owens that the doxastic judgment, I ought to believe that *p*, has no intrinsic authority over the belief that *p*. Belief is a judgment-sensitive attitude by virtue of being sensitive to judgments of the form, (it is the case that) *p*. Desire is a judgment-sensitive attitude by virtue of being sensitive to judgments of the form, (it is the case that) *F*ness, or *φ*ing, or whatnot, is desirable. Arguably, intention is a judgment-sensitive attitude by virtue of being sensitive to judgments of the form, I ought to *φ*. Different types of attitudes are judgment sensitive by virtue of being sensitive to different types of judgment. Belief is not judgment sensitive by virtue of being sensitive to judgments of the form, I ought to believe that *p*.

my state, and one which terminates in the formation or endorsement of an attitude.
(*Ibid*: 63)

When asking, ‘Do I believe that p ?’ as a theoretical question, the believer’s stance toward her own belief is just like that of any outside observer. Proper evidence for her answer is evidence that she believes that p , and the answer, ‘I believe that p , but p is false’ does not involve her in any incoherence. When asking, ‘Do I believe that p ?’ as a deliberative question, by contrast, the agent’s stance toward her own belief is unlike that of any outside observer. The deliberative question is equivalent to the question, ‘Should I believe that p ?’ and does not aim merely to yield an accurate attitude report; it aims to yield belief—i.e., an attitude that conforms to what Moran calls ‘the condition of transparency,’ i.e., for belief, the condition is “that I can report on my *belief* about X by considering (nothing but) X itself” (*Ibid*: 84; italics in the original). In the deliberative stance, the agent’s inquiry is not properly focused on evidence that she believes that p ; it is properly focused on evidence that p , that is, on reasons to believe that p .

Moran argues that the capacity to take a deliberative stance toward one’s own beliefs and other judgment-sensitive attitudes is a prerequisite of rational agency. Agents are responsible for their attitudes inasmuch as they have the capacity to avow them, where an avowal is a statement of one’s attitude that conforms to the condition of transparency (*Ibid*: 101). Avowability also sets a normative requirement on rational agency. Inasmuch as they have the capacity to avow their beliefs and other judgment-sensitive attitudes, rational agents are committed to satisfying the condition of transparency in their beliefs and other judgment-sensitive attitudes. For belief, the commitment is to have the belief that p only if one’s report, ‘I believe that p ’ would be a proper answer to the deliberative question, ‘Do I believe that p ?’

The account of rational belief (in the broad sense) I have proposed supports and explains Moran’s view as it pertains to belief. That is, it supports and explains Moran’s claims that, in

non-pathological cases, rational believers should answer the question ‘Do I believe that p ?’ as if answering the question, ‘Should I believe that p ?’

Let us assume that the following cases are non-pathological. Suppose that an individual (let us call her ‘ S ’) who satisfies the minimal requirements for rational belief has the conscious belief that p . Suppose S asks herself, ‘Do I believe that p ?’ In asking this question, she is deploying her concept of belief (**belief**) and her first-person concept (**i**) to **p** , the synchronic cognitive content of her belief that p , in an inquiring mode. Now, we saw that the first-person clause of **belief** is something along the lines of:

- (F) the thinker finds the first-person content that he stands in **R** to **p** primitively compelling whenever he has the conscious belief that p , and he finds it compelling because he has that conscious belief. (Peacocke 1992: 163-164)

We also saw that the third-person clause of **belief** is something along the lines of:

- (T*) the thinker is disposed to believe the third-person content of the form, **aRp** , when he takes a to ..., and
- (T) in judging a thought of the third person form **aRp** , the thinker thereby incurs a commitment to a ’s being in a state that has the same content-dependent role in making a intelligible as the role of his own state of standing in **R** to **p** in making him intelligible, were he to be in that state. (*Ibid*: 164)

We should also note the first-person analogue of (T), i.e.,

- (F*) in judging a thought of the first person form **iRp** , the thinker thereby incurs a commitment to being in a state that has the same content-dependent role in making him intelligible as the role of a ’s state of standing in **R** to **p** in making a intelligible, were a to be in that state.¹³

As I stressed in my discussion of **belief** in Chapter 3, although (F) and (T*) differ by virtue of specifying a distinct cognitive access to the relation that constitutes the ground for deploying **belief**, they cannot differ with regard to the identity of this ground. Accordingly, any first-

¹³ (T*) and (F*) are not part of Peacocke’s original formulation. Note that I do not take these possession conditions to formulate individuating conditions for **belief**. Rather, I take them to adumbrate such individuating conditions, not only because (T*) includes ellipsis marks, but also because (F), (T), and (F*) are likely to require refinement; accurate possession and individuating conditions for concepts are difficult to formulate.

personal deployment of **belief** to a given state *ipso facto* incurs the commitment that this state is liable to third-personal deployments of **belief** (given suitable access), and any third-personal deployment of **belief** to a given state *ipso facto* incurs the commitment that this state is liable to first-personal deployments of **belief** (given suitable access). So, if the conditions for the possession of **belief** should mention (T) as a commitment incurred in deploying **belief** third-personally, then these conditions should also mention (F*) as a commitment incurred in deploying **belief** first-personally. I noted in Chapter 3 that (T) might either be included as an explicit sub-condition of (T*), or merely be entailed by it—I tend to favor the latter option, as it is more in keeping with the general spirit of Peacocke’s approach. *Ditto* for (F*). However this issue is settled, it is overwhelmingly plausible that something like (T) and (F*) state commitments incurred in deploying **belief** in accordance with its deployment conditions.¹⁴

Given that the case is non-pathological and given (F), *S* is in a position to deploy **belief** to her conscious belief that *p* immediately. Given (F*), however, were *S* to deploy **belief** to her own belief that *p*, she would thereby incur a commitment to being in a state that has the same content-dependent role in making her intelligible as the role of *a*’s state of standing in **R** to **p** in making *a* intelligible, were *a* to be in that state. That is, were *S* to deploy **belief** to her own belief that *p*, she would thereby incur the commitment that her own belief that *p* has the same content-dependent role in making her intelligible as the role a belief that *p* she would attribute to another person would have in making this person intelligible, if this person believed that *p*.

Given that it is highly plausible that belief attribution is governed by the Principle of Correspon-

¹⁴ That (F*), in addition to (T), states a commitment incurred in deploying **belief** in accordance with its deployment conditions does not contradict my earlier claim that **belief** is an essentially first-personal concept. **Belief** is essentially first-personal because the cognitive access to its ground of deployment mentioned in its first-person clause is primary, while the cognitive access to its ground of deployment mentioned in its third-person clause is derivative. See also n. 30, Chapter 3.

dence not merely *de facto*, but as a matter of necessity,¹⁵ it is highly plausible that the commitment mentioned in (F*) entails a commitment to judge thoughts of the first-person form ***iRp*** only if *p*. Accordingly, any first-personal deployment of **belief** constitutes a (renewed) commitment to the truth of the belief conceptualized as a belief. For this reason, rational believers are required to answer the question, ‘Do I believe that *p*?’ as if they were answering the question, ‘Should I believe that *p*?’

Now suppose that *S* does not believe that *p* and asks herself, ‘Do I believe that *p*?’ In asking this question, *S* is deploying **belief**, and ***i***, to the content, ***p***, which she does not believe, in an inquiring mode. In this case, *S* is not in a position to determine whether to deploy **belief** until she has determined whether to deploy ***p***. So we get our result directly—in this case, she should answer the question, ‘Do I believe that *p*?’ as if answering the question, ‘Should I believe that *p*?’—this unsurprising result is as it should be, since we already obtained it for rational believers generally while reflecting on the previous case.¹⁶

¹⁵ For a brief discussion of belief attribution, see Chapter 2, Belief Attribution.

¹⁶ If I am right that belief constitutively aims at truth, the first-person clause of **belief** should be reformulated accordingly, along the following lines:

- (F[#]) the thinker finds the first-person content that he stands in ***R*** to ***p*** primitively compelling whenever he finds ***p*** primitively compelling as a content of conscious belief, and he finds the first-person content that he stands in ***R*** to ***p*** primitively compelling because he finds ***p*** primitively compelling as a content of conscious belief.

There is good reason for replacing Peacocke’s (F) with (F[#]). Clearly, **belief** isn’t merely the concept of a state one tends to attribute to herself while undergoing it consciously—desire, fear, intention, conjecture, imagination, etc. are all states one tends to attribute to oneself while undergoing them consciously. Rather, **belief** is the concept of a state one tends to attribute to oneself while undergoing it consciously and endorsing its content as true.

Replacing Peacocke’s (F) with (F[#]) greatly simplifies the foregoing argument, as it would no longer need to appeal to (F*) to reach its conclusion. Given (F[#]), the question, ‘Do I believe that *p*?’ can only be answered in the affirmative if the question, ‘Is it the case that *p*?’ can be answered in the affirmative. Consequently, whenever a believer answers the question, ‘Do I believe that *p*?’ in the affirmative, she incurs a (renewed) commitment to *p*. Accordingly, rational believers should answer the question, ‘Do I believe that *p*?’ as if answering the question, ‘Should I believe that *p*?’

Replacing Peacocke’s (F) with (F[#]) has further advantages. The replacement would enable one to explain the Principle of Correspondence as a conceptual requirement, as it is entailed by the conjunction of (F[#]) and (T). Since the Principle of Coherence can clearly be explained as a conceptual requirement in terms of (T), the switch to (F[#]) would enable one to explain the Principle of Charity as underwritten by **belief**. With (F), the order of explanation seems backward—one has to appeal to the conjunction of (F*) and the Principle of Correspondence to explain the alethic commitment incurred by first-person deployments of **belief**.

This completes my case for my claim that my account of rational belief (in the broad sense) supports and explains Moran's view as it pertains to belief. Note, however, that my reasoning did not require me to invoke the notion of truth (which I provisionally included among the minimal requirements for doxastic reflection). Since I consider the deliberative stance Moran characterizes to be a reflective stance *par excellence*, I take this as conclusive evidence that possessing the notion of truth is not minimally required for doxastic reflection. Examination of Moran's deliberative stance reveals that, in addition to the concepts constituting the individuating and current synchronic cognitive contents of one's belief, the capacity for doxastic reflection and, so, rational belief (in the broad sense) minimally require possession of the concept of belief and of the first-person concept, and the capacity to entertain a question—more precisely, the capacity to entertain the deliberative question, 'Do I believe that...' (which is properly answered as if answering the question, 'Should I believe that...?')

In "The Norm of Truth and The Concept of Belief" (Chapter 2) I discussed Shah's view, according to which the subjection of belief to the norm of truth is a matter of conceptual necessity. By this, he means that the norm of truth governs deployments of the concept belief rather than belief itself. I am now in a position to say a little more in response to Shah. On my view, the subjection of belief to the norm of truth is a matter of metaphysical necessity (i.e., it has to do with belief, rather than the concept of belief—although, of course, a concept wouldn't be the concept of belief if it did not reflect the subjection of belief to the norm of truth¹⁷). A belief is individuated in part by its individuating cognitive content, which is a complex concept. Concepts are individuated by their possession conditions, which determine consistent deployment

However plausible it is that (F) should be replaced with (F[#]), I prefer to show that the argument goes through without building the thesis that belief aims at truth constitutively into the concept of belief. The argument I offer makes a crucial appeal to the Principle of Correspondence as guiding belief attribution. I think that there is sufficient philosophical agreement concerning this matter to warrant the appeal.

¹⁷ This is a further reason to amend the first possession condition of **belief** as suggested in n. 16, this chapter.

conditions. For beliefs, deployment conditions determine truth conditions (more precisely, the deployment conditions of the individuating cognitive content of a belief determine its truth conditions).¹⁸ Hence, beliefs are governed by the norm of truth as a matter of metaphysical necessity. The capacity to form the belief that *p*, whose individuating cognitive content is a complex concept I will designate by '*p*,' however, does not suffice for the capacity to abide by the norm of truth as it governs *p*. The capacity to abide by the norm of truth as it governs *p* requires, in addition to possessing *p* (and perhaps some additional identifying concepts), possession of the concept of belief and of the first-person concept, and the capacity to entertain the deliberative question, 'Do I believe that...?' The capacity to abide by the norm of truth, then, is largely a conceptual matter. *Pace* Shah, however, the norm of truth governs deployments of both belief and the concept of belief (it governs belief by way of governing the deployment of the concepts that constitute the synchronic and individuating cognitive contents of belief).

¹⁸ Ignoring vagueness.

CHAPTER 5 TWO APPLICATIONS

This last chapter is devoted to applying the foregoing account to specific issues—first, the issue raised for semantics by *de re*, *de dicto* and *de se* attitude reports, and, second, the issue raised for epistemology by Gettier’s cases—in order to show that the theoretical approach I have been pursuing holds some promise of being explanatorily fruitful.

De Re, De Dicto and De Se Attitude Reports

In what follows, I show that the account of belief content I proposed in Chapter 3 yields a natural and plausible explanation of the characteristic differences between *de re*, *de dicto* and *de se* belief reports, which can be generalized to explain the corresponding differences among reports concerning other types of propositional attitude.

Some Definitions

I am adopting Steven Boër’s definitions of substitutional and quantificational transparency:

Where

a and b are denoting terms,

v is a variable, and

a occupies the term position p_E in the expression E ,

let $E(b/a[p_E])$ be the result of substituting b for a at the position p_E , and

let $E(v/a[p_E])$ be the result of substituting v for a at the position p_E .

Then, relative to a given reading of a sentence S :

The term position p_E in S is said to be substitutionally transparent or ... opaque according as the inference from S and $\ulcorner b = a \urcorner$ to $S(b/a[p_S])$ is valid or not on that reading of S .

The term position p_E in S is said to be quantificationally transparent or ... opaque according as the inference from S to \ulcorner Some existing thing v is such that $S(v/a[p_S]) \urcorner$ [v being foreign to E] is valid or not on that reading of S . (Boër 2007: 32-33)

I am also adopting Boër’s definitions of purely *de re* and *de dicto* attitude reports:

Where ‘ V ’ is an attitudinal verb, a report of the form ‘ A V s that P ’ is purely *de dicto*/purely *de re* iff every term position in the embedded sentence ‘ P ’ is [a substitutionally opaque/transparent position in ‘ A V s that P ’]. (*Ibid*: 35)

I will diverge from Boër's convention concerning attitude reports in whose complement clause some but not all term positions are substitutionally transparent. While Boër counts these as *de re* reports that are *de re* with respect to any item denoted by an expression occupying a substitutionally transparent position in the report's complement clause, I will call them 'substitutionally mixed reports.' On my nomenclature, substitutionally mixed reports are *de dicto/de re* with respect to any expression occupying a substitutionally opaque/transparent position in their complement clause, but are neither *de dicto*, nor *de re* as such. Any expression occupying a substitutionally opaque/transparent position in the complement clause of an attitude report may be said to be embedded *de dicto/de re* in that report.

Desiderata

I will take the following on board as *prima facie* data that any adequate account of *de re*, *de dicto* and *de se* attitude reports must either elucidate or explain away:

Prima Facie Semantic data:

- (i) Expressions may be embedded *de re*, *de dicto* or *de se* in attitude reports.
- (ii) Any expression that is embedded *de dicto* in some attitude report occupies a position that is quantificationally opaque in that report.
- (iii) Controversial, though plausible by my lights: any expression that is embedded *de re* in some attitude report occupies a quantificationally transparent position in that report.
- (iv) Any reflexive expression embedded in a *de se* attitude report occupies a quantificationally transparent position in that report.
- (v) Logical relations between *de se*, *de re* and *de dicto* attitude reports:

Consider:

(1) John believes that *he himself* is in danger.

(2) John believes that John is in danger.¹

Boër (*Ibid*: 41) notes that (1) seems logically stronger than (2) when (2) is taken *de re* with respect to the embedded name 'John,' but logically independent from (2) when (2) is taken purely *de dicto*.

¹ The example is borrowed from Boër (sentences are relabeled).

Note, however, that *de se* belief reports only seem logically stronger than those belief reports that are obtainable from them by way of replacing an embedded reflexive expression with a non-reflexive but co-referring/denoting expression, taken *de re*.

Consider, by contrast to (1) and (2):

(3) Aristotle believes that he himself is the teacher of Alexander.

(4) Aristotle believes that Aristotle is the teacher of the king of kings.

In this case, (3) only seems to entail (4) when (4) is taken *de re* with respect to the embedded name 'Aristotle' when we also stipulate that (3) and (4) are to be taken *de re* with respect to the complement of 'the teacher'. When no such stipulation is made, (3) and (4) seem to be logically independent. With respect to their embedded non-reflexive expressions, then, *de se* attitude reports appear to have unexceptional logical profiles.

Prima Facie Psychological data:

- (vi) A believer whose current, conscious belief is reported *de dicto* is in a better position to determine whether this report is true than a believer whose current, conscious belief is reported *de re*.
- (vii) The attitudes expressed by true and sincere first-personal *de se* attitude reports play a special role in their speakers' psychological and behavioral economy.

Theoretical Tools

I proposed in Chapter 3 that belief is individuated in part by its individuating cognitive content, which I distinguished from its synchronic cognitive content, i.e., the content that belief contributes at a time to the cognitive processes of which it is a part. What motivated the distinction between synchronic and individuating cognitive contents is the following set of theses²:

Let there be two cognitive contents whose constituents are:

- (i) **a[the X] is F,**
- (ii) **a[the Y] is F,**

where 'a' stands for a simple directly referring concept,
'X,' 'Y' and 'F' stand for distinct predicative concepts,
'X' and 'Y' are tokened in an identifying role, and
'F' is tokened in an attributive role,

² The following differs slightly from what I offered in Chapter 3, as the current formulation incorporates terminology I introduced after I originally presented these theses.

let (i) be the cognitive content of some belief state, b_1 , at some time, t_i ,

let (ii) be the cognitive content of some other belief state, b_2 , some time, at t_j .

Then:

- (a) b_1 and b_2 have the same individuating cognitive content, namely, **a is F** ;
- (b) if some believer, S , has b_1 at t_i , then, at t_j , the canonical way for S to verify that a is F is by way of verifying that the X is F (the same is true, *mutatis mutandis*, for b_2);
- (c) having b_1 and b_2 need not place S in a position to judge that the $X =$ the Y .

I argued that the individuating and synchronic cognitive contents of belief are complex concepts, which are fully constituted by simple concepts in given modes of combination. In our general beliefs, individuating and synchronic cognitive contents always coincide. In our beliefs about particulars, individuating cognitive content is a proper part of synchronic cognitive content.

I proposed that there are different types of simple concept, which differ by virtue of the type of contribution they make to the cognitive content of thought. The cognitive role of a directly referring concept is to contribute a referent for predicative thought. The cognitive role of a predicative concept is to contribute information concerning the instancing of some property. The cognitive role of a logical concept is to operate on elements of thought. Simple directly referring concepts, as such, do not contribute any information about their referent to the cognitive content of thought. Accordingly, they are typically tokened as constituents of complex directly referring concepts, together with predicative concepts tokened in an identifying role—i.e., identifying concepts—used to identify and recognize the referent of one's thought. What identifying concepts one uses to identify and recognize the referent of one's thought might change over time, while one's thought about that referent might nevertheless remain unchanged. For this reason, identifying concepts do not belong to the individuating cognitive content of belief, but only to the synchronic cognitive content of belief. In set-theoretic terms, the difference between the synchronic and the individuating cognitive contents of belief is fully consti-

tuted by identifying concepts—if the synchronic cognitive content of a belief includes no identifying concepts, its individuating and synchronic cognitive contents coincide.

For each of us, however, there is one particular to whom we can refer in thought without needing to deploy any identifying concepts, namely, ourselves. Consider practical deliberation, which I examined in Chapter 4. While deliberating about what to do in a given set of circumstances, the object of my inquiry is to determine what to do, while its subject³ is myself—I am wondering what *I* am to do. In such an inquiry, I am not identifying myself via any given concept. Rather, I am thinking of myself directly. In the course of my inquiry, I'll consider my desire, values, etc, reflecting on thoughts such as 'I want...', 'I value...', etc, and seek to formulate appropriate means-end beliefs. Once I have determined the best course of action available to me in the circumstances, I'll formulate a judgment, 'I ought to ...,' then an intention, 'I will ...' In all these types of thought—reflection on pro-attitudes, second-order practical judgment, intention—the first-person concept is deployed without any accompanying identifying concepts.

It is highly plausible that the causal efficacy and regulating power of practical reflection requires the capacity to refer to oneself in thought directly. Indeed, as I argued in Chapter 4, reflection has regulating rather than motivating power. My practical judgment should not be understood as providing a reason or a motive for intending or for acting over and above my reasons for judging—practical judgment sanctions my reasons, which motivate my intention. If I conceived of myself in one way while pondering my reasons, in another while formulating my practical judgment, and in another yet while formulating my intention, it is not clear that my reasons could function as reasons for my reflected practical judgment and reflected intention without my formulating judgments to the effect that the same individual is involved throughout

³ I am borrowing the term from grammar and its meaning in this context should be understood along the grammatical model.

(given thesis (c) above). It might be argued that my reasons could retain their motivating power even in the absence of any judgment to the effect that the same individual is involved throughout, given that each reflexive thought involves a deployment of my first-person concept (albeit combined with distinct identifying concepts). On this proposal, however, the regulating power of reflection would be inexplicable—reasons could motivate action via a reflected judgment and a reflected intention that could not be viewed as reasons for this reflected judgment or for this reflected intention. Note that I can, at times, reflect on what to do as a person of a certain type. When I do so, however, I am not identifying myself *via* a given concept. Rather, I am thinking of myself directly, but placing certain constraints on my reflection—only reasons that are reasons for a person of that type will count as reasons for me in my reflection. The causal efficacy and regulative power of doxastic reflection can similarly be shown to require the capacity to refer to oneself in thought directly. The capacity for reflection generally, then, requires the capacity to refer to oneself in thought directly. If the argument I offered in Chapter 4 is correct, rational belief (in the broad sense), too, requires the capacity to refer to oneself in thought directly. This serves as an elucidation of psychological datum (vii) above.

Account

***De Re* and *De Dicto* belief reports**

On my nomenclature, a purely *de re* belief report states the individuating cognitive content, while a purely *de dicto* belief report states the synchronic cognitive content, of the belief it reports. This makes good on two of the three semantic data included in (i) above—indeed, I consider it a virtue of the account of belief content I propose that it provides a motivated explanation of the possibility of reporting beliefs either *de re*, or *de dicto*.

Let us start by applying my account to a few specific reports, before briefly examining how it treats different types of belief report. Consider the example Quine offers in “Quantifiers and propositional attitudes”:

There is a certain man in a brown hat whom Ralph has glimpsed several times under questionable circumstances on which we need not enter here; suffice it to say that Ralph suspects he is a spy. Also there is a gray-haired man, vaguely known to Ralph as rather a pillar of the community, whom Ralph is not aware of having seen except once at the beach. Now Ralph does not know it, but the men are one and the same. Can we say of this *man* (Bernard J. Ortcutt, to give him a name) that Ralph believes him to be a spy? (Quine 1956: 179)

Consider the following belief reports (*Ibid*: 179-181, sentences renumbered):

- (5) Ralph believes that the man in the brown hat is a spy.
- (6) Ralph does not believe that the man seen at the beach is a spy.
- (7) Ralph believes that the man seen at the beach is not a spy.
- (8) Ralph believes that Ortcutt is a spy.

Presumably, (7) reports a non-occurrent belief of Ralph, i.e., a belief Ralph has not entertained but which he would accept should he entertain it. Now, let:

- be a simple directly referring concept referring to Ortcutt;
- H** be a complex predicative concept predicating the property, being a man in a brown hat;
- B** be a complex predicative concept predicating the property, being a man seen at the beach;
- S** be a (simple)⁴ predicative concept predicating the property, being a spy; and
- Bel(*A*, [...]) symbolize the relation between some individual, *A*, and a content of belief.

Then, the truth conditions for belief reports (5) – (8), considered first as offered in a *de re* mood, then as offered in a *de dicto* mood, can be represented as follows⁵:

⁴ The parentheses mark my uncertainty as to whether the relevant concept is simple or complex—that a given concept is expressible using a simple expression does not establish that this concept is simple.

⁵ I am not including truth conditions for (8) offered in a *de dicto* mood, since (8) would not plausibly be offered as a *de dicto* report, given Ralph’s poor acquaintance with Ortcutt—see my discussion below of proper names embedded *de dicto* in the complement clauses of belief reports.

<u>De Re</u>	<u>De Dicto</u>
(5 ^R) Bel(Ralph, [o is S])	(5 ^D) Bel(Ralph, [o[the H] is S])
(6 ^R) ~Bel(Ralph, [o is S])	(6 ^D) ~Bel(Ralph, [o[the B] is S])
(7 ^R) Bel(Ralph, [o is ~S])	(7 ^D) Bel(Ralph, [o[the B] is ~S])
(8 ^R) Bel(Ralph, [o is S])	

Intuitively, Ralph has inconsistent beliefs, yet unobjectionably so, since he has no way of telling that the individual he identifies as the man in a brown hat is also the man he identifies as the man seen at the beach. The distinction between the individuating and the synchronic cognitive contents of belief vindicates and explains this intuition (given theses (a) – (c) above).

Further, there is obvious utility in being able to report on precisely these two types of belief content—not least in order to explain why, e.g., Ralph is not liable to blame for having inconsistent beliefs about Orcutt. We offer *de dicto* belief reports when we are interested in presenting matters from the point of view of the believer, and *de re* belief reports when we are interested in presenting matters from an external point of view. The distinction between the synchronic and the individuating cognitive contents of belief supplies truth conditions for such reports that are tailored to these interests.

Let us now examine various types of *de re* and *de dicto* belief report in some greater detail. To start, we should note that, to be used successfully, language, like thought, must identify the particulars to which it means to refer. Thought does so through the deployment of identifying concepts (except in the case of some reflexive thought). Language does so through the use of referring expressions—i.e., proper names, indexicals, definite descriptions. Although a purely *de re* belief report means to state the individuating cognitive content of the belief it reports, it must nevertheless identify the particular(s) to which the belief refers. It may do so in any of the

customary ways. When a purely *de re* belief report identifies every particular in its complement clause via a definite description, the linguistic content of its complement clause may be said to be the linguistic equivalent of a possible synchronic cognitive content of a belief individuated in part by the reported individuating cognitive content.

When reporting beliefs about particulars, both *de re* and *de dicto*, we frequently use proper names. In the complement clause of a *de re* belief report, a proper name should be treated as referring to its referent directly. Matters are different in *de dicto* belief reports. Although proper names have no conceptual equivalents, it is plausible that, in the complement clause of a *de dicto* belief report, a proper name should be treated as referring to its referent as identified via a recognitional concept, that is, a settled identifying concept, which is settled in the sense that the believer has a tendency to deploy it, rather than other available concepts, to identify and recognize the relevant particular. Accordingly, the truth-conditions of *de dicto* belief reports involving proper names include a complex directly referring concept, consisting of a simple referring concept combined with a predicative concept tokened in a recognitional role, i.e., a settled identifying role.

The case of indexicals (specifically, demonstratives and personal pronouns) is similar. Plausibly, if the complement clause of a sentence expressing a belief report contains a demonstrative or a pronoun and is uttered *de re* in a context that fixes the reference of this demonstrative or pronoun on some particular, this pronoun or demonstrative should be treated as referring to this referent directly in the uttered *de re* belief report. However, demonstratives and personal pronouns are also occasionally used in the complement clauses of belief reports that seem offered in a purely *de dicto* spirit, at least inasmuch as they intend to present matters from the point of view of the believer. Such uses tend to occur when the believer is well acquainted with

the particular to which/whom the demonstratives and pronouns refer. The belief report, ‘John missed his train because he believed that she [pointing at John’s mother] meant to give him a ride to the station.’ lends itself to such a reading. If such reports are indeed to be treated as purely *de dicto*, then any demonstratives or pronouns in their complement clause should be treated, I think, as functioning as proper names, i.e., as referring to their referent as identified via a recognitional concept.⁶

We also report the beliefs people have concerning non-existent particulars. In Chapter 3, I considered two alternative ways in which my account might accommodate the possibility of such beliefs. On the first alternative, the relevant component of the synchronic cognitive content of beliefs concerning non-existent particulars includes identifying concepts, but no simple directly referring concepts. Accordingly, such beliefs are not *about* non-existent particulars at all, and cannot be reported purely *de re*. On this alternative, a true *de dicto* report of a belief concerning a non-existent particular reports the synchronic cognitive content of a false belief. On the second alternative, beliefs concerning existent and non-existent particulars have the same logico-semantic form, the difference between such beliefs being that the former, but not the latter, are truth-evaluable. On this alternative, a true report of a belief concerning a non-existent particular, offered *de dicto*, reports the synchronic cognitive content of a non-truth-evaluable belief; offered *de re*, it reports the individuating cognitive content of a non-truth-evaluable belief.

***De Se* belief reports**

Let us now turn to *de se* belief reports. On the account I have endorsed, concepts are fully individuated by their possession conditions, which specify the conditions under which a thinker

⁶ If there are such cases, they are distinct from cases involving mixed reports, in which the pronouns or demonstratives are embedded *de re* (perhaps because of lack of interest in offering an exact account of the entire synchronic cognitive content of the reported belief). Purely *de dicto* utterances of belief reports involving a pronoun functioning as a proper name should be distinguished from cases involving the use of a pronoun whose antecedent is also included within the complement clause of the purely *de dicto* report. In such cases, both pronoun and antecedent should be treated as referring to the same particular as identified in the same way.

possessing the relevant concept finds transitions to and from propositional attitudes containing this concept primitively compelling. We deploy our first-person concept under two types of circumstance. The first type can be characterized generally as involving a conscious attitude or action on the part of the thinker, when the thinker finds the first-person thought that she is having that attitude or performing that action primitively compelling. So, when I have the conscious belief that p , I find a thought of the form, **e believes p** (where ‘**e**’ stands for my first-person concept, which is distinct from anyone else’s first-person concept), primitively compelling. Likewise, when I am perceiving F consciously, and when I am ϕ ing consciously, I find thoughts of the form, **e perceives F** , and **e ϕ s**, primitively compelling. The same can be said for any other type of attitude that can be entertained consciously. For now, let us characterize circumstances of the second type negatively, as circumstances that are distinct from circumstances of the first type. Circumstances of the second type might include, e.g., deploying my first-person concept to the individual whose reflection frightened me as I caught it in a mirror upon recognizing this reflection as mine.

Concepts, however, are not thinker specific, as they are types of thought content. My first-person concept is a simple directly referring concept whose cognitive function is to refer to me. Given that my first-person concept is a simple directly referring concept, tokening it contributes no information about me to my cognitive processes—the function of contributing such information belongs to predicative concepts. Accordingly, there does not seem to be anything to distinguish my first-person concept from the simple directly referring concept other thinkers use when they have thoughts about me. What I have been calling my first-person concept, then, is the simple directly referring concept that refers to me.

The foregoing reflection indicates that, for beings capable of conscious reflexive thought and action, simple directly referring concepts are individuated by at least two possession conditions. The first possession condition pertains to first-personal deployments of the concept, when the thinker has a conscious attitude or consciously performs an action and finds the first-person thought that she is having that attitude or performing that action primitively compelling. The second possession condition pertains to other deployments of the concept, which may but need not be first personal, when the individual is recognized as the same individual to whom the concept applies in the circumstances mentioned in the first possession condition.⁷ For the sake of simplicity, I will continue using the phrase, ‘first-person concept.’ Henceforth, the phrase should be understood as meaning ‘simple directly referring concept, as possessed by its referent.’

Accordingly, the two types of circumstance under which I can deploy my first-person concept can now be characterized as follows: I can deploy this concept under circumstances that fall under the deployment condition determined by its first possession condition, and I can deploy it under circumstances that fall under the deployment condition determined by its second possession condition. Clearly, there is no need for me to deploy any identifying concept when I deploy my first-person concept under circumstances that fall under its first deployment condition. By contrast, under circumstances that fall under its second deployment condition, my first-person

⁷ This is not to say that tokening a simple concept that refers to a being capable of conscious attitudes and actions contribute information to the effect that this being is capable of conscious attitudes and actions to one’s cognitive processes. The point is merely the following: Possessing the concept, **elka**, requires finding first-person thoughts of the form, **elka A’s (that) p** (where **A** stands for an attitude concept and **p** stands for some content) primitively compelling while having, and because of having, the corresponding attitude (the same can be said for conscious actions, *mutatis mutandis*). It also requires being disposed to deploy **elka** in other, third-personally defined circumstances. These conditions hold for anyone who possesses **elka**—I just happen to be the only person who is ever in a position to deploy **elka** in accordance with the deployment conditions determined by the first possession condition. There is no corresponding first-person clause for simple concepts referring to particulars that incapable of conscious attitudes or actions.

Interestingly, although I come to the issue from a different tradition, the distinction between the possession conditions of simple concepts that refer to particulars that are capable of conscious attitudes and actions vs. those that are incapable of conscious attitudes and actions, is clearly reminiscent of Jean-Paul Sartre’s distinction between things that are for-themselves and things that are in-themselves.

concept is deployed as a result of recognition. The causal ancestry of beliefs involving a deployment of the first-person concept in circumstances of the second type includes a deployment of an identifying concept (typically, a recognitional concept), although the synchronic cognitive content of the resulting *de se* belief need not include this identifying concept.

There are, then, two types of beliefs one may have about oneself. First, there are beliefs in which the first-person concept is deployed under circumstances of the first type. Second, there are beliefs whose causal ancestry involves a deployment of an identifying concept, *C*, and a judgment of the form, I am the *C*. The first-personal component of the synchronic cognitive content of both types of belief includes the first-person concept. The first-personal component of the synchronic cognitive content of both types of belief may, but need not, include an identifying concept. So, I may think that I, the Queen of Sheba, deserve special treatment, and I may think that I, the person whose reflection just caused me such a fright, am quite handsome after all. Beliefs in which the first-person concept is deployed under circumstances of the first type are genuine *de se* beliefs. If the first-personal component of the synchronic cognitive content of a genuine *de se* belief includes an identifying concept, the self is not identified *via* this concept.

De se belief reports are offered *de re* with regard to any reflexive expression included in their complement clause. Accordingly, *de se* belief reports may be substitutionally mixed, as other expressions may be embedded *de dicto* in their complement clause. This account is intuitively plausible—*de se* belief reports are offered as beliefs the believer has about herself, not as beliefs the believer has about herself as identified in a given way. Nevertheless, when reporting the beliefs someone has about herself, one may be interested in presenting the non-reflexive elements of her belief from her point of view, i.e., *de dicto*, or from an external point of view, i.e., *de re*. On the account I am proposing, the truth conditions of *de se* belief reports mention

the first-person concept, without any identifying concepts, for any instance in which a reflexive expression appears in the complement clause. In any other respect, *de se* belief reports are treated as any other belief report.

Assessment

Let us assess how well the account I have proposed fares with the *prima facie* semantic and psychological data I listed above. Clearly my account explains (i), i.e., that expressions can be embedded *de re*, *de dicto* and *de se* in attitude reports. On the first alternative account of the content of beliefs concerning non-existent particulars I considered above, (ii), i.e., that any expression that is embedded *de dicto* in some attitude report occupies a position that is quantificationally opaque in that report, and (iii), i.e., that any expression that is embedded *de re* in some attitude report occupies a quantificationally transparent position in that report, are also secured. On the second alternative, (ii) is secured but (iii) isn't (unless one thinks that "non-existent" particulars exist⁸).⁹ My account also clearly supports (iv), that any reflexive expression embedded in a *de se* attitude report occupies a quantificationally transparent position in that report. Finally, let us consider (v). First, it is clear from what I have said above that on my account, *de se* reports have unexceptional logical profiles with respect to their embedded non-reflexive expressions. Let us now consider Boër's two sentences,

- (1) John believes that *he himself* is in danger.
- (2) John believes that John is in danger.

Now, let:

- j** be the simple directly referring concept that refers to John;
- D** be a simple predicative concept predicating the property, being in danger;

⁸ Steven Schiffer (2003) has recently endorsed a thesis along these lines.

⁹ I tend to favor the first alternative because it secures (iii).

J be a complex predicative concept, tokened in an identifying role, that predicates some set of properties John instances

Then, on the account I am proposing, the truth conditions for (1), and for (2), offered first in a *de re* mood, then in a *de dicto* mood, would be:

(1^S) Bel(John, [**j is D**])

(2^R) Bel(John, [**j is D**])

(2^D) Bel(John, [**j[the J] is D**])

(1^S) entails (2^R)—by virtue of being identical with it—but does not entail (2^D). (2^D) entails (1^S), but John need not be in a position to know this (in order to be in a position to know this, he must be in a position to judge that he is the *J*). On my account, (1) and (2), offered in a *de dicto* mood, are not logically independent—the latter is logically stronger than the former. The appearance of logical independence, however, is explained by John’s doxastic situation, which does not enable him infer (1) from (2) directly. Accordingly, my account makes good on (v) also.

Clearly, my account also makes good on psychological data (vi) and (vii). On my view, a believer whose current, conscious belief is reported *de dicto* is ideally placed to determine whether this report is true, while a believer whose current, conscious belief is reported *de re* need not be. Finally, my account explains why the attitudes expressed by true and sincere first-personal *de se* attitude reports can play a special role in their speakers’ psychological and behavioral economy. Under circumstances that fall under the first deployment condition of the first-person concept, such attitudes involve a direct deployment of a directly referring concept to the self, which is not identified *via* any identifying concept even when the self is conceived as bearing given properties. It is this capacity of beings capable of conscious reflexive thought and action to refer directly to themselves in reflexive thought which, I think, secures their continuity as persons through reasons, thought, reflection, intention and action. True and sincere *de se*

attitude reports in which non-reflexive expressions are embedded *de dicto* offer us as intimate an access as is available to us into the mind of another person.

Gettier's Cases

In what follows, I show that my account of belief content also yields a natural and plausible explanation of what goes epistemically awry in the cases Gettier presents in “Is Justified True Belief Knowledge?”

There is a tendency among epistemologists to call any case that casts doubt on the true-and-justified-belief analysis of knowledge (henceforth, the ‘classical analysis’) a ‘Gettier-style case.’ So, epistemologists often refer to Goldman’s Barn County case (Goldman 1976) as a Gettier-style case. The following does not speak to Goldman’s case—generally, it does not speak to any case that shows that knowledge might be defeated by evidence unavailable to the believer. Since my aim here is neither to rescue the classical analysis, nor to offer an alternative analysis, I need not answer every Gettier-style case, be it on behalf of the classical analysis or of my alternative. It is often assumed that that the “Gettier problem” requires a single solution. Since the term, ‘Gettier-style case’ is now applied to a fairly disparate family of cases, whose common characteristic is that they cast doubt on the classical analysis, it might very well be that solving the “Gettier problem” will involve a family of solutions, each tailored to the type of case it seeks to answer. My interest here is with the original members of this family of Gettier-style cases, i.e., the cases offered by Gettier.

In “Is Justified True Belief Knowledge?” Gettier submits two cases intended to show that true and justified belief is not sufficient for knowledge. So far, I have avoided undertaking any commitment concerning the nature of knowledge. The following minimal requirement will suffice to discuss Gettier’s cases. I have said that the synchronic and individuating cognitive contents of belief are complex concepts, and that complex concepts are fully constituted by

simple concepts in given modes of combination. I have also said that concepts are individuated by their possession conditions, which determine consistent deployment conditions. The deployment conditions of the individuating cognitive content of a belief determine the truth conditions of that belief.¹⁰ Further, in my view, the norm of truth governs belief by virtue of governing the deployment of its synchronic and individuating cognitive contents. I propose that having knowledge requires satisfying the norm of truth as it applies to one's belief. Now, satisfying the norm of truth as it applies to any belief is a matter of satisfying the deployment conditions of the simple concepts, in the relevant modes of combination, that constitute the synchronic and individuating cognitive contents of that belief. The deployment conditions of any concept mention the types of circumstance under which this concept is correctly deployed. It is highly plausible that the norm of truth sets a dual requirement on deployment of any concept: First, it sets a correctness requirement, by requiring that circumstances of the relevant type obtain and, second, it sets an evidence requirement, by requiring that one who deploys this concept has evidence that circumstances of the relevant type obtain. I am proposing, then, that having knowledge requires satisfying the correctness and evidence requirements for the simple concepts, in the relevant modes of combination, that constitute the synchronic and individuating cognitive contents of that belief.

We are now in a position to examine Gettier's first case:

Case I:

Suppose that Smith and Jones have applied for a certain job. And suppose that Smith has strong evidence for the following conjunctive proposition:

d. Jones is the man who will get the job, and Jones has ten coins in his pocket.

... Proposition (d) entails:

e. The man who will get the job has ten coins in his pocket.

¹⁰ Once again, ignoring vagueness.

Let us suppose that Smith sees the entailment from (d) to (e), and accepts (e) on the grounds of (d), for which he has strong evidence. In this case, Smith is clearly justified in believing that (e) is true.

But imagine, further, that unknown to Smith, he himself, not Jones, will get the job. And, also, unknown to Smith, he himself has ten coins in his pocket. (Gettier 1963: 122)

Let us examine (d) and (e), understood as the synchronic and individuating cognitive contents of beliefs endorsed by Smith. Let:

j be the simple directly referring concept that refers to Jones;

J be a complex predicative concept, tokened in an identifying role, that predicates some set of properties Jones instances;

M be a complex predicative concept predicating the property, being the man who will get the job;

T be a complex predicative concept predicating the property, having ten coins in one's pocket.

Then (d) and (e), understood as the synchronic and individuating cognitive contents of Smith's beliefs, can be represented as follows:

Synchronic Cognitive Content

d. *j*[**the*J***] **is** *M* \wedge *j*[**the*J***]
is *T*

e. *j*[**the*M***] **is** *T*

Individuating Cognitive Content

d. *j* **is** *M* \wedge *j* **is** *T*

e. *j* **is** *T*

Notice that I differ from Gettier as to what referent is properly assigned the expression, 'the man who will get the job,' as it figures in (e) understood as reporting the content of Smith's belief. In my view, Gettier misinterprets (e) as a descriptive statement. One tends to overlook this misinterpretation because (e) is first introduced as a descriptive statement. In the context of our inquiry concerning the justification of Smith's belief, however, (e) reports the content of Smith's belief. Accordingly, 'the man who will get the job' does not refer to the man who will get the job, i.e., to Smith, but to the man Smith believes will get the job, i.e., to Jones—remember, Smith reasons from (d) to (e); in this case, the referent the simple directly referring concept in

‘*j*[**them**] **is T**’ is established by the cognitive history of that belief, as Jones is not the man who will get the job.¹¹ Further, in the context of this inquiry, (e) is clearly offered *de dicto*. So, on my account, (e) reports the synchronic cognitive content of Smith’s belief. Now, this belief is individuated in part by the individuating cognitive content, *j is T*, and can also be reported *de re* as the belief that Jones has ten coins in his pocket. By hypothesis, this belief is true and Smith has strong evidence for it (presumably, one can’t have strong evidence for a conjunctive belief unless one also has strong evidence for believing each conjunct).

In my view, then, the problem with (e) is not that (e) is true in virtue of the number of coins in Smith’s pocket—it is true in virtue of the number of coins in Jones’ pocket, concerning which Smith has a belief based on strong evidence. The problem, rather, is that Smith misidentifies Jones as the man who will get the job. Misidentification of particulars, I think, is a significant epistemic issue, although to my knowledge it has been entirely ignored.¹² As I noted earlier, although identifying concepts are not part of the individuating cognitive content of our beliefs about particulars, they are crucial to our cognitive lives, since our cognitive access to particulars other than the self is by way of identifying concepts. At any given moment we are, as it were, behind the veil of the identifying concepts we token at that moment and, so, prone to errors of misidentification by our very doxastic condition, which I articulate in theses (a) – (c) in Chapter 3 and reiterate the foregoing section.

Misidentification is an epistemic defeater; that is, misidentification defeats knowledge. This is because a belief affected by misidentification, even when true and well grounded in

¹¹ Matters would have been different if Smith had not reasoned from (d) to (e) and the content of (e) had been, ‘Whoever will get the job has ten coins in his pocket. Then Smith’s belief would have been true, though not based on evidence.

¹² The reason why misidentification of particulars has been ignored is that epistemology tends to focus on propositions believed, rather than concepts deployed. On my nomenclature, this is tantamount to focusing on individuating content, which leaves identifying concepts out.

evidence, is alethically defective, inasmuch as it involves an improper deployment of an identifying concept—although such a belief abides by the norm of truth with regard to its individuating cognitive content, it fails to abide by the norm of truth with regard to its synchronic cognitive content. Misidentification and falsehood are epistemic defeaters for analogous reasons. To evaluate a belief as false is to evaluate it negatively along the dimension of correctness because its individuating cognitive content fails to satisfy its correctness requirement by virtue of being constituted by some concept that fails to satisfy its correctness requirement. Beliefs about existing particulars are often evaluable as false because their synchronic and individuating cognitive content is constituted by some predicative concept, tokened in an attributive role, that fails to satisfy its correctness requirement. To evaluate a belief as subject to an error of misidentification is to evaluate it negatively along the dimension of correctness because its synchronic cognitive content fails to satisfy its correctness requirement by virtue of being constituted by some predicative concept, tokened in an identifying role, that fails to satisfy its correctness requirement. Just as one may, but need not, be rationally culpable for having a false belief, so one may, but need not, be rationally culpable for having a belief subject to an error of misidentification. Whether one is rationally culpable for deploying any concept depends on whether one satisfies the evidence requirement for that concept. Since Smith has strong evidence for his belief that Jones is the man who will get the job, Smith's misidentification, too, is non-culpable.

Let us now turn to Gettier's second case:

Case II:

Let us suppose that Smith has strong evidence for the following proposition:

f. Jones owns a Ford.

... Let us imagine, now, that Smith has another friend, Brown, of whose whereabouts he is totally ignorant. Smith selects three place names quite at random and constructs ... three propositions [including (h)]:

h. Either Jones owns a Ford, or Brown is in Barcelona.

... Imagine that Smith realizes the entailment of each of these propositions he has constructed by (f), and proceeds to accept [them] on the basis of (f).

But imagine now that two further conditions hold. First Jones does not own a Ford, but is at present driving a rented car. And secondly, by the sheerest coincidence, and entirely unknown to Smith, the place mentioned in proposition (h) happens really to be the place where Brown is. (Gettier 1963: 122-123)

Let us now consider (h), i.e., 'Either Jones owns a Ford, or Brown is in Barcelona.' Fortunately, we need not examine the detailed synchronic and individuating cognitive components of Smith's belief in this case, as such examinations are somewhat tedious. For the sake of simplicity, let us assume that (h) states the individuating cognitive content of Smith's belief. Note that, on the account I am proposing, (h) is true (as it should be on any account). Further, Smith is deploying his disjunctive concept appropriately, given that he has every reason to believe that Jones owns a Ford. Now, we should notice the following about the norm of truth as it governs (h), as well as its two constituents, (h₁), that Jones owns a ford, and (h₂), that Brown is in Barcelona:

(T) (h) is governed by the norm of truth by way of its constituents, i.e., (h₁), the concept of disjunction, and (h₂).

Although the concept of disjunction is the governing concept in (h), in the sense that it determines the truth conditions of (h) relative to (h₁) and (h₂), relative to these truth conditions (h₁) and (h₂) are nevertheless beholden to the norm of truth independently of each other. Now, I proposed above that knowledge requires satisfying the correctness and evidence requirements for the simple concepts, in the relevant modes of combination, that constitute the synchronic and individuating cognitive contents of one's belief. Given that (h₁) and (h₂) are combined in (h) via the concept of disjunction, satisfying the correctness and evidence requirements for (h) requires

satisfying the correctness and evidence requirements for (h₁) or for (h₂).¹³ (h₁), however, fails the correctness requirement, while Smith fails the evidence requirement relative to (h₂). So, neither (h₁) nor (h₂) is such that Smith satisfies both the correctness and the evidence requirements relative to it. So, Smith fails the norm of truth relative to (h). So, Smith does not know (h). In this case also, Smith's deployment of (h) is non-culpable, since rational culpability is a matter of failing to satisfy the evidence requirement. Satisfying the evidence requirement for (h) requires satisfying the evidence requirements for (h₁) or for (h₂), and Smith satisfies it for (h₁).

Note that the knowledge requirement I propose above, namely, that knowledge requires satisfying the correctness and evidence requirements for the simple concepts, in the relevant modes of combination, that constitute the synchronic and individuating cognitive contents of one's belief, accords with the spirit, if not the letter, of the classical analysis. I have formulated the requirement as a necessary, but not sufficient, for knowledge in order to allow that evidence one does not possess might function as an epistemic defeater (as the Gettier-style cases I am not addressing here might show). My requirement, however, suffices to offer a plausible explanation of what goes amiss in Gettier's cases—both can plausibly be explained in terms of an incorrect, albeit non-culpable, deployment of concepts. As I noted in n. 12, this chapter, epistemology has tended to focus on propositions, rather than concepts, as the basic constituents of belief content. Epistemic properties, such as rational and irrational, justified and unjustified, or correct and incorrect, have consequently been thought to be basic properties of belief, rather than properties instanced by belief by virtue of being instanced by the constituents of the content of belief. Such accounts are limited to assessing the rationality, justification, or correctness of belief as a whole, as they do not have the theoretical tools to submit the components of belief to corresponding

¹³ Clearly, the proposed requirement does not extend to beliefs whose individuating content has the form, [*Cv ~ C*]. Except for vague concepts, possession of a concept plausibly requires implicit knowledge that this concept either applies, or fails-to-apply, to the items in its domain of application.

assessment. In my view, such accounts are too coarse grained and fall prey to Gettier's cases as a result. Knowledge does require belief that is correct and adequately supported by evidence. Gettier's cases involve belief that is incorrect, albeit true and adequately supported by evidence.

Conclusion

The foregoing study has sought to develop a fine-grained account of belief content. On this account, the content of belief is a complex concept, fully constituted by simple concepts in given modes of combination. This account distinguishes the individuating cognitive content of a belief, i.e., the content that enters into the individuation conditions of this belief, from the synchronic cognitive content of this belief, i.e., the content it contributes at a time to the cognitive processes of which it is a part. This account was shown to yield plausible truth conditions for *de re*, *de dicto* and *de se* belief reports, and a plausible explanation of what is epistemically amiss in Gettier's cases. This offers some indication that the approach is fruitful.

I developed this account as a tool to answer a question I did not address in this study, namely, 'What are the first-person and third-person conditions for the deployment of the concept of epistemic justification?' This question was motivated by the suspicion that many disagreements in epistemology, especially the internalism/externalism debate concerning justification and knowledge, arise because the opposing sides privilege one deployment condition at the expense of the other. The next stage in this project, then, will be to apply the foregoing account to this question. This, however, will have to be done in another context.

LIST OF REFERENCES

- Andrews, Kristin (2003): 'Knowing mental states: The asymmetry of psychological prediction and explanation,' in *Consciousness: New Philosophical Perspectives*, Smith, Q. and Jokic, A. eds., Oxford: Oxford University Press, 201-219.
- Bennett, J. (1990): 'Why Is Belief Involuntary?' *Analysis* 50/2, 87-107.
- Boër, S.E. (2007): *Thought-Contents: On the Ontology of Belief and the Semantics of Belief Attribution*, *Philosophical Studies Series* 104, New York: Springer.
- Bratman, M. (1999): *Faces of intention*, Cambridge: Cambridge University Press.
- Breen, N., Caine, D., Coltheart, M., Hendy, J. and Roberts, C. (2000): 'Towards an Understanding of Delusions of Misidentification: Four Case Studies' *Mind & Language* 15/1, 74-110.
- Burge, T. (1993): 'Content Preservation' *The Philosophical Review* 102/4, 457-488.
- (1997): 'Interlocution, Perception, and Memory' *Philosophical Studies* 86/1, 21-47.
- Cohen, L. J. (1989): 'Belief and acceptance' *Mind* 98, 367-389.
- (1992): *An essay on belief and acceptance*, Oxford: Oxford University Press.
- Davidson, D. (1980): *Essays on Actions and Events*, Oxford: Oxford University Press.
- (1984): *Inquiries into Truth and Interpretation*, Oxford: Oxford University Press.
- (1991): *Subjective, Intersubjective, Objective*, Oxford: Clarendon Press.
- Dretske, F. I. (1981): *Knowledge and the Flow of Information*, Cambridge, MA: Massachusetts Institute of Technology (MIT) Press.
- Edgley, R. (1969): *Reasons in theory and practice*, London: Hutchinson.
- Evans, G. (1982): *The varieties of reference*, Oxford: Oxford University Press.
- Frankish, K. (2004): *Mind and Supermind*, Cambridge: Cambridge University Press.
- Fricker, E. (1994): 'Against Gullibility' in *Knowing from Words*, Matilal, B.K. and Chakrabarti, A. eds., New York: Springer, 125-161
- Funkhouser, E. (2003): 'Willing Belief and the Norm of Truth' *Philosophical Studies* 115, 179-195.
- Gettier, E. (1963): 'Is Justified True Belief Knowledge?' *Analysis* 23/6, 121-123.
- Goldman, A.I. (1976): 'Discrimination and Perceptual Knowledge' *The Journal of Philosophy* 73, 771-791.

- (1979): ‘What is Justified Belief?’ in *Justification and Knowledge: New Studies in Epistemology*, Pappas, G.S. ed., Dordrecht: D. Reidel, 1–23; reprinted in (2000): *Epistemology: An Anthology*, Sosa, E. and Kim, J. eds. Oxford: Blackwell Publishers, 340–53.
- Grice, H. P. (1989): *Studies in the Way of Words*, Cambridge MA: Harvard University Press.
- Harman, G. (1986): *Change in view*, Cambridge: Cambridge University Press.
- Hieronimi, P. (2006): ‘Controlling Attitudes’ *Pacific Philosophical Quarterly* 87/1, 45–74.
- Kavka, G. (1983): ‘The Toxin Puzzle’ *Analysis* 43, 33–6.
- Lackey, J. (2006): ‘The Nature of Testimony’ *Pacific Philosophical Quarterly* 87/2, 177–197.
- and Sosa, E., eds. (2006): *The Epistemology of Testimony*, Oxford: Oxford University Press.
- Ludwig, K. and Ray, G. (2002): ‘Vagueness and the Sorites Paradox’ *Philosophical Perspectives* 16, 419–461.
- Ludwig, K. and Leport, E. (2005): *Donald Davidson: Meaning, Truth, Language, and Reality*, Oxford: Oxford University Press.
- Machina, K. F. (1976): ‘Truth, Belief and Vagueness’ reprinted in Keefe, R. and Smith P. (eds.) (1996): *Vagueness: A Reader*, Cambridge, MA: MIT Press, 174–203.
- Moran, R. (1994): ‘Interpretation Theory and the First Person’ *The Philosophical Quarterly* 44/175, 154–173.
- (2001): *Authority and Estrangement: An Essay on Self-Knowledge*, Princeton: Princeton University Press.
- Owens, D. (2000): *Reason without Freedom: The Problem of Epistemic Normativity*, London; New York: Routledge.
- (2003): ‘Does Belief have an Aim?’ *Philosophical Studies* 115, 283–305.
- Peacocke C. (1992): *A Study of Concepts*, Cambridge, MA: MIT Press.
- Quine, W.V.O. (1956): ‘Quantifiers and Propositional Attitudes’ *Journal of Philosophy* 53, 177–186.
- Sainsbury, R. M. (1990): ‘Concepts Without Boundaries’ reprinted in Keefe, R. and Smith P. (eds.) (1996): *Vagueness: A Reader*, Cambridge, MA: MIT Press, 251–264.
- Schiffer, S. (2003): *The Things We Mean*, Oxford: Oxford University Press.
- Searle, J. R. (2001): *Rationality in Action*, Cambridge, MA: MIT Press.
- Shah, N. (2003): ‘How Truth Governs Belief’ *The Philosophical Review* 112/4, 447–482.

- Vahid, H. (2006): 'Aiming at Truth: Doxastic vs. Epistemic Goals' *Philosophical Studies* 131, 303–335.
- Velleman, J. D. (2000): *The Possibility of Practical Reason*, Oxford: Oxford University Press.
- Wedgwood, R. (2002): 'The Aim of Belief' *Philosophical Perspectives: Language and Mind* 16, 267–297.
- Williams, B. (1973): *Problems of the Self*, Cambridge: Cambridge University Press.
- Williamson T. (1992): 'Vagueness and Ignorance' *Proceedings of the Aristotelian Society, Supp.* 66, 145–162.
- (1997a): 'Imagination, Stipulation and Vagueness' *Philosophical Issues* 15, 215–228.
- (1997b): 'Précis of *Vagueness*' *Philosophy and Phenomenological Research* 57/4, 921–928.

BIOGRAPHICAL SKETCH

Elisabeth K. I. Shortsleeve obtained a Bachelor of Arts, with a major in philosophy, then Master of Arts in philosophy, from Boston College. She then pursued her doctoral studies in philosophy at the University of Florida.

A STUDY OF BELIEF

Candidate's name: Elisabeth K. I. Shortsleeve
Phone number: 352-392-2084
Department: Philosophy
Supervisory chair: Kirk Ludwig
Degree: Ph.D.
Month and year of graduation: August 2009

This study investigates the nature of belief and argues that the authority of the norms to which belief is subject is underwritten by the very nature of belief. Philosophy studies questions that might arise for any reflective individual, but does so carefully and rigorously. In this case, the question is, ‘Why should all my beliefs be beholden to truth when having a false belief can sometimes be of greater benefit to me than having a true belief?’ This study offers a careful, motivated answer to this question. Though technical, it is of interest to anyone for whom this question arises.