

UCLA

UCLA Electronic Theses and Dissertations

Title

Embodied Rational Agency

Permalink

<https://escholarship.org/uc/item/4q92g4w6>

Author

Luthra, Yannig Ashoka

Publication Date

2013

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
Los Angeles

Embodied Rational Agency

A dissertation submitted in partial satisfaction of the requirements
for the degree Doctor of Philosophy
in Philosophy

by
Yannig Ashoka Luthra

2013

Abstract of the Dissertation
Embodied Rational Agency
by
Yannig Ashoka Luthra
Doctor of Philosophy in Philosophy
University of California, Los Angeles, 2013
Professor Tyler Burge, Chair

My dissertation aims to make progress in understanding what it is to be an embodied rational agent. I argue that human agency consists not just in our rational capacities for thought and understanding, but also in non-rational capacities, like our physical ability to put one foot in front of the other and walk. In trying to understand the nature of embodied rationality, my dissertation takes steps toward developing the idea that our activity as free, rational beings consists not just in cognition, but also in bodily action.

I suggest that our success as practical reasoners depends constitutively on our physical abilities. Practical reason's primary function is to guide action, even though acting is not a part of one's practical reasoning. Practical reason differs in this respect from theoretical reason, since theoretical reason functions to guide belief, and forming a belief is a part of one's reasoning. In guiding action, practical reason "reaches past itself" in a way that theoretical reason does not. Since acting is not a part of one's reasoning, practical reason can only succeed in fulfilling its primary function as part of a larger action system. This larger action system crucially includes the physical abilities that we rely on to execute most of our actions. Thus, our success as practical reasoners depends not just on our capacities for practical thought and understanding, but also on the larger action system that those capacities function as a part of, including our physical abilities.

The dissertation of Yannig Ashoka Luthra is approved.

David Blank

Pamela Hieronymi

Gavin Lawrence

Tyler Burge, Committee Chair

University of California, Los Angeles

2013

Dedicated with love and gratitude to Christine, my mother.

Contents

1. Introduction	2
2. Nonrational aspects of agency	20
Two rationalist views about agency	24
Criticizing rationalism about agential control	27
Criticizing rationalism about action	37
3. Interpersonal skill and ethical success	43
Interpersonal skill and technical aspects of ethical success	47
Non-rational aspects of interpersonal interaction	55
Ethical success and non-rational aspects of interpersonal skill	60
4. Conclusions of practical reasoning	75
Conclusions of practical reasoning are not actions	77
Conclusions of practical reasoning are intentions	90
5. Reason's guidance of action	110
Causation	115
Function	120
The management problem	129
Reasons, reasoning, and guidance	135
6. Practical warrant	145
Epistemic warrant	147
Practical justification	151

Perception, sensation, and practical entitlement	157
Reliability	167
Post-intention action guiding capacities	173
7. Entitlement for relying on physical abilities	180
An objection to Anscombe's view about knowledge of action	187
Perceptual entitlement	193
Physical abilities and agency	200
Physical abilities and practical reason	210
Bibliography	221

Acknowledgments

I want to thank my dissertation chair Tyler Burge, and my other wonderful teachers, Barbara Herman, Pamela Hieronymi, AJ Julius, Sean Kelsey, and Gavin Lawrence. I am also indebted to my fellow students, especially Denis Buehler, Sonny Elizondo, and Ingrid Steinberg. I owe my general philosophical education, and much of the material in this dissertation, to my teachers and fellow students.

I wrote this dissertation during a splendid time of life. I owe its blessings to my wife Renee, and my daughter Michele-Christine, as well as to my sister Lakshmi, my father Dev, and my friend Justi. Most of all, I am grateful to my mother Christine, to whom I dedicate this dissertation. She was the kind of person I hope to be.

Vita

Education

B.A. in Philosophy, Wesleyan University Winter 2002

Publications

“Aristotle on Choosing Virtuous Action for its Own Sake”

Pacific Philosophical Quarterly Forthcoming

Presentations

“Comments on Perin: Rational support without temporally extended reasoning”

Aristotle on Deliberation, UC Riverside May 2013

“Aristotle on Choosing Virtuous Action for its Own Sake”

NUSTEP, Northwestern University May 2013

“Perception and Reason”

Perception, Essex University May 2012

“Morality, Reason, and Interpersonal Skill”

What Makes Us Moral, VU University, Amsterdam June 2011

“A Puzzle about Production and Self Realization”

Marx and Aristotle, Institute of Education, University of London June 2011

“Reason and Skillful Coping”

Phenomenology, UWE, West Anglia October 2010

If you are all body

Why do you have this soul?

If you are all soul

Why are you covered with this shape?

Mevlâna Celâddîn Rumi

Dîvân-i-Kebîr, Meter 3, Verse 2202

Chapter 1

Introduction

This dissertation is about the way that embodied rational agency depends on non-rational capacities. It is especially concerned with the role of physical abilities in rational agency, like the physical ability to put one foot in front of the other and walk. We are embodied rational agents in that we engage in physical actions on the basis of practical reasoning. But in what way do our physical abilities figure in our agency? Many object to the image of practical reason steering the body like a captain steers a ship. Our bodies and our physical abilities seem to be *constitutively* involved in our agency, and not to be mere equipment wielded by practical reasoners.

Some attempts to understand human agency as constitutively embodied downplay the role of reason in human agency, especially those in the phenomenological tradition of Heidegger and Merleau Ponty. That tradition emphasizes the role of non-rational capacities in guiding human action. It contains valuable insights about the role of non-rational capacities in our agency. But these insights are obscured by a tendency to conceive of reason in a narrow way. Work in this tradition tends to think of all rational capacities as capacities for *legalistic* reasoning. It tends to assume that exercises of reason are algorithmic applications of rigid general rules, undertaken and acted upon with a disinterested attitude toward the situation.

A plausible conception of reason would eschew such a legalistic conception. Rational capacities are, roughly, capacities for thought and understanding, and capacities to be moved by the practical significance of considerations recognized in thought and understanding—like capacities for volition or intention. Thinking and understanding need not consist in disinterested, algorithmic applications of rigid general rules. Because of its overly narrow conception of reason, this tradition tends to overreach by assigning too marginal a role in our agency to rational capacities. Our ability to act on the basis of reasons is a core feature of our kind of agency. Phenomenology-based views tend to lose sight of this important fact.

This dissertation tries to make progress in understanding the role of non-rational capacities in our agency from a perspective that accords a central role to reason. The overarching motivation for this dissertation is to make progress in understanding the way that our *rational* agency—our ability to act on the basis of reasons—extends beyond our rational capacities themselves. It aims to shed light on the way that practical reason functions in concert with non-rational capacities, as different aspects of a unified overall action system. This is a large topic, and open-ended. There are many possible ways of approaching the topic that promise to advance understanding.

I focus in this dissertation on four main ideas that bear on the constitutive role of non-rational capacities in rational agency. First, our capacities for agential control, through which we determine how we act extend beyond our rational capacities for practical reasoning and volition. We control our actions, including our intentional actions, partly through the exercise of non-rational action-guiding capacities, like capacities for motor control.

Second, intentional actions, including physical actions, are guided by practical reasoning in much the same way intentions are, even though performing an intentional action is not part of the agent’s practical reasoning. Practical reason can guide intention on its own, without the exercise of any capacities outside of one’s capacity for practical reasoning. However, in guiding action, practical reason “reaches past itself” in a way that its guidance of intention

does not. Since action is not contained in one's practical reasoning, practical reason can only guide intention as part of a larger action system. This larger action system crucially includes the physical abilities that we rely on to execute most of our actions. Practical reason's guidance of action requires the exercises of the action-executing abilities in roughly the same way that its guidance of intention requires the exercise of one's capacity to form intentions. In this way, our success as practical reasoners depends not just on our capacities for practical thought and understanding, but also on the larger action system that those capacities function as a part of, including our physical abilities.

The third main idea of the dissertation tries to get at the relation between practical reason and non-rational aspects of agency by looking at the role of non-rational capacities in warranting our actions and intentions. Whether our intentions and action are warranted depends both on exercises of our capacities for practical reasoning, and also on the exercise of non-rational capacities like perception, sensation, and motor control. Warrant for our actions and intentions also depends on the general reliability of our action systems as a whole, including non-rational capacities as well as practical reason.

The fourth main idea is that we have a default entitlement for relying on our agency as a whole, including both practical reason, and the physical abilities involved in carrying out physical actions. We do not need evidence about the effectiveness of practical reason, or our physical abilities, in order to be warranted in relying on those aspects of our agency. We have a default entitlement for relying on our physical abilities because of their constitutive role in our agency.

I develop these ideas over six chapters. I summarize these chapters in the remainder of this introduction. Chapter One criticizes two closely connected rationalist views about human agency. The first of these views, rationalism about *agential control*, claims that the capacities for agential control in normal adult human beings are rational capacities. The second view, rationalism about *action* is a circumscribed version of the first: certain

capacities for agential control provide the “spark” of agency that makes the things we do count as actions. Rationalism about action claims that *those* capacities for agential control are rational capacities, though it does not require that all capacities for agential control be rational capacities.

The arguments of this chapter focus on aspects of technical skills that control details of skillful action, like the details of a violinist’s finger movements. I discuss a variety of considerations that suggest that the capacities that determine many such details are non-rational capacities. First, the details of skilled behavior, and the psychological processes that determine them, are often not accessible to conscious reflection by the agent. Second, these capacities often have to be acquired through practice, and not just through an intellectual grasp of the details of skillful action. Third, mistakes in the details of attempts at skilled action often do not open the agent to rational criticism. Mistakes in the details of skillful action often open the agent to criticism as clumsy or inept, without warranting criticism as irrational, unreasonable, foolish, or stupid. The view finds a fourth source of support in the distinction drawn by psychologists between declarative and non-declarative memory systems. Psychologists hold that skilled action is accomplished largely through “instructions” for how to carry out an action that are stored in non-declarative memory systems. These memory systems seem to be non-rational capacities, since the information they store is not available to conscious awareness, is available for use only in a circumscribed range of cases, and the mental states they “contain” are not assessable as true or false.

I argue that, even though the capacities that determine details of technically skilled action are non-rational capacities, they are capacities for agential control. First, the skills constituted partly by non-rational capacities seem to be appropriate objects of pride and admiration. Action guiding capacities that warrant such attitudes would seem, at least *prima facie*, to be capacities for agential control. Second, one can be directly and non-derivatively at fault for failures in the operation of these non-rational capacities. For example, one can

be directly and non-derivatively at fault for flubbing a musical performance through the failure of non-rational capacities, and one can appropriately apologize to other members of the ensemble for such an error. (Though, to be clear, I do not claim that the error in such cases is *moral* error.) Third, limitations in, and interference with, one's non-rational capacities can excuse an agent from being at fault for wayward attempts at skillful action. I argue that limitations and interference with non-rational capacities would not excuse fault if the relevant capacities were not capacities for agential control.

These arguments oppose rationalist views about agential control. The final arguments in the chapter target rationalist views about action. The details of an individual's way of carrying out an action are themselves manifestations of agency. Exercises of rational capacities leave many integral details of one's way of acting unspecified. As a result, I argue, exercises of rational capacities do not suffice to make those details count as manifestations of agency. They count as manifestations of agency at least partly in virtue of the exercises of non-rational capacities that specify the details.

Chapter Two explores the role of interpersonal skill in ethically good action, with an emphasis on non-rational aspects of interpersonal skill. This chapter makes use of the discussion of non-rational aspects of agency in Chapter One.

I begin by drawing a distinction between "valuational" and "technical" aspects of ethical success. Valuational aspects of ethical success consist, roughly, in having ethically good values and attitudes toward others. For example, treating a grieving friend in an ethically good way requires that one value kindness and compassion, that one care about one's friend, and that one intend to treat the friend in an appropriately compassionate way. But successfully treating the grieving friend in an ethically good way also requires that one do a good job in determining how in particular to act toward the friend, like whether to try to cheer him up with a joke, or whether to sit quietly with him. Determining specifics like these are technical aspects of ethical success. Technical aspects of ethical success consist in doing well

in determining how specifically to behave, so as to bring good values and attitudes to bear in appropriate ways.

I go on to suggest that interpersonal skills, like skill at treating others in loving and compassionate ways, consist partly in non-rational action-guiding aspects of agential control. For example one's ability to treat others in appropriately loving or compassionate ways consists partly in non-rational capacities which determine details of tone of voice, mode of eye contact, body posture, and so on.

Finally, I suggest that these non-rational capacities contribute constitutively to technical aspects of ethical success. I make four points to support this view. First, I distinguish the role of non-rational aspects of interpersonal skill in ethically good action from the role of non-rational aspects of the ability to swim (for example) in the ethically good action of saving a drowning child. The exercise of interpersonal skill helps give the action its loving or compassionate tenor, where the action is ethically good partly in virtue of its being loving or compassionate. No similar point applies to the ability to swim. Second, I argue that we can be culpable for failures to treat others well through failures in the exercise of non-rational aspects of interpersonal skill, since those non-rational capacities are capacities for agential control. Third, I suggest that details of interaction determined by non-rational capacities can appropriately occasion Strawson-style "reactive attitudes." For example, a loving way of looking at one's spouse can appropriately occasion gratitude for being treated in a loving way. Such details of behavior can warrant reactive attitudes because they are aspects of treating others well that are determined by capacities for agential control. Fourth, I argue that details of the way one executes an intention to act compassionately, for example, are integral to the ethical goodness of one's way of acting toward other people. But many such details are left open by exercises of the agent's rational capacities. As a result, those details have to count as ethically good partly in virtue of exercises of the non-rational capacities that specify the details. I put most weight on this last argument.

In technical skill and interpersonal skill, non-rational aspects of agency operate in concert with our rational capacities. The exercise of these non-rational capacities is often part of one's way of bringing reason-based values, plans, and intentions to bear in action. The remaining chapters of the dissertation aim to shed light on the way in which practical reason operates together with non-rational capacities as different aspects of one's overall action system—and the way that our success as practical reasoners depends constitutively on non-rational aspects of one's action system.

The claim that actions are conclusions of practical reasoning, often attributed to Aristotle, suggests a way of approaching this topic. If actions are conclusions of practical reasoning, then the non-rational, physical abilities involved in carrying out an action are integral to the functioning of practical reasoning. Their exercise is necessary to the successful functioning of practical reasoning in the same way that the capacity to form a belief is necessary to the successful functioning of epistemic reasoning. While I ultimately suggest (in Chapter Four) that there is an element of truth in this proposal, there is a strong case to be made against the claim that actions are conclusions of practical reasoning. Chapter Three argues that actions are not conclusions of practical reasoning and that intentions are.

The view that intentions are conclusions of practical reasoning is more widespread than the view that actions are. But there are few sustained arguments for the claim that intentions are conclusions of practical reasoning. There are also few sustained arguments against the claim that the conclusions are actions. Most extant arguments against the view that actions are conclusions of practical reasoning seem to me to miss the mark, and to neglect the core problems with the claim that actions are conclusions of practical reasoning. In the first half of the chapter, I argue that actions are not conclusions of practical reasoning. In the second half, I defend the view that intentions are.

The arguments I present against the claim that actions are conclusions of practical reasoning center on two main points. First, the claim that actions are conclusions of reasoning,

including physical actions, is incompatible with the point that reasoning is a psychological phenomenon. A conclusion of practical reasoning is an element in the reasoning, and so must be an element in a psychological phenomenon. But physical actions are not, or not just, psychological phenomena.

Second, I suggest that reasoning, including both practical and epistemic reasoning, consists, roughly, in “figuring something out” through making use of reasons. As a result, I argue, a conclusion of reasoning must be some kind of committal attitude, some kind of “making up one’s mind,” whereby one purports to have figured something out for oneself. Epistemic reasoning is a kind of figuring out what is the case, and practical reasoning is a kind of figuring out what to do. Accordingly, a conclusion of practical reasoning is a committal attitude, whereby one purports to have figured out what to do. It is one thing to figure out whether to carry out some action, like going for a walk. It is another thing actually to carry out the action, and go for the walk. Thus, an action like going for a walk is not a conclusion of practical reasoning.

In the second half of Chapter Two, I defend the view that intentions are conclusions of reasoning against two sets of objections, due to Sergio Tenenbaum and Joseph Raz, respectively. Tenenbaum argues that actions must be conclusions of practical reasoning, since whether an action is justified is underdetermined by reasoning to an intention. Tenenbaum presents an example where an individual walks across the room to turn on the light, and shoves his computer out of the way so as to take the quickest route to the switch. The individual’s intention to switch on the light is rational, but his action is irrational. Tenenbaum claims that the irrationality of the action is not due to a failure in reasoning to an intention, and is due to the action itself. As a result, the action is part of the agent’s practical reasoning, and can be the conclusion.

Tenenbaum’s example is unconvincing. The example he presents is clearly a case of acting on an intention that was based on poor reasoning. The rationally defective intention is not

the intention to switch the light on, but the intention to shove the computer out of the way. I argue that any mistake in action that is criticizable as irrational reflects some psychological error by the agent in determining how to act.

In the remainder of the chapter, I criticize three arguments that Raz makes against the claim that intentions are conclusions of practical reasoning. Raz's arguments are intended to support his view that conclusions of practical reasoning are not actions, but judgments about what one has reason to do, or about what one must to do.

First, Raz argues that, if intentions are conclusions of practical reasoning, then whether one's reasoning counts as *practical* reasoning depends on whether one forms an intention at the end of the reasoning. He regards the latter claim as an unacceptable consequence of the view that conclusions of practical reasoning are intentions. I argue, first, that this claim does not follow from the view that intentions are conclusions of practical reasoning. Reasoning can *function* to issue in intention antecedently to actually issuing in intention. Reasoning that functions to issue in intention is thereby practical, even if the reasoner never actually forms an intention. Second, I argue that whether a piece of reasoning is practical really can depend on later stages of the reasoning, even according to Raz's own view. Patterns of epistemic reasoning can be elements of practical reasoning in some contexts, and can be non-practical epistemic reasoning in other contexts. Whether epistemic reasoning is an aspect of one's practical reasoning can depend on later "stages" of the reasoning.

Raz's second argument is that intention cannot be the conclusion of practical reasoning because weakness of will can consist in a failure to form an intention, and weakness of will is not a failure of reasoning. I suggest that, at least in the relevant cases, weakness of will is associated with a failure of reasoning. Intention is a committal attitude based on reasons, and a failure to form such a commitment can be considered a failure in reasoning. At least, Raz offers no argument against this view.

Third, Raz argues that a transition from a judgment that one has reason to ϕ to an

intention to ϕ is not a step in reasoning, since it is like reasoning from a belief that p to a belief that p. Raz's idea seems to be that such a transition is not a substantive step in reasoning, whereby one figures out something new. Since the transition is not a substantive step in reasoning, the resulting intention is not a conclusion of reasoning.

I make three points in response to this argument. First, an intention need not be based on a reason-to- ϕ judgment to be based on reasons. Raz claims that some such judgment is required to distinguish between merely having a reason that *could* ground an intention, and actually making use of a reason to ground an intention. I argue that this is an overly intellectualist claim about what is required to make use of practical reasons.

The second point I make is that a transition from a reason-to- ϕ belief to an intention to ϕ is quite different from inferring that p on the basis of a belief that p. To make this point, I draw on an analogy to critical epistemic reasoning. In critical epistemic reasoning, one can reason to a belief that p from a belief that one has reason to believe that p. That transition is a step in one's epistemic reasoning. The transition from a reason-to- ϕ belief to an intention to ϕ seems, at least *prima facie*, to be more like that kind of step in epistemic reasoning than like "inferring" that p on the basis of a belief that p. I also point out that there is a substantive difference in the commitment involved in holding the intention, and the commitment involved in the ought judgment. I suggest, on those grounds, that the transition should be considered a substantive step in reasoning.

The third point is that one can still plausibly count intentions as a conclusions of reasoning, even if forming an intention is posterior to the last substantive step in one's reasoning. To make this point, I consider quasi-transitions in memory, like the transition from a belief that the moon is full today to a belief that the moon was full yesterday. Suppose the initial today-belief is formed in a substantive step in epistemic reasoning. The transition in memory is not a substantive step in reasoning. So the yesterday-belief is posterior to the last step in reasoning. Nonetheless, the yesterday-belief is a commitment that is based on

reasons in the agent's psychology. I suggest that there is a useful conception of a conclusion of reasoning according to which such a commitment is a conclusion of reasoning—even if it is posterior to the last substantive step in reasoning. Even if intentions are posterior to the last substantive step in reasoning, they are still commitments that are based on reasons in the agent's psychology. Thus, they can still count as conclusions of practical reasoning.

Actions are represented in the content of intentions. So the point that practical reasoning concludes with intention shows one way that actions figure in practical reasoning as subject matter. But the point leaves unaddressed the way that actions figure in practical reasoning as something more than the characteristic subject matter of practical reasoning. Practical reason does not just guide our intentions, but also guides our actions. Chapter Four presents an account of the way that actions are guided by practical reasoning, even though they are not conclusions of reasoning.

The account is anchored by two main considerations. First, the account is meant to get at a guidance relation between practical reasoning and action which practical reasoning also bears to intention, and which epistemic reasoning bears to belief. Second, the account is meant to help explain why the quality of practical reasoning bears on the rationality of actions—as it also bears on the rationality of intention, and as the quality of epistemic reasoning bears on the rationality of belief.

The account has four main components. First, I suggest that actions count as guided by practical reasoning partly because practical reasoning figures in causal-psychological explanations of why we act as we do. I suggest that, if our actions (or intentions or beliefs) are undertaken (or formed) in a way that is psychologically independent of any reasoning, then reasoning does not bear on the rationality of the action (or intention or belief). However, the fact that reasoning helps cause actions, intentions, a beliefs is far from sufficient to explain why reasoning rationally supports them. Strenuous reasoning can cause a headache, but the reasoning does not rationally support the headache.

In the second part of the account, I suggest that actions, intentions, and beliefs count as guided by reasoning partly because practical reasoning *functions* to causally determine our actions, intentions, and beliefs. That function subserves our action's, intention's, and belief's satisfying a standard of correctness. The standard on belief is truth. Roughly, the standard on action and intention is that the intended action be good to do. Reasoning does not issue in action, intention, and belief *per accidens*. It is a central part of epistemic reason's "job" to issue in true belief. Likewise, it is a central part of practical reason's "job" to issue intentions and actions that satisfy their standards of correctness. To defend this suggestion I make use of Tyler Burge's claim that success conditions are associated with functions. Rationality of actions, intentions, and beliefs is a kind of success in the way our actions, intentions, and beliefs are determined through reasoning. Accordingly, reasoning must determine our actions, intentions, and beliefs according to its function. I also suggest that it is a conceptual truth that guidance always is associated with a function that subserves something like a standard of correctness (a *telos*) for the object of guidance.

In the third part of the account, I argue that actions count as guided by practical reasoning partly because the practical reasoning and the action are different aspects of the exercise of an over-arching ability. This discussion helps to distinguish practical reasoning's relation to actions from its relation to certain objects of what Pamela Hieronymi calls "managerial control," like the behavior of a trained dog.

The fourth and final part of the account is the core point of the chapter. This part of the account aims to show that in practical reasoning one brings reasons to bear on action in the same way that one brings reasons to bear on intentions—and the same way one brings reasons to bear on beliefs in epistemic reasoning. In epistemic reasoning, one makes use of reasons to show that the belief the reasoning satisfies its standard of correctness (i.e., is true). Analogously, in practical reasoning, one makes use of reasons to show that one's intention satisfies its standard of correctness. At the same time, in practical reasoning one

also makes use of reasons to establish that one's action satisfies its standard of correctness. One brings reasons to bear on both intention and action because the standard of correctness on intention is the same as the standard of correctness on the intended action. The standard for both is, roughly, that the intended action be good to do. The reasons in one's practical reasoning function to show that one's intended action is good to do. In that way, the reasons function to show that both the intention and the action satisfy their standards of correctness. One makes use of reasons in practical reasoning to get one's intention right, and to get one's action right, in the same way. Actions, intentions, and beliefs are guided by reasoning in that reasoning functions to issue in correct action, intention, and belief, through making use of reasons that show that they are correct—where correctness for action and intention consists in the intended action's being good to do, and correctness for belief consists in the belief's being true.

Drawing on the discussion of practical reason's function in Chapter Four, Chapter Five presents accounts of two types of practical warrant: rational justification for actions and intentions, and entitlement for actions and intentions. While rational justification for actions and intentions is widely discussed, there is no work that I know of about aspects of warrant for actions and intentions beyond an agent's rational justification. I use the term 'entitlement' for such aspects of warrant for actions and intentions. This kind of practical entitlement is the primary focus of Chapter Four.

I begin with a review of Tyler Burge's accounts of epistemic justification and epistemic entitlement, which provide a template for my accounts of practical justification and practical entitlement. Epistemic justification consists, roughly, in epistemic reason's success in guiding belief, given limitations in the agent's perspective and epistemic capacities. There are two kinds of epistemic entitlement. The first kind consists in the success of non-rational aspects of one's belief system in guiding belief, like capacities for perception and sensation. The second kind consists in characteristics of one's belief system that make the system count

as a genuine epistemic *competence*—most importantly, this kind of entitlement is provided by the reliability of different aspects of one’s belief system in guiding the individual to true belief.

I then present my account of practical justification. I suggest that an intention’s practical justification consists in practical reason’s success in guiding intention, given limitations in the agent’s perspective and rational capacities. An action has practical justification just in case it is an execution, or attempted execution, of a justified intention.

Next, I discuss two different kinds of practical entitlement. The first consists in the guidance of action by non-rational aspects of the agent’s action system. Some such capacities, like perception and sensation, carry out their action-guiding function by helping to guide the agent’s intentions. I suggest that practical justification depends in a constitutive way on this kind practical entitlement, although practical justification only *consists* in guidance by rational capacities. Other non-rational action-guiding capacities, like motor skills, operate “downstream” of intention, in the execution of intended actions. Practical justification does not depend on exercises of these capacities. The practical warrant they provide is a kind of practical entitlement.

A second kind of entitlement consists in practically good characteristics of our action systems which make our action systems count as a genuine practical competence. Epistemic warrant is a matter of epistemically good exercises of epistemic competence, and not least-bad exercises of epistemic incompetence. Similarly, practical warrant consists in practically good exercises of competent agency, and not least-bad exercises of incompetence. I focus on the reliability of different aspects of one’s action system in successfully guiding the agent to good action. Our action guiding capacities count as practical competences in virtue of their general reliability in facilitating effective agency. If you cross the road, the perceptual and rational capacities involved in guiding your action count as aspects of a practical competence because they work well in general. For example, your perception of oncoming traffic, and

your reasoning about when to cross, counts as aspects of a practical competence because the relevant capacities work reliably according to their design, and do not only work well occasionally and by accident. As a result, the reliability of these action-guiding capacities contributes to practical warrant. The kind of warrant that they contribute to is entitlement, and not justification, since it does not consist in justification by reasons. The reliability of practical reason, of non-rational intention-guiding capacities like perception and sensation, and of post-intention action-guiding capacities like motor skills, all contribute to practical warrant.

I suggest that both of these kinds of practical entitlement conflict with practical analogs of internalist claims about epistemic warrant. Epistemic internalists commonly make three claims. First, any factor that provides epistemic warrant for beliefs must be an exercise of rational capacities to make use of reasons. Second, any such factor must be a consciousness mental state, or accessible for use in conscious reasoning. Third, the individual must be directly responsible for any factor that helps warrant beliefs. I draw on the discussion of practical entitlement to criticize analogous internalist claims about factors that provide practical warrant.

Let me briefly canvas the main points I make against these internalist claims about practical warrant. The first few points concern perception and sensation. First, we are not directly responsible for our perceptions and sensations, even though perceptions and sensations do help warrant our actions and intentions. Second, some sensations and perceptions that help warrant our actions and intentions are not conscious mental states. For example, the kinesthetic sensory registrations that one makes use of in tracking the movements of one's body are often non-conscious. Third, sensation is almost generally regarded as a non-rational capacity. I sketch some considerations suggesting that perception is a non-rational capacity, though that issue is contested. The last few anti-internalist points I make focus on the reliability of our action guiding capacities. First, we are not directly responsible for the

reliability of our action-guiding capacities. Second, reliability is not a conscious mental state. Third, reliability is a characteristic of one's action guiding capacities, and not an exercise of those capacities. Thus, reliability is not an exercise of rational capacities to make use of reasons. Nor can reliability be an element in conscious reasoning.

Chapter Six discusses a kind of practical warrant not discussed in chapter Five. It discusses practical warrant for relying on one's physical abilities. Normally, when one performs some physical action, one relies on the physical abilities through which one carries out the action. When one grasps a cup, or plays Bach's first cello suite, one relies on the physical abilities involved in performing those actions. Chapter Six argues that we have a defeasible default entitlement for this kind of reliance. In cases when an agent has no defeating evidence that her physical abilities are not to be relied on, she has a default, non-observational warrant to rely on her physical abilities. For example, warrant for relying one's ability to walk need not depend for its force on evidence that one is able to walk, or that one is successfully making use of one's ability to walk.

This claim is a variant of Anscombe's claim that knowledge of physical action is non-observational. Traditionally, we are thought to have non-observational warrant concerning self-evident truths, conceptual truths, and facts about our own psychologies. A provocative feature of Anscombe's view is that it claims we have non-observational warrant concerning overt physical action. The variant of Anscombe's view I defend also holds that we can have non-observational concerning physical actions.

There is a reasonable *prima facie* case to be made against the claim that we can have such knowledge of our physical actions. Whether you are successfully carrying out a physical action, like extending your arm, or juggling, or wiggling your ears, is a contingent, extra-mental matter of fact. It depends on what body parts you have, on whether you have the relevant motor skills, whether you are tied up or not, and so on. Except in special cases, one does not have apriori knowledge of contingent, extra-mental matters. Knowledge of whether

it is raining or sunny, for example, must be based on observational grounds. Knowledge about such matters depends on observational registrations of the world, which provide information about what is going on in the world. It is not clear why we should have license to believe we are successfully acting as intended without grounding by observational information that indicates what is actually going on with our physical bodies.

My account of our default entitlement to rely on physical abilities aims to address worries of this kind. Warrant for reliance on physical abilities does not require observational evidence of the reliability and proper functioning of our physical abilities. I model my account of a default entitlement to rely on physical abilities on Tyler Burge's account of our epistemic entitlement to rely on perception. Roughly, according to his account, we are epistemically entitled to rely on reliably accurate perceptual capacities because reliably accurate perception is partly constitutive of the normal proper functioning of our belief-guiding capacities, and not accidental with respect to their proper normal functioning. Similarly, I suggest that we are non-observationally entitled to rely on physical abilities because our reliance on physical abilities, and their reliably effective functioning, is partly constitutive of the normal proper functioning of our agency. The reliably effective functioning of our physical abilities is non-accidental with respect to the normal proper functioning of our agency. In this way, our reliance on physical abilities is not like a "shot in the dark."

I add to the case for the claim that we are non-observationally entitled to rely on our physical abilities by considering the relation between physical abilities and practical reason. As discussed in Chapters Three and Four, practical reason has a core function of guiding action, in addition to its narrower function of guiding intention. I suggest that the exercise of our physical abilities is integral to the full action-guiding function of practical reason. Without physical abilities, practical reason could guide intention, but it could not guide action. I suggest that we have a default entitlement to rely on practical reason to carry out its full action-guiding function. Since our physical abilities are integral to that function, this

entitlement to rely on practical reason to guide action implies that we have a corresponding default entitlement to rely on our physical abilities. Both entitlements are non-observational.

I conclude by arguing that these entitlements are not due exclusively to the nature and functioning of rational capacities. In this way, the entitlements are not apriori, even though they are non-observational. The entitlements are due in part to the reliable effectiveness of our physical abilities. I suggest that this conclusion conflicts with certain rationalist views about the application of rational norms on action. Rational norms apply to action in that we genuinely ought to act on the basis of good reasons. Some rationalists claim that rational norms apply to actions apriori, just in virtue of our status as practical reasoners, and the nature of our capacities for practical reasoning. (Other rationalists only claim that rational norms apply apriori to intentions. The discussion in this chapter does not bear on that issue.) The grounds for our entitlement to rely on practical reason to guide action conflicts with this rationalist claim about rational norms on action. I argue that rational norms apply to actions only if we are warranted in relying on practical reasoning to guide action. If warrant for that reliance is not apriori, then rational norms do not apply to our actions apriori.

Chapter 2

Non-rational aspects of agency

Introduction

It is common to conceive of human agency as consisting in the rational control we exert over our doings. According to this view, we are active agents, and in control of our doings, through exercises of our rational capacities for thought, understanding, and volition. This rationalist conception of human agency goes hand in hand with a rationalist view about human nature. What typifies us as the kind of being we are is reason. If what we *are* is rational beings, it makes sense that our control over our actions would be rational control, control by our capacities to respond to reasons in thought, understanding, and volition. Traditionalist rationalist conceptions of our agency neglect the important role of non-rational capacities in our agency. This chapter argues that our agency, our control over what we do, consists partly in non-rational action-guiding capacities—capacities to determine how we conduct ourselves outside of our capacities for thought, understanding, and volition.

There have been two main approaches to challenging traditional rationalist views about human agency. One approach emphasizes the role of emotion and desire in human action, where reason is thought to contrast with emotion and desire. But reason does not contrast

with emotion and desire in a straightforward way. Emotions and desires can be based on operative reasons, and they can be subject to standards of rationality and irrationality. An individual might be angry at his spouse on the basis of his belief that she had been unfaithful. And, also on the basis of that belief, he might desire to end the marriage. That anger and that desire may be irrational if, for example, the agent's belief that his spouse had been unfaithful is irrational. These points suggest, at least *prima facie*, that emotions and desires can be exercises of rational capacities.¹ So in acting on the basis emotion and desire, one may be acting under the guidance of reason.²

The second approach, with roots in the phenomenology of Heidegger and Merleau Ponty, emphasizes the role of competences that operate below the level of reason in guiding human action.³ This line of thought contains valuable insights about the role of non-rational capacities in human action. But these insights are obscured by a tendency to conceive of reason in a narrow way. Work in this tradition tends to mistake reasoning for *legalistic* reasoning.⁴ It tends to assume that exercises of reason are algorithmic applications of rigid general rules, undertaken with a disinterested attitude toward the situation. A plausible rationalist view would eschew such a legalistic conception of reason. Rational capacities are, roughly, capacities for thought and understanding, and capacities to be moved by the practical significance

¹See, e.g., De Sousa (1979), Gilligan (1985), Nussbaum (1992), Goldie (2000), Scanlon (1998), and Damasio (1999) for discussions of the connection between reason and emotion and desire.

²Hursthouse (1991) argues that an action can be performed as an expression of emotion. She argues that such actions can count as intentional, even though they are not motivated by belief-desire pairs of the sort described by Davidson (1980). Though she describes such actions as “arational,” her view is compatible with rationalist conceptions of agency. Accepting her criticism of Davidson's view, one might argue that such emotion-expressing behavior manifests agency in that the behavior executes a conceptualized intention, e.g., to smash the vase—even if the intention is not based on a “rationalizing” pair of beliefs and desires. Or, alternatively, one might argue that the emotion-expressing action manifests agency in that the emotion that is expressed by the action is guided by reason.

³Dreyfus is the flag bearer for this approach. See, e.g., (2002), (2006), (2007). Also see, e.g., Kelly (2000), Ennen (2003) and Rietveld (2008) for kindred discussions of skillful action.

⁴The gerundive form of the word ‘reasoning’ suggests that it refers to an interruptible process. I use the word to refer to exercises of rational capacities to hold a commitment (like a belief or an intention) on the basis of reasons. I do not assume that exercises of these capacities must be interruptible processes.

of considerations recognized in thought and understanding—like capacities for volition or intention. Thinking and understanding need not consist in disinterested, algorithmic applications of rigid general rules.⁵ Because of its overly narrow conception of reason, this approach tends to overreach by assigning too marginal a role to our rational capacities in guiding action. Our ability to act on the basis of reasons is a core feature of our kind of agency. Phenomenology-based views tend to lose sight of this important fact. This chapter aims to oppose overly rationalist views about human agency, but without losing sight of the important role that reason has in human agency, and the different forms that exercises of reason can take.

The chapter criticizes two closely connected rationalist views about human agency. The first of these views, rationalism about *agential control*, claims that the capacities for agential control in normal adult human beings are rational capacities. I argue that non-rational capacities are among our capacities for agential control. The second view, rationalism about *action*, claims that the capacities for agential control in virtue of which the things we do count as our actions are rational capacities. I argue that non-rational aspects of agential control can help make an event count as an action. I characterize these two rationalist views about agency in more detail in section 2.

The arguments of this chapter focus on aspects of technical skills that control essential details of skillful action, like the details of a violinist's finger movements. These aspects of technical skills are largely non-rational capacities. They operate below the level of thought and understanding. I discuss a variety of considerations that suggest these capacities are non-rational capacities. First, many details of skilled behavior, and the psychological processes that determine them, are inaccessible to conscious reflection by the agent. Second, these capacities often have to be acquired through practice, and not through an intellectual grasp of

⁵See, e.g., McDowell (2007b) and (1979), Herman (2007), Gilligan (1985), Wiggins (1975 - 1976) for resistance to legalistic conceptions of reason.

the details of skillful action. Third, mistakes in the details of attempts at skilled action often do not open the agent to rational criticism. Mistakes in the details of skillful action often open the agent to criticism as clumsy or inept, without warranting criticism as irrational, unreasonable, foolish, or stupid. Fourth, empirical research on non-declarative memory and motor control suggests that details of physically skilled action are determined largely by psychological states that are available for use only in a circumscribed range of tasks, are not assessable for truth or falsity, and are not accessible to consciousness. The capacities associated with these psychological states seem to be non-rational capacities.

Even though the capacities that determine details of technically skilled action are non-rational capacities, they are capacities for agential control. These capacities are integral, constitutive aspects of technical skills. A violinist who lacks capacities that control the details of finger placement needed for good intonation is thereby lacking in skill. Unlike, say, the ability to digest food or regulate one's heartbeat, an individual's ability to play the violin in tune is attributable to the agent herself. It is *she* herself who plays in tune. Playing in tune seems to be an exercise of agency. But her playing in tune is largely a matter of the details of finger placement that are determined by non-rational capacities. Thus, those non-rational capacities seem to be capacities for agential control.

The claim that they are capacities for agential control finds support in their relation to certain evaluative attitudes. First, the skills constituted by non-rational capacities seem to be appropriate objects of pride and admiration. Second, one can be at fault for failures in the operation of these non-rational capacities. One can be at fault for flubbing a performance through the failure of non-rational capacities, and one can appropriately apologize to other members of the ensemble. (Though, to be clear, I do not claim that the error in such cases is *moral* error.) Third, limitations in one's non-rational capacities can excuse an agent from being at fault for wayward attempts at skillful action. I argue that they would not do so if they were not capacities for agential control.

I develop these arguments in section three, to criticize rationalism about agential control. In section four, I build on these arguments to criticize rationalism about action. But before turning to those arguments, I want to present a more detailed characterization of the two rationalist views about agency that I will be criticizing.

Two rationalist views about agency

The first of the two rationalist views is about agential control. Capacities for agential control are capacities through which we determine, as agents, how we act.⁶ Our rational capacities for agential control include our capacity to determine what we shall do in forming conceptualized intentions to act, and our capacity to form intentions on the basis of reasons, among others. Rationalism about agential control claims that our ability to determine as agents how we act consists solely in guidance of our actions through the exercise of rational capacities like these.

Some remarks McDowell makes help to bring out the main thrust of this rationalist view about agential control. He says, “the capacities that are operative in ordinary perceptual engagement with the world, and in *ordinary bodily action*, belong to a subject’s rationality in that strong sense,” where rationality in the strong sense is “responsiveness to reasons as such” (2007a, p. 366, emphasis added).⁷ Raz also expresses a rationalist view about agential control in claiming that it is through reason that “we direct our lives, we are in control” (1997, p. 227). And Velleman gives voice to a rationalist view about agential control in claiming that rational capacities that cause an agent’s behavior are the “locus of his agency.” He says, “As Aristotle put it, ‘Each person seems to be his own understanding.’”

⁶One might point out that a kind of epistemic agential control is involved in theoretical reasoning, and not just reasoning that determines how we act. I set aside that kind of epistemic agential control here.

⁷Also compare McDowell’s claim that human beings are animals “whose natural being is permeated with rationality” (1996, p. 85), and that “conceptual rationality is everywhere in our lives, in so far as our lives are distinctively human” (2007b, p. 349).

Hence causation via a person's rational faculties qualifies as causation by the person himself" (2000, p. 6). When exercises of non-rational capacities are integral to the determination of one's behavior, the agent herself is not determining what she does in making use of those capacities, since the agent is not to be identified with her non-rational capacities.

This rationalist view about agential control does not hold that all endogenous capacities to determine activity are rational capacities. First, the view is about *human* agency, and does not claim that action guiding capacities of non-human animals are rational capacities. Second, the claim does not concern human capacities that regulate digestion, normal blinking, heartbeat, and basic flinching reflexes. It is difficult to characterize exactly which human capacities are supposed to count as capacities for agential control. Roughly, capacities for agential control are capacities to guide action, through which the agent *herself* determines what she does. The rationalist view about human agential control is that, when we ourselves control what we do, our control consists in the exercise of rational capacities.

The rationalist view about agential control goes hand in hand with a rationalist view about the nature of action. Rationalism about action advocates a rationalist solution to what Frankfurt calls "the problem of action." In Frankfurt's words, "The problem of action is to explicate the contrast between what an agent does and what merely happens to him, or between the bodily movements that he makes and those that occur without his making them" (1978, p. 1). The problem of action asks after the way in which characteristically human actions are produced by their agents, which makes them count as actions. The rationalist answer to this question is that characteristically human actions count as actions because they are produced through the exercise of certain rational capacities.

Rationalism about the nature of human action is a restricted version of rationalism about agential control. What makes an event count as an action is some relation between the event and an exercise of capacities for agential control. It is that element of control in virtue of which the event is something the agent made happen. The rationalist view about action is

that *those* capacities for agential control are rational capacities. Rationalism about action does not claim that *all* capacities for agential control are rational capacities. It makes the weaker claim that those capacities for agential control whose exercise helps make an event count as an action are rational capacities. As it were, the “spark of agency” that makes an event count as an action is claimed to consist exclusively in the exercise of rational capacities.

A strong version of rationalism about action claims that *any* event rightly described as an action counts as an action because of its relation to exercises of rational capacities—including animal behavior. This very strong view requires that the behavior of pre-rational animals does not count as action.⁸ But accounts of the nature of action usually have a more narrowly delimited subject matter, seeking primarily to explain human action. Some human behavior is outside the scope of these accounts, like Freudian motivated slip-ups, or non-intentionally drumming one’s fingers on the table.⁹ These sorts of behavior are thought to fall short of being “full-fledged” human action, even though they are actively produced. The rationalist view about action that is at issue in this chapter is not about marginal cases of human activity like slips and fidgets.

Accounts of the nature of human action are often given in terms of an event’s relation to the agent’s rational capacities. Prominent examples include Davidson’s, Velleman’s, and Bratman’s accounts of action. Davidson’s (1980) proposal is that an event is an action because it is intentional under some description; and it is intentional because it is done on the basis of reasons that rationalize the action. Velleman’s (2000) account is that an event is a “full-blooded” action because it is done on the basis of reasons that constitute the agent’s understanding of why it makes sense to do what she does. Bratman’s (2007) account is that an individual is the “full blown” agent of an action because the action is based on a desire,

⁸See Burge (2009b) for criticism of this strong rationalist view.

⁹For discussion of human activity that does not count as “full fledged” action see, e.g., Frankfurt (1988), p. 58, and Velleman (2000), pp. 2-4.

and basing the action on the desire is a result of the agent's ongoing policy of treating desires of that sort as justifying action.¹⁰ Capacities to act on the basis of reasons, to understand the rationale for acting in certain ways, and to determine one's actions through meta-level prescriptions concerning practical justification, are all rational capacities. So according to each of these views, what makes an event count as an action is its relation to the exercise of rational capacities. These accounts of action do not assign any role to non-rational capacities in making the things we do count as actions.

Criticizing rationalism about agential control

I turn now to arguing against rationalism about agential control. The main idea of the overall argument is that technical skills consist partly in non-rational action-guiding capacities, and are capacities for agential control. I develop the argument in two parts. First, I argue that technical skills can consist partly in non-rational capacities. Second, I argue that these non-rational aspects of technical skills are capacities for agential control.

Non-rational aspects of technical skills

The kinds of technical skills I have in mind are physical skills, like those of musicians, sculptors, dancers, athletes, and so on. Technical skills are “know-how” aspects of what makes a person count as, for example, a good dancer, or a good athlete. They determine certain aspects of the agent's skillful way of acting, like the specific location of a violinist's fingers.¹¹

¹⁰The policy also has to be “reflexive,” and the agent has to be “satisfied” with the policy. A policy is reflexive when it is part of the content of the policy that the policy is to be enacted. An individual is satisfied with a policy when the policy does not conflict with any other policy the agent has.

¹¹The ability of skillful practitioners comprises more than technical skill. Being a good dancer or athlete does not consist solely in capacities to determine how one acts. Being good at these sorts of things can consist partly in “purely physical” attributes, like flexibility and strength. Serena Williams's strength, for

Technical skill can be partly constituted by non-rational capacities. Before presenting my arguments, I want to note that it is *prima facie* plausible that technical skills operate in large part below the level of thought and understanding. The point is conceded even by McDowell—a concession which seems in tension with his rationalist conception of agency. Speaking of the technique involved in throwing a baseball properly, McDowell says, “the skill itself provides for the movements to be as they needed to be...without the agent’s means-end rationality being called on to intervene” (2007a, p. 367).

McDowell’s remark accords with common sense. For example, specific details of a fielder’s action of throwing to first base do not seem to be grasped by the agent at the level of thought and understanding. It is natural to doubt that the agent grasps at the level of thought and understanding exactly how far back to draw her arm, or the specific point in the trajectory of her throwing motion where she releases the ball. Similarly, it seems *prima facie* unlikely that the capacities for balance that are part of a gymnast’s skill are rational capacities. A gymnast does not keep her balance by acting on beliefs that specify the way she should orient her body in order to keep from falling off the beam. That ability to keep her balance, developed through practice, seems to operate below the level of reason.

However, these abilities are constitutive of technical skills. The baseball player’s ability to determine exactly how far back to draw her arm, and the specific point at which to release the ball, are partly constitutive of her throwing skill. If the ball had been released just a little bit earlier or later in the trajectory of the throwing motion, the throw would have been inaccurate. And, similarly, the gymnast’s ability to keep her balance is part of her skill in gymnastics. It is one of the abilities that her skill-developing gymnastic training aims to cultivate.

instance, is part of what makes her a great tennis player. But these sorts of physical attributes are not part of technical skill, since they are not action-guiding capacities. Exercises of strength do not provide “instructions” about the manner in which an action is to be carried out. I reserve the phrase “technical skill” for capacities which guide the actions of skillful practitioners in that way.

Let me support this *prima facie* plausible suggestion that technical skills consist partly in non-rational capacities. First, the inability competent performers often have to explain certain aspects of how they do what they do is an indication that technical skill can consist partly in the possession of non-rational capacities. For example, people often cannot fully explain how they tie their shoes, without observing themselves do it a few times. In my own case, I could give a kind of schematic description from memory. But after observing myself tie my shoes a few times I noticed several integral details about how the action is carried out, which I had not been aware of before observation—how I propped up the first loop with my middle finger while forming the second loop, for example. One would expect that, if the details of how one ties one’s shoes were determined by exercises of rational capacities, the competent shoe-tie-er could explain those details, at least if she were appropriately prompted immediately after tying her shoes.

The view that technical skills can consist partly in non-rational capacities also finds support in the failure of articulable propositional knowledge to secure technical skill for an individual.¹² An enthusiastic baseball fan might have detailed, thorough propositional knowledge about the mechanics of pitching, and still possess no technical skill at pitching a baseball. One can acquire detailed propositional knowledge about pitching by watching and studying baseball. But in acquiring that sort of knowledge, one will not thereby become skilled. To become skilled, one must actually practice pitching. If technically skilled action were achieved through the exercise of rational capacities for thought and understanding, one would expect that this sort of detailed propositional knowledge would constitute technical

¹²Stanley and Williamson (2001) argue that this kind of know-how is a type of propositional knowledge, on the basis of a semantic analysis of the term ‘knows how.’ I am not convinced by their arguments, but I do not address them here. It is worth noting that there might be a way to reconcile their position with the claim that know-how consists partly in non-rational capacities. The views might be reconcilable if propositional knowledge were allowed to consist partly in sub-personal representations, or in the possession of non-rational capacities.

skill.¹³ The fact that this sort of detailed propositional knowledge does not confer technical skill suggests that certain aspects of technical skill consist in the possession of non-rational capacities.

Another source of support for the view that technical skill does not work just through reason is the relation of technical skills to rational criticism. If a pitcher lapses in her performance, and throws a wild pitch, she need not thereby be criticizable as irrational, unreasonable, foolish, or stupid.¹⁴ The same goes for an able musician who misses some notes, or an able dancer who botches a performance. In many cases, it seems mistaken to charge the agent with irrationality or stupidity on account of these sorts of errors. The action might be described as clumsy, or inept, but not irrational.

When errors in technical skill consist in failures of one's rational capacities to perform well, it is correct to charge the agent with irrationality, or some closely related kind of rational failure like being foolish or unreasonable. One's rational capacities would fail to perform adequately in determining appropriate means to the agent's goal of, for example, throwing a strike, or playing in tune. It would be a failure of reason to determine appropriate means to the agent's ends. This sort of failure in the operation of one's rational capacities would constitute a failure to satisfy norms of instrumental rationality, and open the agent to rational criticism. Since many failures in technical skill are not subject to rational criticism, many such failures are not failures in the exercise of rational capacities. Rather, they are failures in the exercise of non-rational aspects of technical skills.¹⁵

¹³Distinguishing between competence and performance is unhelpful to the rationalist here. The studious fan is unskilled. Her problem is not that there is some interference preventing her from deploying the skill she possesses.

¹⁴Such a lapse could consist in errant determination of details of one's action, or in a failure to determine details of one's action—that is, in errors of omission. Either kind of lapse would be a failure in the operation of the agent's technical skill.

¹⁵To be clear, my claim is only that some aspects of technical skill are non-rational capacities. I do not deny that some other aspects of technical skill, like a quarterback's ability to recognize the opponent's defensive formation, are rational capacities.

This conclusion finds further support in evidence from psychology. A picturesque example of the role of non-rational capacities appears in Gigerenzer's (2007) description of research showing that baseball players catch fly-balls by using strategies they are unaware of. Gigerenzer reports that, when a ball comes in high, a player fixates his gaze on the ball, starts running, and adjusts his speed so that the angle of gaze remains constant (p. 10). This strategy enables players to get themselves to the point where the ball will land, without requiring their cognitive subsystems to perform complex calculations about its trajectory. Most players are unable to become consciously aware that they are using this interception strategy.¹⁶ Many competent baseball players would find it informative to learn that they make use of this strategy. One would expect that users of this strategy would know about it, if the strategy were carried out through exercises of rational capacities.

Research about the role of different types of memory in skilled action also supports the view that human technical skill can be partly constituted by non-rational capacities. Psychologists distinguish between the declarative memory system and a group of non-declarative memory systems. Non-declarative memory systems seem, *prima facie*, to be non-rational capacities. Representations stored in declarative memory are available for conscious thinking, are available for use in guiding performance in a wide range of conditions, and have truth-assessable content. In contrast, information stored in non-declarative systems is not available for use in conscious thinking,¹⁷ is available for use only in a circumscribed range of conditions, and is not truth-assessable (Squire (2004)). Rational capacities are usually con-

¹⁶So are dogs, who use the same technique to catch frisbees (ibid., p. 12).

¹⁷A group of experiments by Pawel Lewicki et al. (1987, (1988) illustrate how strategies used in skillful activity can be thoroughly inaccessible to conscious reflection. Subjects were asked to locate the digit '6' on a screen filled with numbers as quickly as possible. With practice they improved at the task, by making use of a complex pattern predicting the location of the digit. But subjects were surprised to learn there was a pattern predicting the digit's location. Moreover, subjects were offered a \$20 reward for figuring out what the pattern was, and were given unlimited time repeating the task to solve the problem. But none had success. The researchers claim that, because of its complexity, normal human beings are incapable of consciously apprehending the pattern without making use of external heuristics, like diagrams. It seems very unlikely that the subjects made use of these complex patterns at the level of thought and understanding.

ceived of as capacities to operate with thoughts that are accessible to consciousness, that are available for general use by the agent in her thinking about any subject, and that have truth assessable content. Non-declarative memory systems seem to be “non-rational” capacities.

A lot of research about non-declarative memory concerns amnesia patients. Amnesia patients retain much of their know-how, and much skill-learning ability, despite their inability to explain what they are doing and how. In normal subjects too, the ability to successfully carry out certain tasks requiring know-how does not much covary with the ability to report facts and strategies they make use of in completing the task.¹⁸ Psychologists take these findings to suggest that skillful action is accomplished largely through the exercise of “non-declarative” memory systems, which seem to be non-rational capacities.

Accounts of motor control also posit a large role for specialized, inaccessible representations in guiding physical action. Psychologists claim that the motor control involved in, e.g., hitting a good shot in tennis is achieved largely through representations of sequences of movements that cannot be called up to consciousness, and are only available for use in carrying out a circumscribed range of tasks. For example, Willingham (1998) suggests that *goals* for skilled movements are determined by representations accessible to awareness and verbal description. But, he claims, representations involved in selecting targets for the movement are mostly not available to awareness. And the sequencing of movements involved in determining just how the body carries out the action is also determined through operations with representations that are mostly not available to awareness.¹⁹

The dominant view among researchers concerned with declarative and non-declarative memory, and with motor skills, is that much know-how consists in non-declarative memory

¹⁸See, e.g., Reber (1994), Willingham (1989), Squire (2009). Also compare Ennen (2003), who uses research on non-declarative memory to support a phenomenological account of agency as almost completely non-rational.

¹⁹Also see the work of Daniel Wolpert’s lab on motor control for accounts of computations involved in the functioning of motor skills, which are not plausibly attributable to the agent at the whole-individual level (see, e.g., Franklin and Wolpert (2011)).

systems and sub-personal motor capacities. It seems likely that these motor capacities and non-declarative memory systems are not capacities for thought and reasoning attributable to the agent at the whole-individual level. As a result, this research helps to support the claim that technical skill is partly constituted by non-rational capacities.

I want to emphasize that the case I have been making for the claim that technical skill consists partly in non-rational capacities does not depend on a legalistic conception of reason. The arguments I have made do not depend on the claim that reason is a capacity for emotionless, detached, rigid applications of general rules. Admittedly, the arguments do depend on certain assumptions about the nature of reason. Psychological states involved in exercises of reason are, in suitable circumstances, accessible to consciousness, and are articulable. Rational capacities are associated with rational criticism, like criticizing an action as irrational or foolish. Rational capacities are available for use in a general, non-circumscribed, range of tasks. Psychological states involved in exercises of reason are attributable to the agent at the whole-individual level. These traditional assumptions about reason are not legalistic. They are accepted by those who challenge overly legalistic conceptions of the nature of reason.

Technical skills and agential control

I have been arguing that technical skills consist partly in non-rational capacities. This conclusion threatens rationalist views about agential control only if these non-rational aspects of technical skill are capacities for agential control. I want to argue that they are.

The esteem which seems to be warranted by mastery of some technical skills provides an initial indication that they are capacities for agential control. The control that master practitioners exert over their crafts is an extraordinary human achievement. The mastery of craft possessed by a virtuoso violinist or tennis player or ballet dancer is an appropriate object of admiration and ambition for others, and an appropriate source of pride and fulfillment for

the practitioners themselves.²⁰ The appropriateness of attitudes like these toward technical mastery suggests that it is possessed by the agent “insofar as she is human.”

The technical skill of a tennis player or a ballet dancer is an ability to determine the way one plays tennis, or dances a ballet. So technical mastery is an ability to determine how one acts, which the agent seems to possess “insofar as she is human.” An ability to determine how one acts which one possesses “insofar as one is human” would seem to be a capacity for agential control. If that is right, then technical mastery is a capacity for agential control. The technical mastery that merits pride and admiration consists largely in non-rational capacities to determine essential details of skillful performance. Thus it seems that non-rational capacities are partly constitutive of capacities for agential control.

Another indication that technical skills are capacities for agential control is that they are capacities to determine how one acts that are genuinely attributable at the whole-individual level. It is the individual herself who is skilled at tennis or ballet dancing. The ability is not primarily attributable to sub-personal capacities, like the capacity for digestion, or the capacity to regulate one’s heartbeat. It is the agent *herself* who possesses the ability. An ability to determine how one acts that is possessed by the agent herself, at the whole-individual level appears to be a type of agential control.²¹

A further source of support for the view that technical skills are capacities for agential control comes from the relation between technical skill and being at fault. An individual

²⁰I assume that the object of these attitudes can be the mastery itself, and not only the effort that goes into acquiring it.

²¹It seems likely, though, that technical skills that are attributable at the whole individual level are constituted in part by sub-capacities that are sub-personal, and are not genuinely attributable at the whole-individual level. This fact may seem puzzling. One might expect that a capacity that is partly constitutive of a whole-individual capacity must itself be a whole-individual capacity. To see that this thought is mistaken, it may help to consider capacities attributable to groups. An ensemble’s ability to play well is attributable to the ensemble as a whole. The ensemble’s ability to play well might consist partly in the cellist’s ability to play well. The cellist’s ability to play well is “modular,” in that it is not attributable to the ensemble as a whole. Nonetheless, the cellist’s ability is still partly constitutive of the ability of the ensemble as a whole. Though of course there are important differences between group abilities and abilities of individuals, the analogy helps show how abilities attributable at the whole-individual level might consist partly in sub-personal capacities.

can be at fault in virtue of failures of non-rational aspects of technical skills. Consider an able musician who misses some notes during a concert, or an able athlete who botches a routine play, and costs her team the game. The athlete and the musician are at fault for their mistakes. In many cases their teammates, or fellow ensemble members, may blame them for their mistakes.

Of course, the blame is not *moral* blame. The agent is not at fault for a moral error. There is no moral error in performing poorly in sports or music. One tends to think of notions of blame and being at fault as associated with moral error. But these notions can apply to non-moral mistakes as well. They seem to apply to failures of technical skills. Mistakes in the exercise of technical skills can warrant non-moral blame, as well as other “reactive” attitudes and responses, like regret, and apology.²² The attribution and acknowledgement of fault is commonplace in sports, in musical performance, and in other contexts where there is something at stake for other people in the success of the exercise of technical skills. Bill Buckner, for instance, is famously at fault for his flubbing of a routine play which could have won the 1986 world series for the Boston Red Sox.

The fault in cases of technical mistake can be direct. It need not derive from an earlier failure to train adequately, or to get enough sleep the night before the match or concert, for example. The fault also does not have to derive from a failure to concentrate adequately, or to care enough about performing well. We can imagine that the agent has practiced adequately, is well rested, is concentrating, and so on, but still fails to perform up to her abilities, and flubs her execution of the action. She is still at fault for her mistake. She is at fault in virtue of a failure of the exercise of the technical skill itself. The evidence for this claim comes from reflecting on examples of poor performance. It is commonplace to attribute and acknowledge fault for poor performance, even when a mistake does not appear

²²See Strawson’s “Freedom and resentment” (2003) for discussion of the relation between freedom, agency, and attitudes like these.

to derive from an error beyond the exercise of the technical skill itself, like a lack of diligence in preparation, or a failure to concentrate.

The circumstances under which agents are excused from fault for mistakes in technical performance also suggests that non-rational aspects of technical skill are capacities for agential control. If an individual errs in her performance of some technical action, limitations in her skill and obstructions to the exercise of her skill can excuse her from being at fault. For example, if a baseball player's throw misses its target, she will not be at fault for missing if the target was too far for her to be expected to hit it, given limitations in non-rational aspects of her capacities to aim her throw. In contrast, a baseball player whose skill-level permits her to hit the target without much difficulty would be at fault for missing. Or, to take another example, if a baseball player's throw misses its target because her pain medicine interferes with her motor system, she will not be at fault for the miss, whereas she would have been at fault for the mistake if not for the drug's interference.

The concept of excuse is tied to the concept of agential control. If limitations or interference with capacities to determine how one acts excuse fault for mistakes, then those capacities are capacities for agential control. When limitations and interference with capacities excuses fault, it is because the limitations or interference mean that the agent's mistake is not due to a failure that was in her control as an agent. The failure is out of the agent's control in that the capacities for agential control through which she could have acted successfully are limited, or interfered with. Thus, if limitations or interference with one's capacities excuses fault, then those capacities are capacities for agential control over what one does. Since limitations and interference with non-rational capacities can excuse fault, non-rational capacities can be capacities for agential control. For example, if the capacities of the baseball player in the preceding paragraph were not limited in the excusing way, or were not interfered with, the individual would be at fault for her mistake. She would have been at fault because those capacities would put it within her control as an agent to avoid

the mistake without too much difficulty. Limitations in non-rational capacities can excuse one from fault for errors in performance because they are limitations in the agent's control over her action.²³

Criticizing rationalism about action

I have been arguing that the capacities for agential control that guide our actions include non-rational capacities. I now want to argue that these non-rational capacities for agential control can help make the things we do count as actions.

Details of the way a skilled action is carried out are actions, or aspects of actions that manifest agency. Admittedly, at least in some cases, it stretches ordinary usage to describe aspects of actions as actions themselves. A pitcher's releasing the ball at the specific point she does may be more naturally thought of as an aspect of her pitching action, than as an action in its own right. But what is important for present purposes is that these sorts of details of skilled action manifest human agency. Earlier I argued that non-rational aspects of technical skill that fix such details are capacities for agential control. Variants of most of those arguments apply to the claim that the details of behavior themselves are manifestations of agency.

First, the skillfulness of our actions warrants pride and admiration. That skillfulness consists largely in details of actions that are not specifically conceptually represented in

²³As a clarification about these last two arguments, concerning fault and excuse, it will help to distinguish two different ways in which one's action might constitute a mistake. First, an agent might make a mistake in carrying out her physical action unsuccessfully. Second, an agent might make a mistake in *attempting* an action. For example, one might make a mistake in attempting a long throw to home plate from the outfield, if one's lack of ability makes success very unlikely. The attempt might constitute a mistake even if the throw happens, by luck, to be successful. Being at fault for this sort of mistake is a matter of whether one has and makes proper use of information about the limitations of one's abilities. The exercise of abilities involved in *executing* the difficult throw are irrelevant to whether the agent is at fault for *attempting* the throw. The mistakes at issue in the last two arguments concern mistakes in the execution of one's action, and not mistakes in attempting an action.

the agent's intentions. Because those details of action are part of the skillful action that warrants attitudes like pride and admiration, they seem to be things we do as agents. Second, it seems to be the agent *herself* who acts skillfully. Unlike digestion, an individual's acting skillfully seems to be attributable to the agent at the whole individual level. And, unlike fidgets and Freudian slips, there is no evident lack of control in cases of skillful action. Since the action's skillfulness largely consists in details of her action, those details seem to be aspects of activity that are due to the agent *herself*. Activity that is due to the agent herself would seem to be a manifestation of agency. Third, an individual can be at fault for failure to get these details right. A pitcher can be at fault for throwing inaccurately, even though the inaccuracy of her throw is due to a mistake in details of finger placement determined below the level of thought and understanding. If such details of how a technically skilled action is executed were outside the scope of what the individual does "as an agent," the agent would not be at fault for getting those details wrong.²⁴ The agent manifests agency not only in attempting a pitch, but also in throwing the pitch in a skillful manner, in getting these details of skillful execution right.

Details of skillful action count as manifestations of agency partly in virtue of the exercise of the non-rational capacities that determine those details. To make my case for this claim, I want to begin by considering a plausible, but ultimately mistaken, rationalist view about what makes details of skillful action count as manifestations of agency.

A rationalist about the nature of human action might point out that details of our skilled behavior do have some connection to exercises of rational capacities, at least in most cases. For example, a pitcher's releasing the ball at a certain point in her throwing motion is part of the execution of a conceptualized intention to throw a pitch. That intention, which is an exercise of the agent's rational capacities, is part of the psychological explanation for the

²⁴Or, more precisely, she would not be *directly* at fault for getting the details wrong, rather than being derivatively at fault—in the way she might be derivatively at fault for poor digestion that results from failing to take her medicine regularly.

agent's releasing the ball at that point. She would not have thrown the ball at all if not for the intention. So, *a fortiori*, she would not have thrown the ball in that particular way, releasing it at that particular point in the throwing motion. A rationalist might claim that what makes the agent's releasing of the ball at that certain point a manifestation of agency is the fact that it is a result of the agent's intention to throw the pitch.²⁵

This view might seem to find support in the fact that an agent can exercise her skills outside the context of intentional action. A skilled tabla player, for example, might absent-mindedly and non-intentionally, or unconsciously intentionally, drum her fingers on the table, in the same way she would if she were playing her instrument. A rationalist might argue that the details of the non-intentional drumming are not manifestations of agency of the right kind. They are aspects of a complex kind of fidgeting, and are not attributable to the agent *herself*, as an agent. When details of skilled behavior *are* attributable to the agent "as an agent," it is because of their relation to exercises of her rational capacities.

I want to resist the claim that details of the tabla player's non-intentional table drumming are not attributable to her as an agent. There is certainly some kind agency involved in the drumming, since the agent controls the details through exercises of her skill. And the activity is attributable to the agent herself, at the whole-individual level. But it does seem plausible that, when the drumming occurs as part of the execution of a conceptualized intention,

²⁵This proposal is similar to one made by Randolph Clark (2010) about the nature of intentional action. Clarke suggests that an event might count as an intentional action because of its relation to an intention, even if it is not represented in the content of the intention. For example, he suggests that a particular stroke made in the course of shaving might count as an intentional action because it is caused in a certain way (i.e., non-deviantly) by the agent's intention to shave, even if the particular stroke is not represented in any intention. (I think it is clear that particular strokes in shaving are intentional, but he could have used a better example, like releasing a ball at a specific point in one's throwing motion.) Clark's proposal is about whether causal connection to intention is what makes an event count as an intentional action. He presents this view in response to a criticism of Davidsonian causal theories of action, due to David-Hillel Ruben (2003). Though related, this issue is not exactly the same as the issue under discussion here, which concerns the role of rational capacities in making an event count as a manifestation of agency. If intentions can have non-conceptual content, then the Davidsonian view that events count as actions because of their causal connection to intentions is compatible with the claim here—that events count as actions partly because of their connection to non-rational capacities for agential control.

the details of the drumming are manifestations of agency of a different, “fuller” kind. The rationalist might claim that if an action, or aspect of an action, is a manifestation of agency of this “fuller” kind, it is because of the rational control the agent exerts over it.

I want to argue that when details of skillful action are manifestations of agency—of this “fuller” kind, or of any other kind of agency—it is partly in virtue of the agent’s non-rational control over the action. The core point is that exercises of the agent’s rational capacities do not specify the details of skilled action. As a result, an agent’s rational control over the details of skilled action is not sufficient to explain why the action counts a manifestation of “full” agency, or any other kind of agency.

Details of skillful action count as manifestations of agency because of some aspects of agential control that determine that the action shall be carried out in *that* way, rather than some other way. Generally, if an individual’s capacities for agential control leave it open whether or not she shall ϕ , instead of determining that she shall ϕ , then her ϕ ing is not a manifestation of agency. Her ϕ ing counts as a manifestation of agency partly in virtue of the exercises of capacities for agential control that determine that she shall ϕ . For example, if an individual’s stepping on the hose is a manifestation of agency, it is at least partly because exercises of her capacities for agential control—like an intention to step on the hose—determine that she shall step on the hose. If she does not intend to step on the hose, and no other capacities for agential control specify that she shall step on the hose, then her stepping on the hose is an accident, and not a manifestation of agency.

Suppose the agent has some more general intention, like an intention to stop the water flowing. If she executes that intention by stepping on the hose, her stepping on the house does not count as a manifestation of agency just because of its relation to her intention to stop the water flowing. That more general intention leaves it open whether or not she will step on the hose. Her stepping on the hose counts as a manifestation of agency in virtue of exercises of capacities for agential control that determine that she will do that action in

particular, rather than, say, turning the water off at the faucet.

As argued above, exercises of rational capacities leave open many details of how technically skilled action is to be carried out. An intention to throw a baseball leaves open the details of how the ball is to be thrown, in the same way that an intention to stop the water flowing leaves open whether the agent will step on the hose, or stop the water flowing some other way. An intention to stop the water flowing does not suffice to explain why stepping on the hose counts as a manifestation of agency. Similarly, an intention to throw a ball does not suffice to explain why the details of the action count as manifestations of agency. The details count as manifestations of agency partly in virtue of exercises of capacities for agential control that specifically determine those details. The capacities for agential control that specifically determine those details are non-rational capacities. Thus, the details of skillful action count as manifestations of agency partly in virtue of the exercise of those non-rational capacities for agential control.

It may be true that when details of skillful action are part of the execution of a conceptualized intention, they are manifestations of a “fuller” kind of agency than activities like non-intentionally drumming one’s fingers on the table. It is plausible that exercises of rational capacities are essential to our actions’ counting as manifestations of this fuller kind of agency. But that view is compatible with the view that non-rational capacities can also help make an action count as manifestations of that fuller kind of agency. Whether details of skillful action are attributable to the agent herself, “as an agent,” is a matter of agent’s exerting control over those details. Non-rational capacities are essential to our control over those details, since it is those capacities, and not rational capacities, that specifically determine those details. So exercises of non-rational capacities help make those details count as manifestations of agency.

The arguments in this chapter support the conclusion that human agency does not consist exhaustively in rational capacities, and that non-rational agential control can help make the

things we do count as actions. These claims are compatible with the proposal that all manifestations of the fullest kind of human agency count as manifestations of agency *partly* in virtue of their relation to exercises of rational capacities, like intentions with conceptual content.²⁶ The idea that rational capacities are essential in this way to the fullest kind of human agency would be a way of respecting the plausible and widespread view that there is a fundamental connection between human agency and reason.

The conclusions argued for here, concerning the constitutive role of non-rational capacities in human agency, are only meant to *temper* overly rationalist views about agency. The point of arguing against these overly rationalist views is not to deny that reason is fundamental to human agency and human action. Instead, the point is to suggest that human agency, our ability to determine for ourselves how we act, comprises both rational and non-rational capacities; and the things we do count as actions in virtue of both rational and non-rational aspects of our agency. The plausible claim that reason is fundamental to human agency should not lead us to neglect the importance of non-rational capacities in human agency and action.

²⁶The claims are also compatible with the view that all manifestations of human agency, of any kind, count as manifestations of agency partly in virtues of the relation to exercises of rational capacities. But I do not endorse that view. Non-intentional human activities, like idly drumming one's fingers, count as manifestations of agency, even though the agent exerts no rational control over such activities.

Chapter 3

Interpersonal skill and ethical success

This chapter investigates the ethical significance of certain aspects of interpersonal skill. The discussion is intended to apply friendly pressure to rationalist views about ethical goodness, in order to help clarify what notion of ethical goodness those rationalist views aim to capture. According to the rationalist views I have in mind, an action is ethically good just in virtue of its being based on a rational recognition of reasons for the action. Kant is widely interpreted as having a view along these lines. For Kant, if an action is morally good, it counts as morally good just in virtue of whether it is performed because it conforms to the moral law. Acting because of the moral law means being motivated to act by an apriori rational recognition of the authority of the moral law. An action's connection to this rational recognition of the moral law is what makes the action ethically good. Aristotle's view of ethics is often taken to be rationalist as well. According to Aristotle, the human good, including ethically good action, consists in activity carried out through good exercises of the rational part of the psyche.¹ These rationalist views are picked up by contemporary representatives of

¹The practical part of the rational part of the psyche, in particular. Compare books 6 and 9 of the Nicomachean ethics for discussion of the difference between practical and non-practical kinds of rational success. Aristotle's remarks about the role of pre-rational dispositions toward pleasure and pain in ethical character pose at least a *prima facie* difficulty for a fully rationalist interpretation of his ethics. I set that interpretative issue aside here.

Kantian and Aristotelian traditions in ethics, like Korsgaard (e.g., (1996) and Foot (2003). Rationalism about ethics is not confined to these two traditions. It is also present in work that seeks to understand what is ethically required of us in terms of the reasons our actions ought to be based on, like the work of Parfit (2011), Scanlon (1998) and Raz (1999).²

These different views in ethics have in common that they characterize the ethical goodness of our actions in terms of an underlying rational appreciation of things—of what is good and worthwhile in life, of the significance of certain other people in our lives, and so on. This, then, is the rationalist view this chapter aims to challenge: that our actions count as ethically good just in virtue of their connection to the rational perspective underlying our actions. The defining commitment of this rationalist view is that what makes an action ethically good is just the rational source of the action—that is, the operation of the agent’s rational capacities that explain why she acted as she did. The rationalist view does allow that non-rational capacities can play an important role in *enabling* us to perform ethically good actions. The non-rational capacities involved in digestion and respiration help enable us to act in ethically good ways by helping to keep us living. But our actions do not count as ethically good in virtue of their operation. Rationalism claims that only rational capacities play a *constitutive* role in providing for the ethical goodness of our actions.³

There is much to recommend this rationalist conception of ethical goodness. The ability to guide one’s actions according to a thoughtful appreciation of what is right, and what matters in life, is clearly indispensable to living an ethically good life. To live an ethically good

²Some views of this last kind are inspired by Kant, but many of them differ in significant ways from Kant’s own view.

³There are other rationalist claims about ethics. One concerns the epistemology of ethical norms: all ethical norms, or a core subgroup of ethical norms, are knowable a priori. According to another rationalist claim, the applicability and authority of some or all ethical norms does not depend on contingent “empirical” vicissitudes of time place and person, like what one happens to desire, or what a culture happens to value. Neither of these is the rationalist view at issue in this document, though they are related. The rationalist view at issue here is the view that what makes our actions count as ethically good is a kind of success in the operation of our rational capacities.

life, one must have good values, and one must have ethically good reason-based attitudes toward other people, like respect and compassion. For our actions to be ethically good, they must be rooted in these sorts of values and attitudes.

However, in this document I develop a challenge to the rationalist conception of ethical goodness. There is room to doubt that ethical goodness consists *exclusively* in exercises of reason. Our non-rational capacities to determine our behavior also seem to play a constitutive role in providing for the ethical goodness of our actions. My discussion focuses on a range of behaviors involved in interpersonal interaction, which stem from non-rational aspects of interpersonal skill. I argue that certain aspects of the ethical goodness of our actions can depend constitutively on exercises of non-rational aspects of interpersonal skill, and not just on exercises of rational capacities.

This point should not be taken to overturn rationalist views about ethics. Instead, the point should be treated as an occasion to get clearer about the kind of ethical success that traditional rationalist views in ethics aim to capture. In my view, there are core aspects of ethical goodness which do consist exclusively in exercises of rational capacities. These core aspects of ethical goodness are keyed to whether one has good values, good reason-based attitudes toward others, and the like. Non-rational capacities contribute to what I call “technical” aspects of ethical goodness. These aspects of ethical goodness are keyed to whether one does well in determining how in particular to act given the vicissitudes of the context of action, so as to bring to bear one’s good values and attitudes in an appropriate way. Traditional rationalist views about ethical goodness should be understood as aiming to capture the core aspects of ethical goodness, and not to have these technical aspects of ethical goodness within their purview.

Rationalism vs. legalism about ethical norms To make it clearer what the rationalist view is that I mean to put pressure on, I want to distinguish my aim here from that of

Dreyfus and Dreyfus (1990). They argue against a certain intellectualist conception of ethical goodness. The view they oppose is that ethically good action is action based on sound judgment about how best to act in the circumstances, where (i) the judgment is based on algorithmic, quasi-deductive inference from general principles, (ii) the agent is occurrently conscious of her reasoning, and (iii) her judgment, and the reasoning underlying the judgment, are marked by a detached attitude toward the situation. They argue that “experts” in ethical action usually do not rely on judgments about what to do in the situation, do not make use of inferences from general principles to determine what to do, do not have occurrent conscious awareness of their reasoning, and usually experience themselves as “immersed” in the situation rather than detached from it.

They present their view as aligned with a tradition that includes Hegel, Charles Taylor, and Carol Gilligan, and which they claim opposes an intellectualist tradition that includes Kant, John Rawls, and Lawrence Kohlberg. The view they mean to oppose seems better described as legalistic than intellectualist or rationalist. It is better described as legalistic because of its emphasis on settling questions about what to do by adverting to abstract general principles in a way that is detached, consciously self-aware, analytical, and algorithmic.⁴

The tradition Dreyfus and Dreyfus align themselves with is intellectualist and rationalist, even though it is not legalistic. It is clear, for instance, that Gilligan’s arguments in “In a different voice” are aimed at correcting a mistaken tradition about developed capacities for thought, understanding, and judgment. They are not aimed at excising the notions of thought, understanding, and judgment from a proper account of moral development. Gilligan describes her arguments as “challenging the current assessment of women’s moral judgment” and aimed at encompassing “the thinking of both sexes” (p 2). She thinks of the contrast between the moral expertise she identifies as common in women, and the Kohlbergian account

⁴This legalistic view might fit Kohlberg’s conception of moral development. But I doubt that it accurately captures the views of Kant and Rawls.

of moral development, as a contrast between “two modes of judging” (p. 8).

A large part of the significance of Gilligan’s work consists in what it has to say about what it is to be a thoughtful, intellectually sophisticated being. Thinking well about what to do is not a matter of the algorithmic application of general principles. Nor is thinking about what to do detached in a way that is rightly contrasted with acting from emotion. Emotions are plausibly integral to moral thinking. Anger, sadness, sympathy, frustration and so on are all plausibly aspects of our intellectually sophisticated perspective on things. Rather than undermining the traditional view of a moral being as a ‘thinking thing’ or a ‘rational animal,’ Gilligan’s work sheds light on what it is to be a thinking, reasoning being.⁵

The contribution to ethical goodness made by capacities for thought and understanding remains an intellectual contribution, even if it is recognized that thought and understanding do not require detachment, algorithmic inference from general principles, or the excision of emotion. The exercise of intellectual capacities makes certain kinds of evaluations appropriate, like evaluating an individual as wise, thoughtful, or insightful—or, on the other side, obtuse, confused, or thoughtless. In many normal cases where there is no detached algorithmic deliberation, evaluative notions like these are still in place.

In this chapter, I am interested in the ethical significance of capacities that are not rational capacities even in this wider sense. I claim that the ethical success of treating others well is due partly to the exercise of “low-level” action-guiding abilities, which do not warrant this sort of intellectual evaluation. The action-guiding abilities I have in mind are non-rational aspects of interpersonal skill.

⁵Also see Williams (1973), Nussbaum (1992) and Damasio (1999) on the role of emotion in thought.

Interpersonal skill and technical aspects of ethical success

As a first step in supporting this claim, I want to discuss the role of interpersonal skill in ethically good action in a more general way, without focusing specifically on non-rational aspects of interpersonal skill. My aim in this section is to suggest that the ethical success of treating others well consists partly in the successful exercise of a certain kind of “technical” know-how concerning specifics of how others ought to be treated given the vicissitudes of the context of action, so as to bring to bear one’s good values and attitudes in an appropriate way.

I would like begin by characterizing the sort of success I have in mind in talking about ethical success and ethical goodness.⁶ But it is difficult to specify which varieties of practical success are distinctively ethical. It is a large task to map the terrain covered by notions of moral goodness, ethical goodness, living well, human flourishing, and so on. In my view, to map that terrain requires substantive philosophical work about what different kinds of practical success consist in. Indeed, part of the aim of this chapter is to make progress in getting clearer about how ethical success might be divided off from other kinds of practical success.

But I do want to offer a very rough, initial working characterization of ethical goodness. In talking about ethical goodness, I mean to be tracking a notion that is familiar and commonplace. The notion I have in mind is closely connected to the phrase ‘being a good person’—as in “I find myself attracted to people who are stylish and successful, but what I really want is to be with someone who is a good person.” Popular wisdom has it that the characteristics associated with being a good person include being fair, honest, compas-

⁶For present purposes, I see no advantage to distinguishing between ethical success and ethical goodness. I use the terms according to which allows for less awkward and less cumbersome expression.

sionate, acting with dignity, being adequately moved by the needs of others, nurturing one's relationships with friends and family, and so on. According to the notion of ethical goodness I have in mind, an action is ethically good roughly when the action manifest one's quality of being a good person. This notion of ethical goodness does not encompass every meaningful, valuable aspect of human life. Appreciating music does not constitute ethical success, but it can be an enriching, meaningful part of life. In this respect, the notion of ethical success at issue here is narrower than Aristotle's notion of *eudaimonia*, which encompasses non-ethical aspects of living well, like contemplation.

Treating other people well in interpersonal interactions constitutes a certain kind of ethical success. Treating others well can consist in being appropriately kind, loving, compassionate, friendly, and the like. Treating others well can also consist in expressing disapproval or resentment in a measured and respectful way, in being appropriately contrite in making an apology, or expressing sincere gratitude for a good turn.

Treating people well in these ways involves some kind of recognition that, in the circumstances, one ought to act lovingly, contritely, or compassionately, for example. In addition, to succeed in treating others well in these way, the agent must also determine *how* to act in an appropriately loving, contrite, or compassionate way toward the other person. She must determine what sort of behavior constitutes, for example, compassionate treatment in the circumstances. If the agent is culpably mistaken about whether her action constitutes compassionate behavior, there is something ethically defective in her way of treating the other person. For example, a person intending to act in a compassionate way toward a friend grieving the loss of his mother might mistakenly suppose that her friend will find comfort in stories about the deceased mother, when in fact the stories are painful to hear. An agent who culpably makes a mistake like this culpably fails to treat her friend well.

In some cases, mistakes like these can be due to a problem in the agent's values, or in her reason-based attitudes toward the other person. The agent in the example above might

make her mistake about how to act in a successfully compassionate way because she is not sufficiently attentive to cues from her friend, about whether he likes hearing the story. She might be inattentive because she is somewhat self-absorbed, and so distracted by her own comparatively minor problems. Or she might be inattentive because she does not value her friendship, or appreciate the importance of compassion, with sufficient seriousness—even if her relatively shallow positive attitude toward her friendship and toward compassion are enough to get her to *intend* to act in a compassionate way toward her friend. Problems with the agent’s values and reason-based attitudes might lead the agent to unintentionally treat her friend in a hurtful, non-compassionate way.

In other cases, however, the mistake need not trace back to corrupt values and attitudes of this kind. Such a mistake can be purely epistemic. The inadequately compassionate behavior in the example above might be due to an epistemic mistake of failing to realize that the stories would be painful rather than comforting, and not to any antecedent ethical problem with the agent’s psychology. Given that one wants to be treated another person in a compassionate way, there is an epistemic question as to what particular behaviors are appropriately compassionate. We have a kind of agency over at least some of the ways we make use of epistemic reasons in reasoning about matters of fact. Accordingly, the agent can be epistemically culpable for a mistake in figuring what kind of behavior would be appropriately compassionate, without the mistake being due to some problem with her values, or how much she cares about her friend. Even in a case like this, where the primary locus of the mistake is epistemic, the agent culpably fails to treat her friend in an appropriately compassionate way.

Cases like this one bring up a question about what kinds of practical success count as *ethical* success. If an individual culpably fails to treat others well, but the locus of the failure is primarily epistemic, one might ask whether or not her failure counts as an *ethical* failure. One might want to reserve the term ‘ethical failure’ for cases like the one in the first example,

where the agent fails to treat her friend well because of something corrupt about her values, or her attitude toward her friend. On the other hand, in the second case, the agent does culpably fail to treat her friend well, even though the mistake is primarily epistemic. It is plausible that a culpable failure to treat someone as he ought to be treated is an ethical failure.

This issue complicates certain ways of thinking about the relation between ethical goodness and having a good will. According to the Kantian view sketched in the introduction, whether an action is ethically good depends on whether it is motivated by a good will, where that amounts to something like being motivated by one's respect for people as autonomous rational agents. According to this view, an action counts as ethically good treatment of another person because the action is motivated by the agent's respect for the other person as an autonomous rational agent. It is not immediately clear how this view squares with cases of primarily epistemic failure to treat others well. A culpable failure to treat a person in an appropriately compassionate way *can* be due to a failure to respect the person in the right way. But it need not be. A culpable failure to treat a person in an appropriately compassionate way can be due in the first instance to an epistemic mistake of failing to recognize exactly what kind of behavior is called for in the circumstances. It need not be due to an antecedent failure to respect the other person.

In my view, cases of the kind described here should not be taken as grounds to reject views like Kant's. Instead, they should be taken as occasions to clarify what kind of kind of ethical success such views aim to capture. We can distinguish in a rough and ready way between "valuational" and "technical" aspects of ethical success. Views like Kant's should be interpreted as aiming to account for valuational aspects of ethical success, and not for every kind of ethical success. "Valuational" aspects of ethical success include different kinds of success involved in having ethically good values, and ethically good attitudes toward others, and toward oneself. Examples of successes of this kind include caring about one's children,

valuing fairness, valuing compassionate treatment of others, and respecting all people as having equal worth.⁷ “Technical” aspects of ethical success include different kinds of success involved in determining how in particular to act, so as to bring to bear good values and attitudes in an appropriate way. Epistemic success in determining what particular behavior would constitute compassionate treatment of a grieving friend is an example of a technical kind of ethical success.⁸

Different kinds of ethical mistakes correspond to these different aspects of ethical success. A culpable failure to value fairness, to respect others, or to care about one’s children, is a failure with respect to valuational aspects of ethical success. If one makes no mistake of this kind, but still fails to bring to bear one’s good values and attitudes in an appropriate way, that constitutes a failure with respect to technical aspects of ethical success.⁹

⁷These values and attitudes themselves constitute valuational ethical successes. Actions constitute valuation ethical successes when they are based on ethically good values and attitudes of these kinds.

⁸The contrast drawn here between valuational and technical aspects of ethical success is different from the contrast sometimes drawn between Ancient conceptions of ethical success and modern or contemporary conceptions of moral success. The Ancient conception of ethical success is sometimes claimed to be grounded in the agent’s own happiness or flourishing, whereas modern and moral conceptions of moral success are not grounded in the agent’s own happiness (see Annas (1992)). If the Ancients conceived of ethical success as grounded in happiness, the conception of ethical success at issue in this chapter is different. In my view, happiness is grounded, in part, in ethical success, and not the other way around. The contrast drawn here, between valuational and technical aspects of ethical success does not concern that issue about the relation between ethical success and happiness.

⁹Many serious ethical failings are regarded by their agents as consistent with what most of us would regard as good values. Torturers profess to act justly, regarding their victims’ rights as forfeit by their misdeeds, or considering the torture just in light of its potential benefits to others. Violent treatment of children is often regarded by its perpetrators as stern-but fair, or as character-building. Colonizers justify their paternalistic treatment of natives by claiming the treatment helpful to the moral development of backwards people. Aristotle regards the institution of slavery as consistent with the equal treatment of rational agents on the grounds many slaves are not actually rational agents.

If one takes the avowed attitudes and values of wrongdoers at face value, then their mistakes would seem to be technical, and not valuational. Aristotle’s mistake, for example, would be a primarily empirical epistemic mistake, about the cognitive abilities actually possessed by slaves, and not a valuational failure to regard all rational agents as having equal worth. But of course, in many cases we should not take wrongdoers avowals at face value. Failures in recognizing what is required by the equal worth of persons, or by justice, are often due, at least in part, to failures to adequately value justice or the equality of persons.

As I note below, there are constitutive connections between genuinely and seriously holding certain values, and certain views about what kinds of behaviors accord with those values. For example, holding the view that torture is compatible with justice might itself constitute a failure to fully understand and accept the value of justice. In other cases, mistaken views about what behaviors accord with one’s values do not constitute

As long as one is alert to the difference between technical and valuational aspects of treating others well, I can see nothing wrong with using the term ‘ethical success’ to cover the technical aspects, as well as the valuational aspects. There is some way in which culpably failing to treat others well falls short of an ethical ideal, even if the mistake is primarily epistemic. Such a mistake is ethically significant in a way that culpably making an epistemic mistake in an ordinary game of trivial pursuit is not.

The view that there are technical aspects of ethical success is consistent with some central historical accounts of ethical success. The view shows up in Aristotle’s claim in book II of the *Nicomachean Ethics* that ethical virtue involves getting things right with respect to a variety of parameters. According to Aristotle, ethical virtue is an ability to do things (and feel emotions) at the right times, toward the right objects, toward the right people, with the right motive, to the right degree, and in the right manner.¹⁰ Valuational aspects of ethical success would have to come in under the “right motive” parameter for Aristotle. But having the right motive for an action is only a part of virtue. Many of the other parameters seem to concern technical matters, like determining how great a gift would constitute acting in an appropriately generous way, instead of being prodigal or tight-fisted (cf. EN 1107b9-22, and chapters 1 and 2 of book IV). It is one thing to value generosity, the needs of others, and one’s own well-being. It is another thing to discern what degree of charitable giving accords with these values. The capacity for the latter kind of technical discernment is an integral part of ethical virtue, according to Aristotle.

The link Peter Strawson (2003) draws between ethical success and reactive attitudes helps support the claim that there are technical aspects of ethical success—although that claim seems to conflict with Strawson’s own view about what warrant reactive attitudes.

a failure in one’s values, but are due to a failure in one’s values. For example, the view that children suffer from having to parents of the same sex probably does not constitute a failure to fully value gay people as having equal worth as others. But in most cases the view is due to such a failure in values, at least in part.

¹⁰See, e.g., EN 1106b15, 1107b18, 1109a27.

According to Strawson, reactive attitudes like resentment and gratitude are warranted by the good will or ill will (or indifference) of others toward us.¹¹ But cases of the kind described above suggest that reactive attitudes can be warranted by actions that do not manifest ill will or indifference. A culpably botched expression of good will, like an inadvertently hurtful attempt to be compassionate to a grieving friend, can warrant reactive attitudes of a certain kind, even if the mistake is primarily epistemic. The mistake is something to regret, and to apologize for. And the mistreated person may appropriately feel something like resentment, or disappointment in his friend, on account of her behavior.

The point that reactive attitudes are appropriate in cases when mistreatment is due to primarily epistemic mistakes dovetails with the point that such mistreatment constitutes a *culpable* failure to treat the other person as he ought to be treated. In such a case, it is well within the agent's power to recognize how the other person ought to be treated, and to act accordingly. Even if the agent's will is good, still, she should have known better than to act as she did toward her friend. She is accountable for her inadvertent mistreatment of her friend, even though her mistake is primarily epistemic. The fact that her primarily epistemic mistake warrants reactive attitudes suggests that it is a genuinely ethical mistake. If so, there are technical aspects of ethical success, in addition to valuational aspects like bearing others good or ill will.

Now, one may deny that reactive attitudes are reasonable in cases where mistreatment is due to a primarily epistemic mistakes. But that denial would be an error theory, and conflict with widely accepted norms governing ordinary practices of holding and expressing attitudes of resentment and disappointment, making and accepting apologies, and so on.

I am claiming that there are both technical and valuational aspects of ethical success. But I want to emphasize that valuational aspects of ethical success are more central. Technical aspects of ethical success depend on valuational aspects of ethical success. If one has

¹¹Ibid. p. 8.

bad values, there is nothing ethically good about the “technical” achievement of competently bringing those bad values to bear in one’s particular way of acting. If one values the subjugation of others, like Thrasymachus in *The Republic*, there is nothing ethically good about competently subjugating others. In contrast, if one has good values, but is incompetent in bringing them to bear, having the good values still constitutes an important kind of ethical success. The fact that valuational aspects of ethical success are more central also comes through in the difference between the kinds of reactive attitudes warranted by valuational and technical aspects. The kind of resentment, disappointment, anger, regret, and contrition appropriately occasioned by primarily technical failure to treat others well is mild in comparison to the sort occasioned by valuational failure.

In my view, there are constitutive connections between genuinely and seriously holding certain values, and certain views about what kinds of behaviors accord with those values. In the *Republic*, Thrasymachus claims that it is just for the strong to subjugate the weak. I doubt that it is possible for someone to genuinely and seriously value justice, while holding a view like that. Holding Thrasymachus’s view constitutes a fundamental misunderstanding of what justice is, which would undermine one’s claim to genuinely value justice. Similar points apply to other valuational aspects of ethical success, like valuing kindness, and caring about one’s friends.

Nonetheless, genuinely and seriously valuing justice and kindness, caring about one’s friends, and the like, are consistent with certain kinds of culpable mistakes about what behaviors accord with those values and attitudes. A mistake about whether a grieving friend will want to hear a story about his deceased parent can constitute a failure with respect to technical aspects of ethical success, without undermining the agent’s claim to genuinely and seriously care about her friend, to value compassion and kindness, and so on.

Non-rational aspects of interpersonal interaction

Acknowledging that there are technical aspects of ethical success provides an opening to consider the role of non-rational capacities in providing for ethical success. Much of our ability to determine how we act consists in non-rational action-guiding capacities—including integral aspects of the interpersonal skills we make use of to treat others well. Exercises of non-rational aspects of our abilities to determine how we act are partly constitutive of our agential control over how we act. Non-rational aspects of interpersonal skill are capacities for agential control, which are integral to our ability to treat others in ethically good ways in interpersonal interaction. I want to suggest that the successful exercise of these non-rational aspects of interpersonal skill is partly constitutive of technical aspects of ethical success.

Let me sketch the case for the claim that our agential control consists partly in exercises of non-rational capacities. The appeal of the claim might be clearest in the cases of specialized technical skills outside of interpersonal interaction, like the ability to juggle, to play a violin in tune, or to throw a curve ball. A violinist's capacity to determine fine details of how to move her fingers, and a baseball player's capacity to determine details of how to release the ball, are integral to what makes them skillful practitioners. These capacities are partly constitutive of the agent's ability to exert agential control over her actions. The contribution these capacities make to agential control is reflected in the fact that one can be at fault for failures in the exercise of these capacities. It is commonplace for an athlete or a musician who botches a performance because of the failure of these sorts of capacities to acknowledge fault for the mistake, even if she makes no mistake in the exercise of her agency antecedent to the mistake in non-rational aspects of the exercise of her technical skill.

Failures in the exercise of the capacities that determine these details do not warrant rational or intellectual criticism, like charges of irrationality or foolishness. A violin player who misses some notes, or a baseball player who throws a wild pitch, may be charged with

clumsiness or ineptness, but not foolishness or irrationality. The view that details of skilled action are determined in large part by non-rational capacities is confirmed by research on motor skills and non-declarative memory systems.¹² According to this research, much of our behavior is determined by specialized sub-personal cognitive systems that the individual does not have conscious access to, and which psychologists regard as modular and sub-personal. The exercise of these capacities does not amount to thought or reasoning by the agent.¹³

Interpersonal skills are partly constituted by non-rational capacities in a similar way as skill in sports and music. Details of interpersonal interaction play an important role in skillful interactions, just as details of hand movements play an important role in skillfully playing the violin, or skillfully playing baseball. Let me sketch the kinds of details of interpersonal interaction that I have in mind.

Some integral details of interpersonal interaction are aspects of one's manner of speech. These include, for example, levels and contour of volume and pitch, the timbre of one's voice, the speed and rhythm of one's speech, and the placement and duration of pauses in one's speech that give other participants an opening to take a turn in the conversation. Details of what one does with one's eyes also play an important role in interpersonal interactions. These aspects of interaction include gaze direction (looking down, or to the side while speaking or listening), when and how often one initiates eye contact, or looks away, how long one holds eye contact, and "soft" or "harsh" gaze. Social interaction also involves aspects of facial expression that are determined by mouth, eye, and brow position. Such expressions include smiles and frowns of different kinds, pursing the lips, frowning or raising the brow, and so on. Various aspects of body position are also involved in interpersonal interaction. Examples include forward lean, postural mirroring, degree of relaxation in one's posture, head tilt,

¹²See, e.g., Willingham(1998), Squire (1992) and (2009), Wolpert (2011), Gigerenzer(2007). See Ennen (2003) for a helpful discussion of the significance of this kind of research for anti-intellectualist accounts of skillful action.

¹³I discuss these issues in more detail in Chapter One, "Non-rational aspects of human agency."

crossing one's legs, and one's distance from and orientation toward other people.¹⁴

These details of interpersonal interaction have a variety of practical functions. I will mention a few examples to illustrate. Details of this kind can help signal a dominant or a submissive role in a social interaction. For example, superiors and subordinates often exhibit and reinforce the power relationships between them partly through differences in degree of relaxation in body posture, and length of turns in conversation. Similar sorts of behaviors can also help one to embody (often pernicious) cultural norms of masculinity and femininity. Details of interpersonal interaction express different degrees of like or dislike for a person. Liking can be expressed in part by smiling easily, leaning toward one's interlocutor, and frequently initiating eye contact with one's interlocutor. These sorts of behaviors can also exhibit friendliness, or interest in what another person is saying. Orienting oneself somewhat away from one's conversation partner, toward the door, can help one to end a conversation in a smooth, polite way. Not initiating too much eye contact, looking downward with some frequency, speaking at a relatively slow pace, and at a relatively low volume, can all be part of an individual's way of expressing contrition in making an apology. Similar behaviors can also be involved in expressing one's own sadness, or in expressing sympathetic sorrow at another's bad news. The placement and duration of pauses in one's speech can provide openings in conversation, and encourage reasonably equal turn-taking. In a group conversation, looking with some frequency at a diminutive participant during openings in the conversation can help to include him in the group.

Know-how concerning the embodiment of norms of masculinity, or being seductive, or charming or friendly, is partly constituted by non-rational capacities. Many details of one's tone and inflection in speech, body posture, gait, and so on, are not determined at the

¹⁴Several of the examples described below, illustrating the role of subtle behaviors in social interaction, come from an overview of nonverbal communication in Burgoon (1995). Some others come from Goffman (1959), Robinson (1998), Irish and Hall (1995), and Street (1987). See Wolpert (2003) for discussion that explicitly addresses the connections between interpersonal skill and motor control.

level of thought—just as many details of throwing a baseball or playing a violin are not determined at the level of thought. We should expect that interpersonal skill should have a substantial non-rational component, since, like skill in athletics and music, it is a skill in physical performance, which partly consists in the ability to determine details of one’s way of acting. If physical skills in other areas consist partly in non-rational action-guiding capacities, it is difficult to think of a reason why interpersonal skill would be different.

We do exert a certain amount of rational control over the development of the capacities that determine details of interpersonal interaction. A child might be told that, though she says hello to a visiting aunt, she ought to do it in a more friendly way. The child might intentionally look to the behavior of a sibling as a model of how to greet the aunt, on the basis of reasons. (Though I assume the usual case is that children “pick up” on the example of others to learn such things, and do not look to examples intentionally.) However, the capacities she develops from looking to the example are still largely non-rational capacities. In a similar way, a baseball player exerts rational control over her throwing accuracy by intentionally practicing her throwing, even though the capacities she acquires through practice are largely non-rational capacities.

One sign that interpersonal skills are not different is that we are often unaware of details of our behavior that exhibit our interpersonal skills. We often find it informative to have them pointed out to us. In my own case, for example, I found it informative to read that relatively upright, unrelaxed body posture is part of signaling deference and subordination, even though that had long been a part of my own way of acting in deferential and subordinate ways. Another sign is that it is possible to have detailed articulable knowledge of how to interact with others in a skillful way, but still lack interpersonal skill. One might have an intellectual grasp of how to act in a seductive or flirtatious way, but lack the skill to act in that way—just as one might have an intellectual grasp of how to dance a tango well, but lack the physical skill to actually do it.

The fact that details of interpersonal skill are often determined by non-rational capacities is also reflected in the kinds of evaluative attitudes that are appropriate in response to making mistakes with respect to these details. Consider an individual who means to act in a friendly, charming way at a party, but due to a lack of social skill, expresses uneasiness and a lack of interest in her interlocutor through her tone of voice, body posture, and so on. Such an individual would rightly be charged with a kind of ineptness, but not foolishness or irrationality.

Some of the functions of these kinds of details of interpersonal interaction are associated with positive ways of treating others, like being loving, compassionate, and considerate. The softening of one's eyes while talking with one's spouse may have the quality of being loving. Or changing one's direction of gaze so as not to rest one's eyes on an embarrassing deformity in one's interlocutor can manifest sensitivity to her feelings. Treating one's spouse lovingly, or being sensitive to another's feelings, are positive ways of treating others.¹⁵

Ethical success and non-rational aspects of interpersonal skill

I want to suggest that the non-rational capacities that determine integral details of these positive ways of treating others can contribute constitutively to technical aspects of ethical success. I present two lines of argument in this section. The first is based on the claim that non-rational aspects of interpersonal skill are capacities for agential control, and, accordingly,

¹⁵In some cases, particular details of interpersonal interaction seem to manifest positive qualities individually. In other cases, particular details are elements of a more over-arching way of treating others well. For example, consider the relation between an action of making an apology, and a particular detail of that action, like speaking at a relatively low volume, or at a relatively slow pace. Speaking at a slow pace may be one of several aspects of the agent's behavior which together make up her way of apologizing contritely. In such a case, that detail of her behavior (namely, speaking at a slow pace) does not itself possess the quality of being contrite. Instead, her pace of speech is one part of her over-all way of apologizing contritely.

and the claim that failures in the exercise of such capacities can warrant reactive attitudes that mark ethical failure. The second line of argument is based on the claim that details of interpersonal interaction determined by non-rational capacities can be integral to the ethical goodness of one's way of acting toward another person. Rationalist views about ethical success cannot adequately account for the ethical goodness of those details.

Non-rational aspects of agency and reactive attitudes Valuation aspects of ethical success, like one's values and reason-based attitudes toward another person, might indicate that one ought to treat the other person in, for example, a loving, or compassionate, or apologetic way. In such a case, it remains for the agent to determine what way of acting toward the other person would *constitute* acting in a loving, or compassionate, or apologetic way. Details of one's behavior that are determined by non-rational aspects of interpersonal skills can be integral aspects of appropriately loving, compassionate, or apologetic behavior. Those non-rational capacities are thus constitutive aspects of our ability to act in ways that are appropriately loving, compassionate, or apologetic. For example, in speaking at a relatively slow pace, at low volume, and with such-and-such intonation, we make use of know-how concerning the way to act apologetically. That know-how consists partly in the non-rational capacities we make use of to determine those details of the way we act. Another way of speaking could go wrong by sounding too casual, or overwrought. In looking at a child with "soft" eyes, or stroking his head a certain way, we make use of partly non-rational know-how concerning the way to manifest one's love toward a child. Another way of looking at the child, or stroking his hair, might be suited for romantic love, and inappropriate for parental love.

The ability to treat others in loving ways is ethically significant. Acting in a loving way, like acting in a generous or fair way, is partly constitutive of treating others in ethically good ways. The ability to treat others lovingly is part of one's ability to bring certain

ethically good attitudes and values to bear in one's conduct—like one's love for one's child, and one's desire to treat him with warmth and affection. Thus, abilities to bring to bear good attitudes and values in determining how to act are partly constitutive of our ability to treat others in ethically good ways. Non-rational capacities like those that determine the ethically appropriate tone of voice to use with a child are part of our ability to bring to bear good values and attitudes in determining how to act. Thus, non-rational capacities of that kind are partly constitutive of our ability to treat others in ethically good ways.

As suggested above, non-rational aspects of skill in music or baseball are capacities for agential control. If one accepts that point, it is not easy to find grounds for denying that non-rational aspects of interpersonal skill are likewise capacities for agential control. They are action-guiding capacities that are partly constitutive of interpersonal skill, where interpersonal skill is attributable to the agent at the whole-individual level. It is the agent herself who possesses skills concerning the way to treat people with compassion, to act sympathetically, and so on. In determining details of interpersonal interaction through the exercise of these non-rational capacities, we exert a kind of agential control over the way we act toward other people. Accordingly, we can be directly, non-derivatively culpable for failures in the exercise of capacities for agential control. When we fail to treat others well because of failures in the exercise of these non-rational capacities, we can be directly, non-derivatively culpable for those failures. It is plausible that a culpable failure to treat others in ethically good ways constitutes a kind of ethical failure. Thus, it is plausible that ethical success depends constitutively, and non-derivatively, on the exercise of non-rational aspects of interpersonal skill.

To support this claim, it helps to note that, unlike some other technical skills involved in ethically good action, skill in treating others in loving or compassionate ways is not ethically “neutral.” As an example of an ethically neutral technical skill, consider skill at swimming. One might make use of that skill in carrying out an ethically good action of saving a drowning

person. Skill at swimming, including its non-rational aspects, is partly constitutive of one's ability to carry out the ethically good action. There is a big difference between the role of skill at swimming and the role of interpersonal skill in ethically good action. The fact that an act of saving a person from drowning is an instance of proficient swimming does not figure in an explanation of the ethical goodness of the action. Aspects of skill at swimming that make the action an instance of skillful swimming do not thereby help to make the action ethically good.

In contrast, the fact that one's interaction with another person is an instance of treating that person lovingly does help to explain the ethical goodness of the action. The interaction is ethically good partly in virtue of its loving quality. So one's competence at getting the details of loving behavior right is part of an ability to provide for a quality of action that helps make the action count as ethically good. In this way, skill in treating others in loving ways is not ethically neutral like skill at swimming.

The claim that non-rational aspects of interpersonal skill contribute constitutively to ethical goodness finds some more support in the fact that details of behavior determined by those capacities warrant "reactive" attitudes. Details of interpersonal interactions may appropriately provide occasions for the sort of gratitude Strawson opposes to resentment. A loving look may appropriately occasion gratitude toward another person for treating one in a loving way. Or similar details of interaction may provide an occasion for forgiveness and rapprochement between people, as when the softening of a spouse's tone of voice, or her facial expression, helps to assuage one's resentment toward her. These details of behavior can warrant these sorts of reactive attitudes when they are non-intentional, and are determined by non-rational aspects of interpersonal skill. The fact that these details of behavior warrant reactive attitudes suggests that those details are ethically good aspects of one's way of acting. It is not reason, but non-rational aspects of interpersonal skill that determine these ethically good details of behavior. If so, that conclusion suggests in a *prima facie* way that ethical

goodness can consist partly in exercises of those non-rational capacities, and not just in exercises of rational capacities. To make the case more convincing, it is useful to look at an example in more detail.

Consider again an individual interacting with a close friend who is in need of comfort and support, because he is experiencing a strong wave of grief about the loss of a recently deceased parent. Suppose the individual expresses an appropriate degree of sympathy and empathetic sorrow, partially through details of interaction determined by non-rational capacities, like tone of voice, mode of eye contact, body language, and so on. Her way of interacting with her friend is compassionate, sympathetic, and supportive. In contrast, consider an individual in the same situation whose behavior differs only in details which, in the other case, are determined by non-rational capacities. Her words are the same, but her voice is tight, her body stiff, and her facial expression coldly blank. As a result of these features of the individual's subtle behavior, she fails to act sympathetically, in the way she should act toward her distressed friend. In a case like this, the agent's behavior warrants negative reactive attitudes like resentment, whereas in the first case the behavior warrants positive reactive attitudes like gratitude to the agent for being supportive and sympathetic. The appropriateness of reactive attitudes marks the fact that details of behavior are ethically defective in the second case, and constitute an ethical success in the first case.

To support the claim that non-rational capacities contribute constitutively to ethical success, we have to consider why the details of the agent's behavior differ in the second case. The behavior might differ because in the second case the agent does not care very deeply about her friend, does not value compassion, or is self-absorbed. In a case like that, the ethical failure is valuational, and is not due in the first instance to a failure in the agent's technical ability to get details of interpersonal skill right.

The behavior might also differ because in the second case the individual has an ongoing problem of making similar mistakes in her interpersonal interactions. In a case like that,

the agent might be culpable for a failure to deal with that ongoing problem, either by practicing to correct it, or by telling her friends about it and finding alternative ways to express sympathy, compassion, and the like. The failure to deal with the ongoing problem could be a failure in the exercise of the agent's rational capacities. If so, one can accept that there is a kind of ethical failure in the case, without accepting that ethical success depends constitutively on non-rational capacities, even in part. The agent's inadvertent mistreatment of her friend is ethically defective in a derivative way. It derives from a failure in exercises of the agent's rational capacities, which is antecedent to her inadvertently hurtful way of acting toward her friend.

But the locus of the ethical failure need not be antecedent to the failure in the exercise of interpersonal skill in either of the ways described in the last two paragraphs. As suggested above, interpersonal skill, including non-rational aspects of interpersonal skill, constitutes a kind of agential control—like non-rational aspects of skill in baseball and music. Accordingly, we can be directly and non-derivatively culpable for failures in the exercise of non-rational aspects of interpersonal skill.

Consider a case where the agent fails to act sympathetically because she is tired after a long day of work. Suppose she is not so tired that it is unreasonable to expect that she can act properly if she makes a little extra effort, but tired enough to make failure a real, intelligible possibility—even if the extra effort is made. Suppose she makes the required effort, but flubs in her attempt to convey compassion through the exercise of her interpersonal skill. Since she is making the required effort to compensate for her fatigue, and it is reasonable to expect that she can make use of her interpersonal skill to act in an appropriately sympathetic way, she is directly culpable for her failure if she errs. It is well within her power to get the details right and treat her friend in an appropriately sympathetic way. As a result, she is culpable for her failure to treat her friend as he ethically ought to be treated.

It is commonplace for a musician who misses notes, or a tennis player who makes an

unforced error, to acknowledge fault for the mistake, even if she makes no mistake antecedent to failing in the exercise of non-rational aspects of her technical skill. Similarly, a person who fails to act in an appropriately sympathetic way at the end of a long day would rightly acknowledge fault for her mistake, even if she makes no mistake antecedent to failing in the exercise of non-rational aspects of her interpersonal skill. In such a case, the agent culpably fails to treat her friend in an ethically good way. Accordingly, her failure warrants reactive attitudes, like feeling hurt or let down by the way her friend acts toward her. The failure consists in a failure in the exercise of non-rational aspects of her interpersonal skill. Thus, ethical success can depend constitutively and non-derivatively on the exercise of non-rational capacities.

To be clear, the failure in the case described is technical, and not valuational. It is a failure in the exercise of the agent's ability to bring to bear her ethically good values and attitudes in action. An agent's intention to treat a friend in a compassionate way leaves open how in particular the agent will treat the friend. Her ability to determine how in particular to treat the friend, so as to act compassionately toward him, consists partly in rational capacities, and partly in non-rational capacities. She determines how in particular to treat her friend partly through exercises of rational capacities, like capacities to make use of reasons to determine whether or not to tell a story about the friend's deceased parent. But she also determines how in particular to treat her friend through the exercise of non-rational aspects of interpersonal skill, like abilities to control details of tone of voice, and body language.

When one makes a mistake in the exercise of these abilities, one does not thereby fail to have good values, or to care adequately about a friend, or the like. Ethical failures like those are more serious, and warrant different, more dire reactive attitudes. But technical failures to determine how in particular to bring one's good values to bear in action—including failures due non-derivatively to mistakes in the exercise of non-rational aspects of interpersonal

skill—can constitute culpable failures to treat others as they ethically ought to be treated. It is reasonable to count such a failure as a kind of ethical failure, so long as it is not confused with a valuational failure to have good values, to have good attitudes toward others, or the like.

The ethical goodness of details of interpersonal interaction I turn now to another way of supporting the claim that non-rational capacities can contribute constitutively and non-derivatively to ethical success. The argument I want to make here proceeds from the assumption that the details of interpersonal interaction that make an action appropriately loving, sympathetic, or the like, are ethically good aspects of ethically good ways of acting. If one accepts that assumption, there is a strong case to be made that non-rational action-guiding capacities contribute constitutively and non-derivatively to ethical success.

According to the traditional rationalist view about ethical success, an action is ethically good in virtue of its grounding in the agent's rational recognition of reasons for that action. This view implies that, if an action is ethically good, the ethically good way one acts must be conceptually represented by the agent at the level of reason. If an individual treats her friend in an appropriately sympathetic way, she might conceptually represent that she is treating her friend in a sympathetic way, that she is speaking to her friend, that she is holding her friend's shoulder, and so on. But many integral details of her appropriately sympathetic action are not represented at the level of reason. Many such details are determined by non-rational capacities. According to the traditional rationalist view, these details of her action would not be ethically good, since they are not represented at the level of reason. Since they are not represented at the level of reason, those details cannot be grounded in a rational recognition of reasons for those details of her action. If one accepts that non-rationally determined details of interpersonal interaction are ethically good aspects of ethically good action, then one should revise the traditional rationalist view.

Let me consider a response to this argument. A rationalist might concede that the agent in the example does not conceptually represent all the integral details of her way of acting which help to make the action count as an ethically good, appropriately sympathetic way of treating her friend. But the rationalist might emphasize that the agent does conceptually represent her action as an instance of treating her friend in a sympathetic way. The agent succeeds in acting in an appropriately sympathetic way. Her sympathetic way of acting consists partly in details of interpersonal interaction that she does not represent at the level of reason. The rationalist might then claim that the agent's particular way of acting, which constitutively involves those details, is ethically good because it constitutes sympathetic action, and the action is based on a rational recognition of grounds for acting in a sympathetic way. The action is ethically good because the action "matches" (i.e., is a successful execution of) an ethically good intention—where the intention is ethically good exclusively in virtue of exercises of the agent's rational capacities. That connection to exercises of rational capacities is what makes the action count as ethically good. In this way, one might try to preserve rationalism about ethical goodness while accepting that details of action can be ethically good even if they are determined by non-rational action-guiding capacities.

This line of reply seems to me unconvincing. According to traditional rationalist views, an ethically good action counts as ethically good because of the rational control we exert over those actions, by recognizing rational grounds for that particular way of acting. The agent's rational control over her action leaves open integral details of her particular way of acting. The agent intends to act in an appropriately sympathetic way, and she acts in a way that is in fact appropriately sympathetic way. But she does not have a substantive grasp, at the level of reason, that the particular way she is acting constitutes an appropriately sympathetic way of acting. Thus, she does not exert rational control over her particular way of acting, beyond forming a relatively generic reasons-based intention to act sympathetically—as well as reasons-based intentions that partially specify particulars concerning how to act

sympathetically, like an intention to hold her friend's shoulder.

Since the agent's rational control does not go beyond that relatively generic intention to act sympathetically, and a partial specification of the particulars of how to carry it out, the rationalist's description of the agent's ethically good action should be correspondingly limited. For the rationalist, the agent's ethically good action would have to be that she attempts to act in a sympathetic way, that she holds his shoulder, and so on. But particulars of the action that are not specified at the level of reason should not count as ethically good, according to the rationalist view, since the agent does not exert rational control specifically over those details. But the particulars of the agent's way of acting sympathetically do count as ethically good. So it cannot be only exercises of rational capacities that make the sympathetic agent's action count as ethically good.

To bring out the point here, it will help to consider a related point about a rationalist view about action. According to this rationalist view about action, an event counts as an action, or as a manifestation of agency, exclusively in virtue of its relation to exercises of rational capacities for agential control. So, for example, raising one's hand might count as an action in virtue of its relation to a reasons-based intention to raise one's arm. This rationalist view about action conflicts with the fact that many actions, or details of actions, that manifest agency, are determined by non-rational action-guiding capacities, and are not specifically represented at the level of reason. For example, certain details of a violinist's finger placement are manifestations of agency even though they are not represented by the agent at the level of reason. If exercises of one's rational capacities leave open whether one will act in a particular way, rather than in some other way, then it cannot be just exercises of rational capacities that make one's particular way of acting count as a manifestation of agency.¹⁶

Rationalism about ethical success can be understood as a variant of this rationalist view

¹⁶I develop this argument in more detail in Chapter One, "Nonrational aspects of agency."

about action. The rationalist view about ethical success holds that an action counts as ethically good in virtue of its relation to exercises of rational capacities for agential control, wherein the agent recognizes rational grounds for acting in the ethically good way. The problem is that details of action that are left open by exercises of one's rational capacities cannot count as ethically good just in virtue of their relation to one's rational capacities. The agent does not specifically determine those details of her action through reason, and so does not specifically exercise rational control over those particular details of her action. If one accepts that details of interpersonal interaction can be ethically good even though they are not specified by exercises of rational capacities, then the rationalist view about ethical success should be revised.

I want to conclude by considering another way to try to vindicate rationalism about ethical success. A rationalist might claim that non-rationally determined details of interpersonal interaction count as ethically good because of the "normal and natural" connection between the details of the action, and exercises of the agent's rational capacities. If the agent possesses the required interpersonal skills, then the agent's reason-based intention to treat her friend in a sympathetic way issues in action that gets the particular details right *through* the normal proper functioning of the agent's action system. The rationalist might claim that the ethical goodness of exercises of the agent's rational capacities can "transmit" or "permeate" to the particular details of the action because of that connection between the details and the exercises of rational capacities. That is, the agent's particular way of acting counts as ethically good because it issues from ethically good exercises of rational capacities, through the normal proper functioning of the agent's action system. In this way, the rationalist might say that a particular way of acting can count as ethically good in virtue of its connection to exercises of reason, while accepting that the agent does not represent integral details of that way of acting at the level of reason.

According to the proposal on offer, the details of interpersonal interaction count as ethi-

cally good because they issue from the agent's ethically good reason-based intention through to the normal proper functioning of the agent's action system. I think that proposal is true, as far as it goes. But it is not ultimately helpful for defending a rationalist position about ethical success.

As a first step in responding to the proposal, note that exercises of non-rational aspects of interpersonal skill are integral to the functioning of the agent's action system that connects the details of the action to the intention. As a result, according to the proposal on offer, those exercises of non-rational capacities help make it the case that the details of the action count as ethically good.

Now, there is one way that non-rational capacities can help make our actions count as ethically good, which is compatible with rationalist claims about ethical success. But that is not the way that non-rational aspects of interpersonal skill help make our actions count as ethically good.

To develop this point, let me start by explaining the first way that non-rational capacities can help make our actions count as ethically good—the way that is compatible with rationalist claims about ethical success. Non-rational capacities can help make an action count as ethically good by “gluing” her overt action to her ethically good intentions, reasons, and values, so that “inner” ethical goodness “transmits” to the agent's overt behavior. Consider an individual's ethically good action of raising her hand, say to signal that she volunteers to take on a dangerous task for the benefit of others. Part of what makes the action ethically good is that it issues from the agent's ethically good intention to volunteer through the normal proper functioning of the agent's action system. If the arm movement were a spasm, and not appropriately connected to the intention, then the ethical goodness of the intention would not “transmit” to the arm movement. The exercise of the agent's capacity to move her arm is integral to the functioning of the agent's action system that connects the arm movement to the intention in the right way. Thus, the exercise of the agent's non-rational

capacity to move her arm helps to make it the case that the arm movement is ethically good action.

This connection between non-rational capacities and ethical goodness is consistent with rationalist claims about ethical goodness. The rationalist view is that actions are ethically good exclusively because of their connection to exercises of rational capacities. According to this rationalist view, it is not just exercises of rational capacities that help make an action count as ethically good. The *connection* between the action and the exercises of rational capacities also plays an integral role. According to the rationalist view, the primary locus of the ethical goodness is the exercises of rational capacities. The connection between the action and the rational capacities is, as it were, a “glue” that binds the goodness of the exercises of rational capacities to the overt action, by making the overt action an *execution* of the agent’s ethically good intention.

However, the role of non-rational aspects of interpersonal skill in providing for the ethical goodness of our actions goes beyond this “glue” role. Those non-rational capacities do not help make an action count as ethically good just by helping to make one’s bodily movement count as an execution of an ethically good intention. Exercises of interpersonal skill help make it the case that one is acting in one way rather than another. They help make it that case that one is acting, for example, in an appropriately sympathetic way, rather than in an inadvertently hurtful way, by guiding one to act in the one way rather than the other.

To bring out the significance of this fact, it helps to note that, if an individual has poor interpersonal “technique,” the wayward details of her behavior can be “glued” to an ethically good intention in much the same way as if the agent were interpersonally competent. A person who grew up with bad examples of interpersonal interaction may have picked up poor technique in expressing sympathy and compassion to others. When she makes use of that poor technique, her particular way of acting is her way of trying to express sympathy and compassion toward the other person. The details of her behavior are not accidentally

related to her intention to act sympathetically in the way that an arm spasm might be accidentally related to an intention to volunteer.

This point suggests that, if the contribution of non-rational aspects of interpersonal to ethical goodness were by “gluing” the action to the agent’s reasoning, then the details of action in the competent case would not genuinely be ethically good. Appropriately details of interpersonal interaction, determined by interpersonal skill, are “glued” to the agent’s intention to act in a sympathetic way. But misguided details of interpersonal interaction, determined by non-rational aspects of poor interpersonal technique are equally “glued” to one’s intention to act in a sympathetic way. So, if details of sympathetic action are ethically good just in that they are “glued” in the right way to ethically good exercises of reason, then it should not matter whether the details are competent or misguided. In that case, the particular details in the ethically good case would not genuinely be ethically good.¹⁷

Non-rational aspects of interpersonal skill help to make one’s action ethically good by guiding the agent to act in an ethically good, appropriately sympathetic way (to continue with that example), instead of acting in some other way, that is ethically defective. Non-rational capacities help make our actions count as ethically good through their role as part of the agent’s technical ability to determine how in particular to act, so as to act in accord with one’s ethically good values and attitudes.

A rationalist might want to deny that particular details of interpersonal interaction, determined through non-rational capacities, really are ethically good aspects of ethically good ways of acting. The rationalist might claim that the agent’s ethically good action, perspicuously characterized, is to make a good-faith attempt to treat the agent in an appropriately

¹⁷Perhaps exercises of non-rational aspects of poor interpersonal technique do help make actions count as ethically good in a certain way. A misguided attempt to treat another person in a sympathetic way is ethically good in certain important respects. Such an action manifests, or expresses, valuational aspects of ethical success, and is ethically good in that respect. Exercises of the agent’s poor interpersonal technique help make it the case that the action is an expression of those valuational aspects of ethical success, by helping to make the action count as an execution of an intention that is based on good values and a caring attitude toward one’s friend.

sympathetic way, where that attempt is based on good values, a caring attitude toward her friend, and so on. The rationalist might claim that particular details of how the agent carries out this attempt are ethically neutral concomitants of her ethically good action. This rationalist response is plausible if one thinks of ethical success as consisting exclusively in what I have been describing as *valuational* aspects of ethical success. But if one accepts that there are also technical aspects of ethical success, then one's particular way of bringing to bear ethically good values and reason-based attitudes is partly constitutive of ethical success. If so, it is more difficult to find convincing grounds for denying that non-rational capacities contribute constitutively and non-derivatively to ethical success.

Traditional rationalist views in ethics neglect the constitutive role of non-rational capacities in treating others in ethically good ways. This neglect is due in part to the fact that rationalist ethics, at least since Kant, has been primarily focused on valuational aspects of ethical success. This focus is reasonable. The core of ethical success is valuational. But there is more to ethical success than its core. To succeed in treating others well, more is required than having good values and attitudes toward others. One must determine how in particular to act so as to bring one's good values and attitudes to bear. Our agency over these technical aspects of ethical success consists partly in non-rational capacities. As a result, non-rational capacities contribute constitutively to technical aspects of ethical success.

Chapter 4

Conclusions of practical reasoning

Introduction

In *On the Movement of Animals* Aristotle describes actions as conclusions (*sumperasma*). He is widely interpreted as having in mind that actions are conclusions of practical reasoning. Many are attracted to the idea that there is some constitutive connection between practical reasoning and action, which goes beyond action's role as the subject matter of practical reasoning. It is part of the nature of practical reasoning that it issues, or functions to issue, in action. Practical reason's relation to action is different from the relation between zoology and the animals that zoologists study. One appealing feature of the claim that actions are conclusions of practical reasoning is that it bears out this attractive view. If actions are conclusions of practical reasoning, then actions are not just the characteristic subject matter of practical reasoning.

A more common view is that conclusions of practical reasoning are intentions. But sympathy for the claim that conclusions of practical reasoning are actions is fairly widespread.¹

¹Compare, for instance, Dancy (2004), Tenenbaum (2007), Streumer (2010), Anscombe (1965), Vogler (2001).

In this chapter, I argue that conclusions of practical reasoning are intentions, and not actions. While this view is common, there are few sustained arguments for it. There are also few sustained arguments against the claim that the conclusions are actions. Most extant arguments against the view that actions are conclusions of practical reasoning seem to me to miss the mark, and to neglect the core problems with the claim that actions are conclusions of practical reasoning.

Many arguments against the claim describe cases in which an agent fails to act even though there is supposed to be nothing incomplete or defective about her reasoning. Failure to act due to forgetfulness and weakness of will are supposed to be central examples of this kind.² In my view, these failures to act are due to failures in the exercise of capacities constitutively involved in reasoning. Forgetfulness is due to a failure in memory, which is constitutively involved in reasoning. Weakness of will is, in my view, due either to a failure to form an intention, or to form an intention with enough strength. As I argue below, forming an intention is an element of practical reasoning. Forming an intention with sufficient strength is also, in my view, a constitutive aspect of practical reasoning.

The arguments I present against the claim that actions are conclusions of practical reasoning center on two main points. First, the claim that actions are conclusions of reasoning, including physical actions, is incompatible with the point that reasoning is a psychological phenomenon. Second, the claim that actions are conclusions of reasoning is incompatible with the point that reasoning is a kind of figuring something out by making use of reasons. The first section of this chapter criticizes the view that actions are conclusions of practical reasoning by developing these two points. In the second section, I support the view that intentions are conclusions of practical reasoning by defending it from several objections due to Sergio Tenenbaum and Joseph Raz.

²See, for example, Streumer (2010), Raz (2010).

Conclusions of practical reasoning are not actions

As noted above, it is common to focus on the idea that practical reason is not sufficient for action in criticizing the claim that actions are conclusions of practical reasoning. It will be helpful to make explicit why that claim requires that practical reason is sufficient for action. The claim implies that requirement if one assumes that the conclusion of a piece of reasoning is an element in that reasoning. Elements in one's reasoning include steps in one's reasoning, and mental states that one "steps" to or from in one's reasoning. Transitioning to the conclusion is the terminal "concluding" step in that reasoning. Since a conclusion of a piece of reasoning is an element of the reasoning, exercises of one's capacities to reason are sufficient for arriving at the conclusion. So if action is the conclusion of practical reasoning, then one's capacity to reason should be sufficient for action. Accordingly, any failure to act should be due to a failure of one's capacity to reason.

Now, if one assumes that a conclusion of practical reasoning is an element of the agent's reasoning, that assumption implies that the conclusion is an element of a psychological phenomenon.³ Reasoning is a psychological phenomenon. One reasons through making use of mental capacities. One does not reason with one's arms and legs.⁴ Since the conclusion is an element of the reasoning, the conclusion is an element of a psychological phenomenon, and so it is itself an element in the agent's psychology. So if actions are conclusions of

³This point calls for clarification. One might suggest that a proposition can be an element of reasoning, since propositions are contents of psychological states involved in one's reasoning. As I use the term here, such psychological states are elements of one's reasoning, but the contents of the psychological states are not elements of one's reasoning. This terminological choice helps make exposition less cumbersome. The terminology does not make any substantive difference in the discussion here. The view that actions are conclusions of practical reasoning holds that the relation between action and practical reasoning is analogous to the relation between belief and epistemic reasoning. It does not hold that the relation between action and practical reasoning is the same as the relation between epistemic reasoning and the content of one's belief. (Philip Clark's (2001) view is an exception. He holds that actions are contents of psychological states that practical reasoning terminates in, and expresses that view by saying that actions are conclusions of practical reasoning.)

⁴Andy Clark and David Chalmers (1998) dispute this claim in advocating their 'extended mind' view. I consider the bearing of their view on the present issue on pages 18-20.

practical reasoning, actions based on reasoning, including intentional physical actions, are psychological phenomena. But a physical action like walking is no more a psychological phenomenon than is a dog, or Mont Blanc with its snow fields. Walking consists, at least partly, in putting one foot in front of the other. One's feet are not a part of one's psychology. Since physical actions are not psychological phenomena, they are not conclusions of practical reasoning.⁵

Of course, our physical actions often do have deep links to our psychologies. They are often based on certain elements of our psychologies, like our beliefs and desires. And they are often expressive of certain elements of our psychologies, like our values and our emotions. Indeed, it is plausible that many physical actions are individuated partly by reference to underlying elements in our psychologies, like intentions, and one's reasons for action. Whether an event counts as an intentional action, and what intentional action one has carried out, can be constitutively determined in part by whether the event executes an intention, and what intention in particular it executes. One's falling might or might not count as an action, and might count as a joke or an act of deceit, depending on the event's relation to one's psychology. Some physical actions may even be partly *constituted* by elements of our psychologies. Perhaps the psychological functioning involved in negotiating turns when moving a couch up the stairwell is part of the overall physical action of moving the couch.

However, psychological functioning is at most *part* of the physical action. A physical action like moving a couch also constitutively involves the movement of one's arms and legs. Psychological functioning does not consist, even in part, in the movement of one's arms and legs. Of course, some kinds of physical functioning do play certain roles in psychological functioning. The physical processes of one's neural system subvene psychological functioning,

⁵John Broome makes a similar point. He says, that action—at least physical action—“requires more than reasoning ability; it requires physical ability too” (2002, p. 1). Broome's point seems to me correct and helpful, but his discussion of the point does not go beyond that one sentence. In addition to developing this point about physical ability and reasoning in greater detail, the discussion in this section goes beyond Broome's remark by arguing that intentional mental actions also are not conclusions of reasoning.

including practical reasoning and epistemic reasoning. But it is clear that physical phenomena that *subvene* a piece of reasoning cannot be a conclusion of reasoning. At most, such a physical phenomenon could subvene the concluding step of one's reasoning.⁶ According to materialist reductionist views, psychological functioning *is* physical neural functioning. But even from a materialist reductionist point of view, it is quite implausible to claim that the physical movement of one's arms and legs is psychological functioning. As a result, even for a material reductionist, physical actions, like actions that constitutively involve movements of one's arms and legs, cannot be conclusions of practical reasoning.

Mental actions, practical reasoning, and figuring out what to do

The preceding argument specifically concerns physical actions, and does not apply directly to mental actions. As a rearguard maneuver, one might concede that physical actions are not conclusions of practical reasoning, but insist that mental actions are conclusions of practical reasoning.

Let me make a clarification about this rearguard view. Forming a judgment or an intention is an active psychological phenomenon, at least in some cases. One might consider the active formation of those mental states to be a kind of mental action.⁷ However, if practical reasoning concludes in the formation of some kind of practical judgment, or if it concludes in the formation of an intention, that does not vindicate the view under discussion here. The view is that *intentional* mental actions, which one's reasoning is *about*, are conclusions of reasoning. For example, suppose an individual intentionally imagines herself on a beach on the basis of practical reasoning about how to cope with a wave of anxiety. The rearguard

⁶Moreover, except perhaps in unusual cases, the physical functioning constitutive of physical action does not subvene psychological functioning. The movements of one's arms and legs do not subvene psychological functioning.

⁷I prefer to say that forming such a mental state is a mental "act" instead of an "action." I prefer to reserve the term 'action' for intentional activity, and pre-rational analogs of intentional activity.

view does not claim that the conclusion of the reasoning is an intention, where the intention constitutes a kind of mental action. The rearguard view claims that the agent's intentional mental action of imagining herself on the beach is the conclusion of her reasoning.

The argument in the last section does not apply directly to intentional mental actions like intentionally imagining oneself on a beach. Physical actions cannot be conclusions of reasoning because physical actions are not (or not just) psychological phenomena. But intentional mental actions are psychological phenomena, so they cannot be ruled out as conclusions of practical reasoning on those grounds. Considering the possibility that intentional mental actions might be conclusions of practical reasoning will help to bring out what is, in my view, the core problem with the claim that actions are conclusions of practical reasoning.

Before presenting that core problem, I want to make a preliminary argument against the claim that mental actions are conclusions of practical reasoning. If it is granted that physical actions are not conclusions of practical reasoning, it is not very plausible that intentional mental actions are conclusions of practical reasoning. A conclusion of practical reasoning occupies a certain role in the structure of practical reasoning. Intentional mental actions and intentional physical actions do not seem to occupy different roles in the structure of practical reasoning.⁸

Intentional mental actions based on reasoning, like intentional physical actions based on reasoning, are executions of something like an intention or a decision, where the reasoning issues in the intention or decision antecedently to issuing in the action. Since intentional mental actions and intentional physical actions seem to be similarly related to practical reasoning, it would seem they are either both or neither conclusions of practical reasoning.⁹

⁸I say "*intentional* physical action" here to emphasize the parallel with intentional mental action. I do not mean to be focusing here on different physical actions from those discussed in the last section, even though I did not explicitly describe the physical actions in the last section as intentional. The physical actions under discussion there are also intentional physical actions, carried under the guidance of practical reasoning.

⁹Streumer (2010) suggests that actions are sometimes conclusions of practical reasoning, and in other cases conclusions of practical reasoning are intentions or judgments. It seems plausible that there might be

If one accepts that physical actions are not conclusions of practical reasoning, one would need a compelling explanation of why an intentional action's status as physical or mental should make a difference to whether it is a conclusion of practical reasoning. It would be unconvincing merely to point out that the main argument against physical actions' being conclusions does not apply to mental actions.

Let me turn to a more important point. A natural assumption about the nature of reasoning provides a strong case for denying that actions can be conclusions of practical reasoning, where that case applies directly to both mental and physical actions. The assumption is that reasoning of any kind, including both practical and epistemic reasoning, is roughly a kind of "figuring something out." Epistemic reasoning is an exercise of certain psychological capacities to figure out what is the case. Practical reasoning is an exercise of psychological capacities to figure out what to do. The fact that reasoning is a kind of "figuring something out" helps shed light on what a conclusion of reasoning is.

The conclusion of a piece of practical reasoning is the terminus of this kind figuring out what to do. As it were, the conclusion of practical reasoning is the "completion" of one's figuring out what to do, whereby she has figured out for herself what to do. It is a step in the reasoning which makes it appropriate to answer "yes" to the question "did you figure out what to do?"—like the terminal phase in building a house makes it appropriate to answer "yes" to the question "did you build the house?"

Pamela Hieronymi's (2006) notion of "settling a question" is useful showing the appeal of this point. The concluding step in figuring out whether to ϕ is the step wherein the

different *kinds* of conclusions of practical reasoning, roughly corresponding to different respects in which an individual's practical reasoning might be considered conclusive. The practical reasoning of an individual who judges that she ought to leave the bar might be conclusive in a way, even if she does not decide to leave the bar. The practical reasoning of an agent who actually decides to leave the bar might be conclusive in another way. Perhaps judgments and decisions (or intentions) can be considered different kinds of conclusions of practical reasoning, which occupy different roles in practical reasoning. But it seems to me very doubtful that the same kind of conclusion of practical reasoning can sometimes be an intention, sometimes and judgement, and sometimes an action.

agent settles for herself the “question” that the reasoning is about. If, for example, an agent engaging in practical reasoning is figuring out whether to go for a walk, the conclusion of the reasoning is the step of the reasoning wherein the individual settles for herself the “question” of whether to go for the walk.¹⁰ One can also put the point by saying that the conclusion of one’s reasoning is the step wherein one “makes up one’s mind” about the matter that the reasoning is about. The conclusion of reasoning about whether to go for a walk would be, then, the step wherein one makes up one’s mind about whether to go for a walk.

The vague-but-intuitive idea that reasoning is some kind of figuring something out can be captured, at least partially, in the following claim: reasoning consists in the exercise of psychological capacities that function to guide committal attitudes, through making use of reasons.¹¹ I think it is a conceptual truth that reasoning consists in this kind of guidance. I hope that the intuitive point that reasoning is a kind of “figuring something out” helps to bring out this conceptual point. Figuring something out is “completed” in arriving at a committal attitude concerning the matter that one’s reasoning is about. The terminal step in practical reasoning, in figuring out what to do, is a step wherein one arrives at some kind of committal attitude concerning whether to perform the action that the reasoning is about. Thus, whatever a conclusion of practical reasoning is, it must be some kind of committal attitude concerning whether to perform the action that the reasoning is about.

The intentional action that one’s practical reasoning is about is not a commitment about whether to perform that action. To form a committal attitude, whereby one figures out for

¹⁰This terminology is useful, but it can mislead. I assume that, in most cases of practical reasoning, there is no psychologically real phenomenon of entertaining and settling a question about what to do.

¹¹Practical reasoning has a function of guiding actions, in addition to its function of guiding committal attitudes concerning what to do. Practical reasoning *consists* in the guidance of committal attitudes concerning what to do—i.e. in one’s figuring out what to do. There are capacities constitutively involved in reason’s guidance of action, which are not constitutively involved in the guidance of committal attitudes about what to do. Motor skills are an example. Practical reasoning does not consist in the exercise of capacities like these, even in part. I discuss the relation between practical reason and capacities like these in “Entitlement to rely on physical abilities,” and “Reasons’s guidance of action.”

oneself whether to perform an action, is different from actually carrying out that action. The difference is analogous to the difference between making a promise and fulfilling it. This point is easy to see when it comes to physical actions. Forming a committal attitude concerning whether to go for a walk is not the same thing as actually going for a walk. They cannot be the same since forming a committal attitude to go for a walk is a psychological phenomenon, whereas walking is a physical action that constitutively involves the use of one's legs. This point is confirmed by the fact that the figuring out can occur without the acting. For example, if one intends to go for a walk, but has unwittingly been immobilized by a dose of curare, one will have figured out for oneself whether to go for a walk, but one will not have carried out the action.¹²

The same points hold for mental actions too. Imagining oneself on the beach is not the same thing as forming a committal attitude about whether to engage in that exercise of imagination. Unlike that action itself, forming the committal attitude is not an act of imagination.¹³ If practical reasoning is a kind of figuring out what to do, then the concluding step of practical reasoning is a *particular kind* of psychological phenomenon. It is a transition to some kind of committal attitude concerning whether to perform the action that the reasoning is about. Though imagining oneself on the beach is a psychological phenomenon, it is not a psychological phenomenon of that kind.

In my view, the core problem with the claim that actions are conclusions of practical reasoning is that carrying out an action involves more than just figuring out whether to carry out that action. The conclusion of practical reasoning about whether to ϕ is part of one's figuring out whether to ϕ . So ϕ ing is not the conclusion of one's reasoning about

¹²One will not even have *begun* to act in a case like that. The argument presented here helps to explain why beginning to act is not the conclusion of reasoning. Beginning to act is not a committal attitude concerning whether to perform that action. Since practical reasoning is a kind of figuring something out, it concludes in some kind of committal mental state. Thus, beginning to act is not the conclusion of practical reasoning.

¹³However, as noted at the beginning of section 2.1, forming the commitment can be an act in its own right.

whether to ϕ .¹⁴

Question begging? One might worry that the line of argument presented here begs the question against those who claim that actions are conclusions of practical reasoning. An advocate of that view might not accept that practical reasoning is a kind of figuring something out. She might accept that practical reasoning consists *partly* in this kind of thinking about what to do, but insist that practical reasoning also consists partly in *acting* on the basis of that kind of thinking.

It is worth making clear that this response implies that practical reasoning is not just a psychological phenomenon. It implies that one can reason, in part, with one's arms and legs, and not just with one's mind. If one accepts that practical reasoning is a psychological phenomenon, then it is difficult to find grounds for denying that it is roughly the sort of psychological phenomenon sketched above—a kind of figuring something out. So, if one accepts that practical reasoning is a psychological phenomenon, then the argument presented above should go through, and neither physical actions nor intentional mental actions that one's reasoning is about would be conclusions of practical reasoning.

The question, then, is whether or not reasoning is a psychological phenomenon. Disagreement about this issue may be terminological. All sides should agree that there is such a thing as practical *thought*—the kind of thought that our intentional actions are based on when they are based on reasons. Both sides should also agree that, in addition to our capacity for that kind of practical thought, we also have a capacity to act on the basis of reasons. That capacity to act on the basis of reasons comprises both our capacity for practical thought, and a capacity to actually carry out the actions “recommended” in our practical thinking. Action may be a concluding stage in exercises of that larger capacity to act on the basis of

¹⁴Audi (1989, p 71) makes a somewhat similar point. He says that the Aristotelian view makes practical reasoning a “hybrid process composed of what is intuitively reasoning and, on the other hand, action based on it.

reasons.¹⁵ But an action is not a conclusion of practical thinking about whether to perform that action.

Keeping in line with standard usage of the word ‘reasoning,’ I use the phrase ‘practical reasoning’ as a term for a kind of practical thinking, wherein one figures out for oneself what to do. One might insist on using the phrase as a term for exercises of a larger capacity to act on the basis of reasons. This terminological choice would mark a dramatic departure from traditional conceptions of reason. Traditionally, ‘reason’ and ‘reasoning,’ like ‘*Vernunft*,’ ‘*cogitare*,’ and, I would argue, ‘*logos*,’ and ‘*bouleusis*,’ refer to a cognitive capacity and its exercise. The point I want to make here is that practical reasoning, understood as a cognitive capacity, does not conclude in the action that the reasoning is about.

It is also worth pointing out another problem with the view that some constitutive aspects of reasoning are non-psychological, besides its conflict with traditional usage. That view conflicts with natural assumptions about the link between rational capacities and culpable rational error. It is natural to assume that an agent can be rationally criticizable in virtue of failures of the exercise of any constitutive aspect of one’s capacities for reasoning. If one carries out a fallacious inferential transition, one’s reasoning can thereby be rationally criticizable. Or, if one fails to form a belief that p when one recognizes that the evidence clearly establishes that p, one can thereby be rationally criticizable. Or, to take a third example, if one fails to bring to bear some piece of obviously salient consideration in one’s reasoning, one can thereby be criticized as irrational, or foolish.

Every constitutive aspect of one’s rational capacities contributes constitutively to the quality of one’s reasoning. Accordingly, one can fail to reason well in virtue of failures of any constitutive aspect of one’s capacities for reasoning. In such cases, one can be criticized as irrational, foolish, or the like, for the failure of that aspect of one’s capacities for reasoning. However, reflection on examples reveals that, if one is rationally criticizable for some mistake

¹⁵I develop a view roughly along these lines in “Reason’s guidance of action.”

in action, it is always in virtue of some psychological failure.¹⁶ If an individual fails to act as she ought to purely because of a failure of her body to function properly, and she made no mistake in her thinking about how to act—her use of reasons in forming commitments about what to do¹⁷—then she is not rationally criticizable.

For example, suppose an individual tries to carry out some well-chosen action, like loading an infirm man’s groceries into his car. Suppose that the agent’s elbow locks as she is trying to load the groceries in, causing her to spill them on the pavement. Suppose further that she made no error in her thinking about what to do in deciding to help load the groceries. In particular, suppose she had no evidence that she had an elbow problem, and might end up spilling the intended beneficiary’s groceries, instead of helping him. In a case like this, the agent is not rationally criticizable. On the other hand, if she did have evidence that she had an elbow problem that could cause the action to go wrong, and thus had made some error in her thinking about how to act, then she could be rationally criticizable.

Reflecting on further cases reveals that this pattern holds in general. If an error in action is due to a pure failure of the body, and there is no psychological error in the offing, then the agent is not rationally criticizable. Conversely, in any case in which one is rationally criticizable, the criticism is warranted by some failure in the agent’s thinking about how to act. As a result, it would seem that an agent can only be rationally criticizable in virtue of psychological errors in the way one makes use of reasons to guide conclusions about what to do.¹⁸ This point suggests that one can only fail to reason well in virtue of failures in the operation of one’s psychological capacities. Thus, one’s capacities for reasoning comprise

¹⁶The point here is closely connected to the criticism of Tenenbaum’s argument below, on pages 91-95.

¹⁷The progressive aspect of the gerund ‘thinking’ suggests that it is a phenomenon that unfolds in steps over time. What I have in mind here is any exercise of one’s capacities to hold commitments on the basis of reasons. I do not assume that an individual who holds a commitment on the basis of reasons must have gone through a process of reasoning to that commitment, which unfolds in steps over time.

¹⁸One might object that cases of weakness of the will are non-psychological failures that warrant charges of irrationality. But weakness of will is a kind of psychological failure, since the will is a psychological capacity. See pages 97-101 for some discussion of the connection between weakness of will and reasoning.

only psychological capacities, and do not include non-psychological physical capacities.

One response to this line of thought is to point out that in cases of failure to act that are due purely to failures of the body, the agent was not *capable* of succeeding where she failed. When the grocery helper's elbow locked, it was not in her power to have it lock or not lock. One might suggest that the agent's failure in such cases is a failure in the exercise of her capacity to reason, but she is not rationally criticizable for her failure because her failure was out of her control. If you fail to form a belief because some psychological malfunction beyond your control makes you incapable of forming that belief—because the wire is cut, so to speak—you are not rationally criticizable for your failure. Your capacity to form beliefs is still a part of your capacity to reason. For the same reason, one might suggest, if you fail to carry out an action because of some physical malfunction beyond your control, like some cartilage lodging in your elbow joint, you are not rationally criticizable because the failure was out of your control. Your capacity to carry out physical actions, including your capacity to move your elbow, can still be part of your capacity to reason.

Now, this proposal concedes that we are only ever rationally criticizable in virtue of failures in the exercise of our psychological capacities. Thus, according to this proposal, non-psychological capacities to carry out physical actions are constitutive aspects of our capacity to reason, even though we are *always* passive with respect to failures of those capacities, and *never* directly culpable for them. It is natural to suppose that the possibility of culpable error is a mark of capacities for reasoning. Capacities for reasoning are traditionally thought of as active. We are, or at least can be, active with respect to the success or failure in the functioning of any of our rational capacities. Accordingly, it must be possible to be directly culpable for failures in the functioning of our rational capacities, and not always to be passive with respect to such failures. If we were never culpable for failures of the exercise of some capacity, that would suggest that we are not active with respect to the success or failure of that capacity. Thus, if we are never culpable for the success or failure in the functioning of

non-cognitive aspects of our capacities to carry out physical actions, then we are not active with respect to the functioning of those capacities, and they are not part of our capacities for reasoning.

Actions are criticizable as unintelligent, thoughtless, stupid, foolish, crazy, and the like. And they can be praised as intelligent, thoughtful, clever, cunning, and so on. These evaluative concepts are clearly keyed primarily to successes and failures of cognitive capacities. Intelligence consists primarily in the quality of one's cognitive capacities for thought, and exercises of those capacities. Physical actions are evaluable as intelligent derivatively, in virtue of their relation to exercises of those capacities. Derek Parfit suggests that, according to ordinary usage, 'irrational' means roughly the same thing as 'foolish,' 'stupid,' or 'crazy.' One respect in which Parfit's suggestion is correct is that, like these other evaluative concepts, irrationality consists primarily in cognitive failure, and applies to physical actions in a derivative way. Indeed, all rational criticism is warranted primarily by cognitive failure, and can only apply to physical actions in virtue of the actions' relations to cognitive failures. If one accepts that we are active with respect to our capacities to reason, and potentially culpable for failures in any aspect of the exercise of those capacities, then our capacities to reason do not include our physical abilities to execute physical actions.

Extended mind I want to end this section with a brief discussion of the bearing of views about "embodied cognition" on the issue here. An advocate of the claim that actions are conclusions of reasoning might hope for support from the family of views that fall under the "embodied cognition" heading.

The "embodied cognition" heading applies to a very diverse set of views. The main commonality between advocates of the various views is that they see themselves as rejecting Cartesian dualism about the mind and the physical world. Such a rejection might be thought to threaten the claim that physical actions are not psychological phenomena. However, most

views that fall under the “embodied cognition” heading are quite compatible with that claim.

One claim commonly associated with the heading is that the development of our rational capacities is continuous with, and influenced by, the nature of more basic cognitive capacities like vision and motor control. Another claim falling under the heading is that much cognition exploits brutally physical facts about our bodies, like the distance between a cognizer’s ears. A third claim that falls under the heading is that it will not work to try to build intelligent machines by coding rules for manipulating representations with propositional form.¹⁹ The first claim seems to me very likely to be true. The second claim is undoubtedly true. No cognitive scientist would deny it. The third claim is more controversial. But none of these claims have much bearing on the issues at hand. Each of these three claims is obviously compatible with the claim that there is more to moving your body than thinking, or any other type of psychological activity.

The only claim under the “embodied cognition” heading that might reasonably seem to threaten the argument I have been advancing is Andy Clark and David Chalmers’s “extended mind” view. Clark and Chalmers (op. cit.) claim that bodily actions, like drawing diagrams, can be part of one’s thinking, or reasoning. They even claim that objects in the external world, like the diagram itself, or a diary, can be part of one’s thinking or reasoning. Advocates of such a view might reject the claim that physical actions cannot be conclusions of reasoning because they are not psychological phenomena.

Clark and Chalmers’s arguments are not at all convincing, in my view. Much of the argument consists in accusing imagined objectors of question begging. That is not a winning strategy when arguing for a view as counterintuitive as theirs. The main positive argument, as I understand it, is based on a very crude functionalist view, according to which anything that regularly plays a useful role in enabling an individual to carry out a cognitive task is an element in her cognition. A diagram, a reminder in a diary, and a friend, can play regular

¹⁹See, for instance, Lakoff and Johnson (1999), and Anderson (2003).

useful roles in enabling an individual to accomplish a cognitive task of figuring something out. This observation is supposed to be evidence that diagrams, reminders, and even one's friends, can be elements of one's cognition.

I will not criticize this argument in any detail here. Instead, I want to point out that Clark and Chalmers's view actually does not favor the claim that actions are conclusions of practical reasoning. The functionalist motivation for their view coheres with the claim about the primary problem with the view that actions are conclusions of practical reasoning. Clark and Chalmers conceive of thinking or reasoning as a kind of figuring something out. Their idea is, roughly, that physical objects and actions play a constitutive role in figuring things out. Since reasoning is a kind of figuring something out, the conclusion of a piece of reasoning must be something like a commitment, whereby one settles on an answer to the problem being figured out. They think that a physical action of drawing a diagram can be part of figuring something out, since it is a useful part of the process that leads to one's answer to the question. But the step in one's reasoning where one arrives at a commitment concerning the matter at hand cannot consist in a physical action like drawing, or a physical action like going for a walk. Such actions are not commitments, whereby one figures out for oneself what to do. So they should not count as conclusions of practical reasoning, even according to Clark and Chalmers's view.²⁰

Conclusions of practical reasoning are intentions

The natural alternative to the view that conclusions of practical reasoning are actions is that conclusions of practical reasoning are intentions. The view that conclusions of practical reasoning are intentions is widespread. This view is not threatened by the considerations

²⁰Though they do not address the issue themselves, their view does allow that certain special physical actions, like writing or speaking the answer to a question could be a conclusion of reasoning. Unlike going for a walk, those physical actions can constitute commitments. I set aside those special cases here.

presented above, which threaten the claim that the conclusions of practical reasoning are the actions the reasoning is about. Forming an intention to ϕ is a psychological phenomenon. The fact that one intends to ϕ is a fact about one's psychology. So intentions are not precluded from being conclusions of practical reasoning in the same way as physical actions. And, crucially, an intention to ϕ is a commitment to ϕ ing, whereby the individual has "figured out for herself" what she shall do. In this section I support the view that intentions are conclusions of practical reasoning by defending the view from two sets of objections. The first set of objections is due to Sergio Tenenbaum, and the second is due to Joseph Raz.

Tenenbaum's argument that intentions are not conclusions of practical reasoning

In order to support his view that actions are conclusions of practical reasoning, Tenenbaum argues that intentions cannot be conclusions of practical reasoning (op. cit., p. 17). His argument has three main parts.

First, he claims that intentions based on practical reasoning do not determine the manner in which an action is carried out. In particular, an intention does not determine whether the corresponding action will be carried out well or poorly, or in an appropriate or inappropriate way. For example, he suggests that if you intend to walk to the light switch and turn it on, that intention does not determine whether or not you shove your computer out of the way so as to take the quickest route to the switch. Whether an intended action is executed well or poorly, or appropriately or inappropriately, depends on details of how the action is carried out that are not represented in the content of the intention.

Second, Tenenbaum claims that whether one's action is justified depends on the manner in which it is carried out. Suppose one intends to walk to the light switch and illuminate the room on the basis of reasonable considerations, like a belief that doing so will enable one to

read. If the agent carries out the action in an inappropriate way, say in a way that involves shoving her computer out of the way to make the path to the switch shorter, than she lacks justification for acting as she does.

In the third part of the argument, Tenenbaum claims that whether an action is justified is determined by success or failure in one's practical reasoning. But, according to the second claim, the way the action is actually carried out contributes to the justification of the action. As a result, carrying out the action in that way must be part of the agent's practical reasoning. If intention is the conclusion of one's practical reasoning, then the action itself is not a part of the practical reasoning. Acting is posterior to intention. Thus it is action, and not intention, that is the conclusion of practical reasoning.

This argument makes at least two mistakes. First, in the light switch case he describes, the agent's lack of justification for her way of acting is attributable to a failure of practical reasoning antecedent to her action. She intended to shove her computer out of the way. Forming that intention on the basis of the reasons she had for doing so was a piece of bad reasoning. Her action of shoving the computer is unjustified because of that failure in forming her intention on the basis of reasons.

Now, there is a difficult issue about whether and why the action of walking to the light switch and turning the light on is justified or unjustified in a case like the one described. One might be tempted to say that the agent's action of shoving the computer was unjustified, but she was justified in her walking to the light switch. There was nothing wrong with her walking to the light. Her mistake was in shoving the computer. This suggestion is problematic if one accepts that shoving the computer is part of her walking to the light switch. (One might deny that, but a better case can easily be found. Perhaps she stomped her way to the light switch, waking the baby in the next room.) On the other hand, one might defend the view that an action can be justified even if an aspect of that action is unjustified. But these niceties do not matter for present purposes. The aspect of the action

that is unjustified is unjustified because of a mistake in reasoning upstream of intention. If the action as a whole is also unjustified, it is because of that same mistake, and not because of the way the action was actually carried out.

Let me turn to the second mistake in Tenenbaum's argument. In some cases, an action can be justified even if the agent executes the action poorly, or inappropriately. Consider an individual who reaches for her glass because she is thirsty, but knocks the glass over, because her reaching action is carried out in a clumsy way. If she was aware that she would most likely knock over the glass, her reaching action might lack justification. But if she had no way of knowing that she would execute her reaching action in a clumsy way, then she is justified in reaching for the glass. It would be bizarre to accuse her of lacking justifying rational grounds for her action of reaching for the glass. She can be criticized for clumsiness, but not for acting without justification.²¹

It is worth emphasizing that in a case like this the agent makes no mistake in forming her intention to act. She only errs in the way she executes her action. Reflection on cases like these suggests that an agent is unjustified in carrying out a poorly or inappropriately executed action only if her mistake in action is due to some mistake in forming an intention.

Tenenbaum considers a response to his argument along these lines. He counters that, if any lack of justification is due to some failure in forming intentions on the basis of reasons, then intentions have to be implausibly detailed.

I must represent the layout of the room, and my path towards the room, make sure that I keep in mind all possible obstacles, represent how I will move my arms and legs so as to avoid the possible obstacles, think about what can happen in my room that can make turning the light on in a certain way problematic, represent more precisely how I am going to turn on the light, etc. It is quite implausible

²¹I discuss cases like these in more detail in "Non-rational aspects of agency."

that this is all even implicitly represented in forming the intention to turn the light on.

This rejoinder is inadequate. It is extremely plausible that agents take account of many of the sorts of details Tenenbaum mentions here. If the agent's successfully navigating the room is due to her agency, and is not a lucky accident, then it must be a result of her taking account of its layout and responding intelligently to her recognition of its layout. Indeed, if an agent happens to carry out her action in an appropriate way, but was in no way cognitively tracking the fact that that particular way of carrying out the action is the appropriate way, then her way of carrying out her action would lack justification, in much the same way that a true belief based on no evidence would lack justification.

The cognition underlying most cases of successful action is immensely complicated. Much of that cognition occurs at the level of thought and intention. That fact is reflected in the large number of Anscombe-style reasons-seeking questions that have application to different aspects of an action: Why did you go around the coffee table rather than over it? Why did you grasp the door handle? Why did you push the door rather than pull? One can supply one's reasons for many such aspects of carrying out a relative simple action like walking to the next room. In providing one's answers, one is not offering after-the-fact rationalizations. The answers genuinely report the reasons for which the agent acted.

It is true that we do not represent all of these details in conscious deliberation. One might be tempted to think that our intentions cannot cover so many details of action because of an identification of reasoning with conscious deliberation. But that identification would be a mistake. One can intend to perform many aspects of an action on the basis of reasons without consciously considering the reasons, or the intentions which those aspects of one's actions execute. Reasoning is, roughly, the exercise of one's ability to hold commitments on the basis of reasons. Most such reasoning does not consist in conscious deliberation. A lack of occurrent conscious awareness of one's intentions and actions does not cast doubt on the

claim that one's intentions, and their grounding by reasons, are psychologically real.

It also may help to acknowledge that, normally, some considerations about exactly how to carry out an action are not taken into account antecedently to one's beginning to act. Most actions execute a complex structure of intentions that change as the action unfolds through time. Some details about how to carry out the action will only be addressed by reason only when the agent arrives at the relevant stage of the action.²² Though typically, even in such cases, an action will be based partly on reasons that warrant an expectation that the action can be successfully executed. That point does not cast doubt on the claim that the justification of any aspect of an action consists in the way the agent's intentions are based on reasons.

Raz's argument that intentions are not conclusions of practical reasoning

I turn now to considering Raz's (2010) case against the view that intentions are conclusions of practical reasoning. The ultimate goal of Raz's arguments is different from Tenenbaum's. Tenenbaum's arguments are meant to support the claim that actions are conclusions of practical reasoning. Raz's arguments are meant to support his view that conclusions of practical reasoning are judgments, and that practical reasoning is theoretical reasoning that has practical subject matter.²³

When is a piece of reasoning practical? One objection Raz makes to the claim that intentions are conclusions of practical reasoning is that, if an intention is the conclusion of practical reasoning, then two instances of the same pattern of reasoning could differ as to

²²Other details are not addressed by reason at all, but are addressed by non-rational action guiding capacities.

²³Audi (1989) also claims that conclusions of practical reasoning are judgments about what one ought to do.

whether they are elements of practical reasoning. Consider an instance of reasoning that issues in a belief that one ought to ϕ , where the agent subsequently transitions from that belief to an intention to ϕ . Now consider another instance of that same pattern of reasoning, issuing in that same belief, but where there is no subsequent transition to intention. Raz suggests that, if conclusions of practical reasoning are intentions, then the first instance of reasoning (not including the transition to the intention) is an element of practical reasoning, and the second instance of reasoning is not. Raz supposes that the second instance of reasoning is not an element of practical reasoning on the grounds that the reasoning does not issue in the kind of conclusion that is characteristic of practical reasoning (namely intention). The two instances of reasoning differ in whether they are practical, even though each is an instance of the same pattern of reasoning, and involves the same mental states. Raz thinks that result is unacceptable, and so conclusions of practical reasoning are not intentions.²⁴

Raz's argument overlooks the fact that a piece of reasoning can count as practical reasoning even if it does not actually issue in the conclusion characteristic of practical reasoning. Both theoretical reasoning and epistemic reasoning can be inconclusive. An instance of reasoning is an instance of practical reasoning only if the reasoning *functions* to issue in its characteristic conclusion—intention, according to the view under discussion. Being an instance of practical reasoning does not require that the reasoning actually issue in an intention. Raz himself notes that reasoning can be “undertaken in order to examine the case for a modification of one’s beliefs or intentions” (op. cit., p. 4) even if the reasoning does not actually issue in a modification of one’s beliefs or intentions. So both pieces of reasoning in the above argument might be pieces of practical reasoning, even if only one of them actually

²⁴Streumer (2010) uses a similar argument to oppose the claim that conclusions of practical reasoning are actions. He suggests that, if actions are conclusions of practical reasoning, then, until the reasoner acts, it will be indeterminate whether a piece of reasoning constitutes practical reasoning. As a result, if one plans a trip a year in advance, the reasoning underlying the plan will not count as practical for the year before the trip is taken. Streumer suggests that that claim is an unacceptable consequence of the view that actions are conclusions of practical reasoning.

issues in intention.

Indeed, it is very plausible that reasoning that issues in a judgement that one ought to ϕ functions to issue in an intention to ϕ . If one forms that ought-judgment, but fails to form the corresponding intention, that is a kind of rational failure. That failure indicates that the reasoning is supposed to issue in intention, according to its function. Thus, it is plausible that any reasoning that issues in a judgment that one ought to ϕ constitutes practical reasoning.

Another response to Raz's argument is to point out the plausibility of the result that he thinks is unacceptable. Raz thinks two instances of the same pattern of reasoning cannot differ in whether they are elements of practical reasoning. But it is very plausible that they can. Theoretical reasoning about what is the case is often an element in practical reasoning about what to do. An instance of theoretical reasoning about some matter of fact might, in one case, be an element in practical reasoning about what to do, while in another case that instance of the same pattern of theoretical reasoning is purely theoretical, and not an element of reasoning about what to do. For example, reckoning a sum might be an element in practical reasoning about how much to tip a waiter in one case, while in another case that same piece of reckoning is carried out as a purely academic matter, and not as an element in practical reasoning about what to do.^{25,26}

²⁵I discuss this point further in "Epistemic aspects of practical reason."

²⁶Raz's own view about conclusions of practical reasoning provides no way to avoid this result. His view is that practical reasoning is reasoning with distinctively practical subject matter. The conclusion of practical reasoning is a belief about what one must do, or what one has reason to do. Much reasoning antecedent to that conclusion would seem to be practical only in certain contexts. One's judgment that one has reason to leave such-and-such a tip might be based partly on reckoning a sum. In that case, the reckoning is an element of practical reasoning. But in other cases the reckoning need not be an element in practical reasoning. The only way I can see for Raz to avoid this result is to insist that practical reasoning consists exclusively in the last step in one's reasoning about what one must do. This view implausibly reduces practical reasoning to a single transition in thought. Moreover, it implies that much reasoning that non-derivatively affects the rationality of actions and intentions is not in fact practical reasoning. If one's action of leaving a tip of a certain amount is based on an obvious mistake in arithmetic, the action is rationally criticizable. If successes and failures in reasoning that affect the rationality of actions are successes and failures in practical reasoning, then much reasoning antecedent to the final step counts as practical reasoning.

Weakness of will Let me turn to a second argument of Raz’s, which brings up more interesting issues about the nature of reasoning. Raz claims that weak willed failure to act shows that intentions are not conclusions of practical reasoning, but are posterior to one’s reasoning. He claims that the locus of the failure in at least some cases of weak willed failure to act is a failure to form an intention to act. He insists, nonetheless, that a failure to act due to weakness of will is never due to a failure to reason successfully. As a result, he concludes that a failure to form an intention is not a failure to reason successfully. Thus, forming an intention is not a conclusion of practical reasoning.

Raz does not defend his claim that weakness of will is never due to a failure to reason successfully. One potential source of motivation for that claim is the view that the locus of the mistake is the *body’s* failure to execute an action that one recognizes one should carry out. One might interpret Matthew 26:41 as offering a diagnosis along these lines of falling into temptation—“the spirit is willing, but the flesh is weak.”^{27,28} That potential motivation is undercut by Raz’s claim that the locus of the failure in the relevant cases of weakness of will is a failure to form an intention. Forming an intention is a psychological phenomenon, and not the same thing as executing a physical action. So a failure to form an intention does not consist in a bodily failure to execute an action. An advocate of the claim that intentions are conclusions of reasoning can hold that, when the locus of weak-willed failure to act is a failure to form an intention, weakness of will is due to a failure to reason successfully.

There is a plausible case to be made that the locus of the failure in the relevant cases is indeed a failure to reason successfully. An intention is a commitment to carry out an action. It is a committal mental state which can be based on justifying reasons. This fact suggests that intentions can be part of one’s reasoning. Plausibly, an exercise of a rational

²⁷Though the standard interpretation of the passage is that one’s desires interfere with the dictates of the “spirit.” Thanks to David Blank for pointing this out.

²⁸I doubt that non-psychological failure of the body is ever the locus of weakness of the will. The failure is always due to a defect in the will, which is a psychological capacity.

capacity to hold a commitment on the basis of reasons is an instance of reasoning.²⁹ Forming the commitment is an integral element in the exercise of such a capacity. Thus, forming an intention can be an element in reasoning. In the cases of weak-willed failure to act that Raz focuses on, the agent fails to carry out this element in reasoning.³⁰

Suppose a weak-willed agent fails to form an intention to ϕ , despite recognizing considerations indicating that she ought to ϕ . In the counterfactual situation, where the agent is not weak-willed, she forms the intention on the basis of reasons that justify (or purport to justify) the intention. According to the line of thought sketched above, the exercise of rational capacities whereby she holds the intention on the basis of reasons is an instance of reasoning. In the weak-willed case, the agent fails to carry out that reasoning, since she fails to form the intention to ϕ . In that way, the failure in the weak-willed case is a failure to carry out that element of reasoning. If the agent does not fail in any way to reason as she ought to, she will not fail to form the intention.³¹

Raz commits himself to a view along these lines in claiming that a failure to form an intention “can render the whole activity of reasoning on a particular occasion irrational.” This claim of Raz’s strongly suggests that forming an intention is not posterior to one’s reasoning, but is a constitutive aspect of one’s reasoning.³² One’s reasoning can be irrational only if

²⁹One might reserve the word ‘reasoning’ for exercises of capacities to hold commitments on the basis of reasons where the exercise unfolds in steps through time. (Compare footnote 17 above.) The points in this and the following paragraph apply *mutatis mutandis* to that conception of reasoning.

³⁰Admittedly, there is something odd about saying that an individual *reasons* from a judgment that one ought to ϕ to an intention to ϕ . In my view the transition from such a judgment to an intention is a step in practical reasoning. But the claim that intentions are conclusions of practical reasoning can be defended without relying on this point. I discuss these issues in the next subsection, titled ‘The transition to intention.’

³¹I am not claiming here that all cases of weakness of will are due in this way to failures to reason successfully. I am only claiming that this is so in cases of weakness of will associated with a failure to form an intention. However, it does seem to me plausible that in all cases of weakness of will there is some commitment to act that the agent either failed to hold, or failed to hold with sufficient strength, despite having conclusive reasons for the commitment, and for holding it with a certain degree of strength. I will not pursue that proposal here.

³²Roughly, a “constitutive aspect” of reasoning is an exercise of an ability in virtue of which one counts as holding a commitment on the basis of reasons. Constitutive aspects of reasoning differ from “elements” of

there is something wrong with the reasoning.³³ If forming an intention is not a constitutive aspect of one's reasoning, then that failure to form an intention does not constitute a failure of the agent to reason well. As a result, the failure would not render the agent's reasoning irrational. Since the failure does render the reasoning irrational, forming an intention is a constitutive aspect of one's reasoning. Such a failure to form an intention constitutes a failure to reason successfully.

Raz claims that the capacity to form intentions is a rational capacity, but not a constitutive aspect of our capacity for reasoning (ibid., p. 3). His claim that the capacity is a rational capacity is motivated by his acknowledgment that a failure to form an intention can warrant charges of irrationality. But he wants to insist that intentions are not elements in reasoning. So he claims that the capacity to form intentions is a rational capacity, but not a constitutive aspect of our capacity for reasoning. But this claim is incompatible with Raz's claim that failures to form an intention can render one's *reasoning* irrational. Raz's position would be strengthened if he revised it so that a failure to form an intention would make the agent criticizable as irrational, without making her reasoning itself count as irrational.

In my view, an agent is not rationally criticizable if she reasons as she ought to, and she does not violate any rational norms by forming ungrounded, or inadequately grounded, committal attitudes through the exercise of her rational capacities.³⁴ It is tempting to

reasoning. Elements in reasoning include, again roughly, reasons that are operative in reasoning, conclusions of reasoning, and transitions in reasoning from reasons to conclusions, like making use of modus ponens. Not all constitutive aspects of reasoning are elements of reasoning. For example, preserving a sub-conclusion in memory in the course of reasoning is a constitutive aspect of reasoning, but not an element of reasoning. (See the discussion in the next subsection, titled "The transition to intention.")

³³To clarify this claim, I want to distinguish two respects in which one's reasoning can be irrational. It can be irrational in that it is a mistake to engage in the activity of reasoning, for example if there is something else one is supposed to do. Alternatively, reasoning can be irrational in that there is something wrong how one reasons. The second sort of irrationality is what is at issue here.

³⁴In some cases, an agent can be rationally criticizable even if she does not engage in any defective reasoning. For example, an individual might believe that she will win the lottery because of wishful thinking, and not on the basis of any reasons. Such a belief is irrational, even though the individual does not engage in any defective reasoning. The agent is criticizable as epistemically irrational because she violates epistemic rational norms in forming her ungrounded belief. This kind of mistake is in certain respects similar to the

speculate that Raz is pulled to say that a failure to form an intention makes one's reasoning irrational because he is drawn to a view along these lines. If one accepts that view, then it is difficult to deny that forming an intention is a constitutive aspect of one's practical reasoning, instead of being posterior to the reasoning. If so, then failures to form intentions in cases of weakness of will are failures to reason successfully.

But the case so far for claiming that intentions are constitutive aspects of reasoning does not depend on the claim that rational criticism is always due to a failure to reason successfully. The main point is that intentions are commitments that are based on reasons. To buttress the case against Raz's view that intentions are not conclusions of practical reasoning, I want to discuss a claim he makes about transitions from a belief that one has conclusive reason to ϕ to an intention to ϕ . He claims that such a transition is not a transition in reasoning (*ibid.*, p. 14). That claim might initially seem plausible. If such a transition is not a transition in reasoning, that might seem to suggest that intentions are not conclusions of reasoning.³⁵ I suspect that the former claim is a large part of what motivates his view that intentions are posterior to reasoning. In any case, addressing this claim helps to support the main point at issue here—that forming an intention is a constitutive aspect of one's reasoning, and so has a strong claim to be the conclusion of practical reasoning. I want to argue that such a transition is a transition in reasoning.

mistake of carrying out defective reasoning. In forming the ungrounded belief, the individual's rational capacities fail with respect to their epistemic function of guiding the individual to true belief. In carrying out defective epistemic reasoning, one's rational capacities fail with respect to that same function.

Similar points can be made concerning practical rationality. Suppose the agent intends to buy a boat on the basis of her irrational belief that she will win the lottery. Her intention is irrational, because of the irrationality of the belief it is based on. In forming the ungrounded belief, the individual's rational capacities fail with respect to their function of guiding the individual to good action, in addition to failing with respect to their function of guiding the individual to true belief. That is why the mistake of forming the belief can render the intention irrational. I discuss these issues further in "Practical warrant" and in "Epistemic aspects of practical reason."

³⁵When an intention is based on reasons, the transition to the intention need not involve a belief about what one must do, or what one has reason to do. Such transitions are rare, in my view. I discuss this point below.

The transition to intention It might initially seem plausible that a transition from a belief that one has reason to ϕ to an intention to ϕ is not a transition in reasoning. Raz seems to find the claim intuitively plausible, since he does not try to defend it. Let me try to characterize the appeal of the claim.

Raz’s remark seems motivated by the view that such a transition is not informative. He says, “We do not deliberate from P to P. Do we deliberate or reason from I must do A to intending to do it?” (ibid., p. 14).³⁶ It does seem that there is some respect in which such a transition to an intention is not informative. The transition might seem not to be an instance of figuring something out. One does not come to a new realization about oneself or the world in making such a transition. Instead, in making such a transition, the conclusion one reaches about what one has reason to do is given some kind of motivational force, or motivational upshot. One has understood all the relevant facts antecedently to forming the intention.³⁷ In making the transition, one is “moved” by what one has figured out, but one does not figure out anything new. One might assume that reasoning is essentially a kind of figuring something out. So the transition is not a constitutive aspect of reasoning.³⁸

I want to make a few points in response to this line of thought. First, one can form an intention to ϕ on the basis of reasons without doing so via a belief that one has reason to ϕ , or a belief that one must ϕ . Young children and probably some higher animals are capable of

³⁶Raz slides between talking about the belief that one has conclusive reason to ϕ and talking about the belief that one must ϕ . I focus on the first kind of belief in the discussion here.

³⁷In my view, one can understand all the relevant facts antecedently to forming an intention even in cases that involve no “meta” belief about what one ought to do. To see the appeal of this view, it helps to consider an analogous claim about belief and epistemic reasoning, made by Burge (2000). In non-critical, “object-level” epistemic reasoning one can distinguish an aspect of the reasoning that consists in one’s understanding of the implications of the basis beliefs, and one’s being moved by the force of those considerations to adopt the conclusion belief. One can possess an understanding of the implications of the basis beliefs, without being moved by them to endorse the conclusion they recommend. Possessing that understanding does not require that one hold a “meta” belief that one ought to hold the conclusion-belief. Similarly, one can possess an understanding of how certain considerations recommend that one adopt an intention, without having a “meta” belief that one ought to adopt that intention.

³⁸Let me be clear that I am interpolating here, in an effort to dig out the appeal of the claim that the transition to intention is not a transition in reasoning here. Raz does not make these points himself.

forming intentions on the basis of reasons, but lack the concept of a reason.³⁹ If an intention is formed on the basis of reasons without a meta-level belief about one's reasons for action, it may be easier to see that the transition to the intention is a transition in reasoning. If one intends to walk to the kitchen on the basis of a belief that there is food there, the transition from the belief to the intention clearly involves some kind of practical "putting two and two together." The transition involves some sort of substantive appreciation of the practical significance of the basis belief.

Second, not every transition in reasoning is an instance of arriving at a new realization in the way described two paragraphs above. It is true that a transition from a belief that one has reason to ϕ to an intention to ϕ is significantly different from a transition in reasoning like making an inference by applying modus ponens. But not every piece of reasoning is like that kind of inference. Consider a step in critical epistemic reasoning, in which one transitions to a belief that p from a belief that one has conclusive reason to believe that p . Plausibly, in making such a transition one is not making use of the basis belief as evidence for the truth of the conclusion belief. Instead, the transition gives the basis belief some kind of non-practical "motivational" force, or implementational upshot.^{40,41} In a certain respect, making such a transition does not involve figuring out anything new. One has recognized all the relevant facts and understood their implications antecedently to making the transition.

³⁹Raz suggests that no such intention could be warranted. But his argument is extremely unconvincing. I will touch on it only briefly. First, he points out that if one has reasons that would warrant an intention to ϕ , one also has reasons to believe that one has reason to ϕ . Even if that is true, it is irrelevant. Just because the reasons for the intention would warrant a meta-belief about reasons does not mean that the intention is based on such a belief. Second, he points out that the existence of reasons to ϕ , or an agent's having reasons to ϕ is not enough to warrant an agent's intention to ϕ . He supposes that if an agent were to intend to ϕ without believing that she has reason to ϕ , then she would merely have reasons to ϕ , but her intention would not be based on those reasons. That is an implausible view, and Raz offers no support for it. One can carry out an inference and hold a commitment on the basis of the premises without making use of meta-level beliefs about one's cognitive situation.

⁴⁰This point is due to Burge (2000).

⁴¹The "motivational" force is non-practical because it does not play a direct role in moving one to action or intention. It occurs in epistemic reasoning, and plays a direct role in moving one to form a belief. The belief could play an indirect role in moving one to act, if an action is eventually based on the belief.

In making the transition, one is “moved” by one’s grasp of the relevant consideration to make the commitment that they recommend. Even so, that transition to the belief that p is indisputably a constitutive aspect of epistemic reasoning. It is a transition in reasoning, and is not posterior to one’s epistemic reasoning. Thus, even if the transition to intention is best understood as a case of being “moved by” one’s belief, without figuring out anything new, the transition could be a transition in reasoning.

However, there is a case to be made that a transition from a belief about reasons for intention to an intention is partly constitutive of a certain kind of figuring out something new. To see this, consider a case of critical epistemic reasoning again. Suppose an individual engages in “meta” reasoning about whether she has conclusive reason to believe that her son is guilty of some ugly crime. Suppose she forms a “meta” belief that she has conclusive reason to believe that her son is guilty, but cannot bring herself to form the “object-level” belief that her son is guilty. It is not altogether clear whether we should say that the individual has not figured out whether her son is guilty. One might reasonably prefer to say that she has figured out that he is guilty, but cannot bring herself to believe it. But on the other hand, one might prefer to think of “figuring something out” as requiring that one actually makes up one’s mind. In reasoning to a “meta” belief that one has reason to believe that one’s son is guilty, one has not yet made up one’s mind about whether it is the case one’s son is guilty. Whether one has reason to believe that one’s son is guilty is a different matter from whether one’s son actually is guilty. In transitioning from the “meta” belief to the belief that one’s son is guilty, one makes up one’s mind about the latter matter in light of one’s belief about the former. If one thinks of figuring something out as requiring that one actually make up one’s mind about the relevant matter, then forming the “object-level” belief that one’s son is guilty is partly constitutive of figuring out whether one’s son is guilty.

Similar points apply to critical practical reasoning. There is a distinctively and irreducibly *practical* matter of whether actually to ϕ , which is different from the matter of whether one

has reason to ϕ .⁴² In reasoning to a belief that she ought to ϕ , an individual has not yet made up her mind about the practical matter of whether actually to ϕ . In transitioning from her “ought” belief to her intention, the individual makes up her mind about that practical matter, in light of her “meta” belief about reasons for action.

One might think of “figuring out whether to ϕ ” as involving only one’s understanding of whether to ϕ . According to that view, forming an intention to ϕ on the basis of a “meta” belief about reasons would be posterior to figuring out whether to ϕ . On the other hand, one might instead think of “figuring out whether to ϕ as requiring that one actually make up one’s mind about whether to ϕ . According to that view, forming an intention to ϕ on the basis of the meta belief would be partly constitutive of figuring out whether to ϕ . Thus, the view that forming an intention is partly constitutive of reasoning need not conflict with the view that reasoning is a kind of “figuring something out.”

This line of thought goes hand in hand with the suggestion I made earlier, that exercises of one’s rational capacity to hold committal attitudes on the basis of reasons constitute reasoning.⁴³ I suggested that forming an intention can be partly constitutive of reasoning because an intention can be a committal attitude that is based on reasons. An intention is a commitment to act. That is different from a commitment to its being the case that one has reason to act. In making that latter sort of commitment, one has not yet made a commitment to act. A transition from a belief that one has reason to ϕ to an intention to ϕ is a transition from one commitment to another, where every constitutive aspect of the transition is an exercise of the agent’s rational capacities.⁴⁴ The transition is a commitment-guiding exercise of one’s rational capacities. If one accepts that exercises of one’s rational capacity to hold committal attitudes on the basis of reasons constitute reasoning, then forming the intention

⁴²The matter of whether one has reason to ϕ is also practical, but in a different way. It is a matter of fact. In making up one’s mind about whether to ϕ , one does not make up one’s mind about a matter of fact.

⁴³See the third paragraph of the sub-section titled “Weakness of will.”

⁴⁴Raz accepts that the ability to form intentions is a rational capacity.

is partly constitutive of one's reasoning. Since the intention is the commitment guided by this kind of reasoning, it is plausible that the intention is the conclusion of the reasoning.

In my view, the transition in critical epistemic reasoning from a meta-belief about reasons for belief to an object-level belief is an element in reasoning. Roughly, elements in reasoning include reasons that are operative in reasoning, conclusions of reasoning, and transitions in reasoning from reasons to conclusions. The critical reasoner makes use of the meta-level belief that she has sufficient reason to believe that p as a reason for her belief that p . Her meta-level belief about reasons is part of her understanding of the case for her belief that p . I have the same view about the transition from a meta-belief about reasons for intention to an intention. An agent's belief that she has reason to ϕ is part of her understanding of the case for ϕ ing. That meta-belief operates as a reason for the agent's intention. So the transition is an element in reasoning.

The claim that the transition to intention is an element in reasoning finds support in the point made above, that the transition constitutes a kind of "figuring something out." In making the transition to the intention, the agent makes use of one commitment to make up her mind about a further commitment (namely whether to ϕ), through the exercise of rational capacities. Such a transition would seem to be an instance of transitioning from a reason to a conclusion based on that reason. Thus, the transition seems to be an element in reasoning.

However, even if that transition to intention is not an element in reasoning, that point would not help Raz's position very much. At most, it would occasion a clarification about the meaning of the claim that intentions are conclusions of reasoning.

To explain the point, I want to begin by considering conclusions of epistemic reasoning that are modified in memory, in order to calibrate indexical elements to changing contexts. For example, suppose that on Monday an individual engages in reasoning that concludes in a belief that there is full moon today. On Tuesday, she transitions from that belief to a belief

that there was a full moon yesterday. The transition from the today-belief to the yesterday-belief is not an element of reasoning. In transitioning to the yesterday-belief, the individual is preserving her earlier today-belief in memory, and updating its indexical elements. Since the yesterday-belief is, as it were, just a continuation of the today-belief, the today-belief does not operate as a reason for the yesterday-belief. Accordingly, the transition in memory from the today-belief to the yesterday-belief is not an element in reasoning.

The transition is, however, a constitutive aspect of an exercise of the individual's ability to hold a commitment on the basis of reasons. On Tuesday, she believes that there was a full moon on the basis of reasons. She does so partly in virtue of her transition from the today-belief to the yesterday-belief. If one thinks of reasoning as the exercise of an ability to hold commitments on the basis of reasons, then the transition to the yesterday-belief is a constitutive aspect of reasoning. The transition is a constitutive aspect of reasoning even though it is not an element of reasoning.

In a similar way, even if a transition to an intention is not an element of reasoning, it can still be a constitutive aspect of reasoning. The transition from a belief that one has reason to ϕ to an intention to ϕ is a constitutive aspect of the exercise of one's ability to intend to ϕ on the basis of reasons. If one thinks of reasoning as the exercise of an ability to hold commitments on the basis of reasons, then that transition to intention is a constitutive aspect of one's reasoning.

One might restrict the term 'reasoning' so that it includes nothing beyond the conclusion of the last element in the reasoning. According to this usage of the term, the transition to the yesterday-belief in the moon case described above is not a constitutive aspect of one's reasoning. The last element in the individual's reasoning is her transition to her belief that the moon is full today, from her reasons for that belief. The transition from the today-belief to the yesterday-belief is posterior to that last element of her reasoning. So the belief that the moon was full yesterday is not the conclusion her reasoning. If a transition from a belief

that one has reason to intend to *phi* to an intention to ϕ is not an element of reasoning, then, according to this narrower conception of reasoning, the transition is posterior to one's reasoning. If so, an intention formed through such a transition would not be a conclusion of reasoning.⁴⁵

Though it is not the terminology I prefer, that narrower usage of 'reasoning' seems to me a mostly unobjectionable terminological choice. But it does invite misunderstanding, at least somewhat. The restricted usage may lead one to think in a too-narrow way about our capacities to hold commitments on the basis of reasons. There is a widespread and plausible view that our capacity to hold commitments on the basis of reasons *is* our capacity to engage in reasoning. But, as explained above, there are constitutive aspects of our capacity to hold commitments on the basis of reasons whose exercise can be posterior to the last element in reasoning—like preservative memory. As long as one does not lose sight of that point, the restricted usage of 'reasoning' seems to me unobjectionable. Commitments that are formed or updated posterior to the last element of reasoning would not be conclusions of reasoning according to that usage. However, they would be very much like conclusions of reasoning, since they would still be commitments held on the basis of reasons.

In any case, as suggested above, it seems to me that a transition from a belief about reasons for intention to an intention is an element of reasoning. Such a transition is analogous to a transition in critical epistemic reasoning from a belief that one has reason to believe that *p* to a belief that *p*. One's meta-level recognition of the way the belief or intention is supported by reasons is part of one's understanding of the case for that belief or intention. If so, even the restricted usage of 'reasoning' would allow for intentions formed through such transitions to count as conclusions of reasoning.⁴⁶

⁴⁵However, these views do not preclude most reasons-based intentions from being conclusions of reasoning, since most such intentions are not based on meta-beliefs about one's reasons for intention.

⁴⁶Some intentions, however, would be precluded from counting as conclusions of reasoning by the restricted usage. The transition from a belief that the moon is full today to a belief that it was full yesterday has

Conclusion

I have argued in this chapter that it is intentions, and not actions, that are conclusions of practical reasoning. Physical actions cannot be conclusions of reasoning because they are not psychological phenomena. More generally, actions that one's reasoning is about cannot be conclusions of the reasoning, because the reasoning consists in figuring out whether to ϕ . ϕ ing is not part of figuring out whether to ϕ . Figuring out whether to ϕ terminates in some kind of committal attitude concerning whether to ϕ . Carrying out the action is not a commitment to carrying out the action, anymore than fulfilling a promise is the same as making the promise. The core point in favor of the view that intentions are conclusions of practical reasoning is that intentions are commitments to act that are based on reasons.

The view that practical reasoning concludes with intention leaves unaddressed the way that actions figure in practical reasoning as something more than subject matter. Actions are represented in intentions, so the view shows one way in which actions figure in practical reasoning as subject matter. But the view does not shed light on the way practical reasoning constitutively involves actions themselves, as something more than just the subject matter of the reasoning. It is clear that practical reasoning does involve action in some such way. Actions themselves, and not just intentions, are based on practical reasoning. And actions are rationally assessable in virtue of the quality of one's practical reasoning. In Chapter Four, "Reason's guidance of action," I develop a view about the way that practical reasoning bears on action. I develop and defend the claim that, even though actions are not conclusions of reasoning, they are guided by reasoning in largely the same way as conclusions of reasoning, like beliefs and intentions.

a close analog concerning intention. On Monday an individual might intend to call her brother tomorrow, and on Tuesday intend to call her brother today. The transition from the tomorrow-intention to the today-intention is not an element of reasoning. Since the transition is posterior to the last element of reasoning, the today-intention would be posterior to the agent's reasoning, and so would not be the conclusion of her reasoning.

Chapter 5

Reason's guidance of action

Introduction

Joseph Raz (2010) claims that practical reasoning is a special case of epistemic reasoning. According to Raz's view, action is related to practical reasoning in the same way that animals are related to zoology, or money is related to economics. Action is just the characteristic subject matter of practical reasoning. Raz's is a minority view. It is widely accepted that there is some constitutive connection between practical reasoning and action, which goes beyond action's role as the subject matter of practical reasoning. Accordingly, practical reason is often described as being practical in its issue, in addition to being practical in its content and its subject matter. Practical reason is practical in its issue because of the way that it issues in action. This chapter aims to make progress in understanding the particular way in which action figures as the issue of practical reasoning.

One might try to approach this issue by claiming that actions issue from practical reasoning as *conclusions* of practical reasoning. This claim is usually attributed to Aristotle, although that interpretation is contested.¹ If action is the conclusion of practical reasoning,

¹See Charles (1984) for discussion of the minority interpretation.

then action figures as something more than subject matter represented in the content of practical reasoning. Action issues from practical reasoning in the same way that belief issues from epistemic reasoning. They both issue from reasoning as the conclusion of the reasoning.

But, as argued in Chapter Three, the view that actions can be conclusions of practical reasoning conflicts with very natural assumptions about the nature of reasoning. Reasoning is a psychological phenomenon. It is a kind of thinking, or figuring something out. Epistemic reasoning is, roughly, an exercise of one's rational capacity to figure out what is the case. Practical reasoning is, again roughly, an exercise of one's rational capacity to figure out what to do.

A conclusion of an individual's reasoning—whether practical reasoning or epistemic reasoning—is a concluding element *in* that reasoning. More specifically, a conclusion of reasoning is a committal psychological state, whereby one purports to have “figured out” the issue that the reasoning is about. A conclusion of epistemic reasoning about whether it is the case that p is a committal psychological state whereby one purports to have “figured out” whether it is the case that p (namely a belief). A conclusion of practical reasoning about whether to ϕ is a committal psychological state whereby one purports to have figured out whether to ϕ (namely an intention, in my view).

A physical action cannot be a conclusion of practical reasoning because physical actions are not psychological phenomena. Physical actions may be individuated by elements of one's psychology, like intentions and one's reasons for action. Actions might even sometimes consist partly in psychological phenomena. Perhaps moving a couch into an apartment can consist partly in figuring out how to angle the couch to get it up the stairwell. But there is more to physical action than psychology. One moves the couch with one's arms and legs. One does not reason with one's arms and legs. Since conclusions of practical reasoning are elements of a psychological phenomenon of reasoning, a conclusion of practical reasoning cannot consist, even partly, in non-psychological phenomena, like movements of one's arms

and legs.

Mental actions, like imagining oneself on a beach, are psychological phenomena. But they are still not conclusions of practical reasoning. The practical reasoning associated with such a mental action consists in one's figuring out whether to perform that action, like whether to imagine oneself on a beach. The conclusion of the reasoning is a committal mental state, whereby one purports to have figured out whether to do so. Imagining oneself on a beach is not a committal mental state of this sort, and so is not a conclusion of reasoning.

In short, actions are not conclusions of practical reasoning because it is one thing to have figured out whether to ϕ , and another thing to actually be ϕ ing. In my view, conclusions of practical reasoning are intentions. An intention to ϕ is a committal mental state, whereby one purports to have figured out whether to ϕ .²

However, the view that practical reasoning concludes with intention leaves unaddressed the way that actions figure in practical reasoning as something more than subject matter. Actions are represented in intentions, so the view shows one way in which actions figure in practical reasoning as subject matter. But the view does not shed light on the way practical reasoning constitutively involves actions themselves, as something more than just the subject matter of the reasoning.

It is clear that practical reasoning does involve action in some such way. Actions themselves, and not just intentions, can be based on practical reasoning. And actions can be rationally assessable in virtue of the quality of one's practical reasoning. If the practical reasoning associated with an action is poor or excellent, it is not just aspects of the agent's reasoning that are rationally criticizable or laudable. Mistakes in practical reasoning can render actions themselves irrational. Likewise, successes in practical reasoning can render actions rational.

²I develop these points about conclusions of practical reasoning in more detail in "Conclusions of practical reasoning."

It is worth emphasizing that the primary locus of rationality and irrationality of actions is not the actions themselves. The rationality and irrationality of actions is derivative. Actions are rational in virtue of successes in practical reasoning associated with the action. Actions are irrational in virtue of failures in practical reasoning associated with the action, or in virtue of one's failure to engage in practical reasoning that one ought to have engaged in. An action is not a part of one's reasoning. As a result, the success or failure that makes an action count as rational or irrational cannot consist in one's carrying out that action. Rational action does not itself constitute a success in the exercise of one's capacities to reason. Nor does irrational action constitute a failure in the exercise of one's capacities to reason. Actions *derive* their rationality and irrationality from successes and failures in the exercise of one's capacities to reason. Although the rationality and irrationality of actions is derivative, it is genuine. Actions really do satisfy and violate rational norms, even if they do so in virtue of successes and failures of reasoning.

In contrast to actions, conclusions of reasoning, like beliefs and intentions, are aspects of one's reasoning. As a result, their rationality and irrationality need not be entirely derivative. If an agent forms a belief without having any grounds for it whatsoever—as in a case of wishful thinking—the formation of the belief constitutes a failure in the exercise of her capacities for reasoning. The belief is irrational in virtue of that failure. Her belief does not derive its irrationality from some other failure in the exercise of her capacities to reason.

However, in many cases, the rationality or irrationality of conclusions of reasoning is derivative. Consider a belief formed through a rational transition from premises that an individual endorses irrationally. The conclusion-belief is irrational, even though the individual makes no mistake in the specific step in her reasoning where she forms the belief. Forming the belief does not itself constitute a failure in the exercise of her rational capacities. The primary locus of the rational failure is elsewhere in her reasoning—perhaps in the formation of the premise-beliefs—and does not consist in the formation of the conclusion-belief itself.

The conclusion-belief derives its irrationality from some failure in the exercise of her capacities to reason antecedent to the formation of the belief. Like actions, such conclusions of reasoning are derivatively rationally assessable.

As a first pass, it is initially plausible that actions, beliefs and intentions are rationally assessable in virtue of reasoning because we act, and form intentions and beliefs, under the guidance of reasoning. We determine what we do, what we intend, and what we believe, through reasoning that enables us to be right in our actions, intentions and beliefs. Actions, intentions, and beliefs are objects of guidance of reasoning. They are guided by reasoning in a way that makes them rationally assessable in virtue of the reasoning. The rationality of actions, and of conclusions of reasoning, is roughly a matter of the quality of the reasoning that guides them.³

The relation between reasoning and action is in certain respects very different from the relation between reasoning and its conclusions. A conclusion of reasoning is an element in one's reasoning, and an action is not an element in one's reasoning. The fact that actions are not part of one's reasoning might seem to suggest that there is no common account of why actions, intentions, and beliefs, all count as guided by reasoning, and are all rationally assessable in virtue of the quality of one's reasoning. The relation between action and reasoning might seem to be *sui generis*, and quite unlike the relation between reasoning and

³The claim that the rationality of a belief or intention is a matter of the guidance of the belief or intention might seem at odds with the point that the rationality or irrationality of beliefs and intentions need not be entirely derivative. One might doubt that a belief or an intention guides *itself*. The locus of irrationality of a belief or intention can be a mistake in forming the belief or intention. As a result, it may seem that the rationality of a belief or intention is not just a matter of the guidance of the belief or intention.

I want to insist that the formation of a belief or an intention is an element in the guidance of the belief or intention. Reason's guidance of a conclusion of reasoning, like a belief or intention, consists in the entirety of the reasoning that issues in the conclusion. That includes the formation of the belief or intention itself. Every integral aspect of the reasoning is an integral aspect of the guidance of the conclusion. As a result, any integral aspect of the reasoning can contribute constitutively to the quality of that guidance, and can thus contribute constitutively to the conclusion's rationality or irrationality. That includes the formation of the belief or intention itself, since the formation of the belief or intention is an integral aspect of the reasoning. The guidance also includes aspects of the reasoning antecedent to the conclusion, which also can contribute constitutively to the rationality or irrationality of the conclusion.

conclusions of reasoning, like beliefs and intentions.

This chapter presents a partial account of what it is to be an object of guidance of reasoning, which applies both to actions and to conclusions of reasoning, like beliefs and intentions. I aim to identify some core constitutive aspects of the relation between action and reasoning which make actions objects of guidance of reasoning. According to the account I propose, reasoning is related in this same way to conclusions of reasoning, like beliefs and intentions. Actions are guided by reasoning in much the same way as intentions and beliefs, even though actions cannot be conclusions of reasoning. In developing the account, I use the rational assessability of actions, intentions, and beliefs as a touchstone. An account of the way that actions are guided by practical reasoning should help to explain why actions are rationally assessable in virtue of the quality of one's practical reasoning—just as beliefs and intentions are rationally assessable in virtue of the quality of one's epistemic reasoning or one's practical reasoning.

Causation

As noted above, it is a commonplace that practical reasoning is practical not just in its content or subject matter, but also in its issue. The practical issue of practical reasoning is supposed to be action. A majority of the recent discussion about the way that practical reasoning issues in action concerns debates about the “causal theory of action.” According to that theory, action is caused in a non-deviant way by psychological states that constitute reasons for action (like beliefs and desires), or by intentions, where those intentions can be based on reasons for action.⁴ A variant of the causal theory of action can help as a first step in understanding the way that action is guided by practical reasoning. Actions count as guided by practical reasoning partly because practical reasoning is a cause of actions.

⁴The *locus classicus* for this view is Davidson (1980).

Actions are caused by practical reasoning, via intention. The fact that an individual intended to ϕ can rightly be cited in a psychological explanation of why she ϕ -ed. When one acts intentionally, one acts because of one's intention to act, and not independently of one's intention. If an intention is based on practical reasoning, that reasoning can rightly be cited in a psychological explanation of why that agent intended to ϕ . If an intention is based on reasoning, her intention is due to her reasoning, and is not formed independently of the reasoning. Psychological explanations of why an individual acted as she did, or intended as she did, are causal explanations. Thus, practical reasoning can be a cause of an action, through its causing an intention.

Epistemic reasoning can be a cause of belief, just as practical reasoning can be a cause of intention and action. When a belief is based on epistemic reasoning, the reasoning can rightly be cited in a causal-psychological explanation of why the individual holds that belief. Thus, being caused by reasoning is a relation to reasoning that beliefs, intentions, and actions have in common.

The quality of one's reasoning bears on the rationality of a belief, intention, or action partly in virtue of the fact that the belief, intention, or action is caused by one's reasoning. If a piece of reasoning played no role in causally determining whether an individual holds a given belief or intention, or performs a given action, then the quality of the reasoning would not bear on the rationality of the belief, intention or action.

In order to clarify what I have in mind here, I want to consider an objection to the claim that reasoning guides belief, intention, and action partly because they are caused by reasoning. One might argue that the reasoning one describes in answering a reasons-seeking why-question can affect the rationality of one's intention (or action, or belief), even if one never went through the reasoning. It is plausible that an individual can form an intention on the basis of reasons, even if she never went through a temporally extended process of reasoning that unfolds in steps. In a case like that, the individual's intention is rational or

irrational in virtue of the way it is based on reasons, even though she never went through a corresponding step-by-step process of reasoning, and no process of reasoning caused her intention.

But even in a case like that, the individual does make use of the reasons in forming her intention. Her reasons are brought to bear in forming the intention through the exercises of psychological capacities to make use of reasons—albeit not through a temporally extended process of step-by-step reasoning. The way the individual makes use of reasons can rightly be cited in a psychological explanation of why she intends to ϕ . An individual might intend to turn left *because* she believes that is an efficient way to get to the grocery store, even if she never went through a step-by-step, temporally extended process of reasoning linking the turn to the grocery store. Such a psychological explanation of why an individual intends to ϕ is a kind of causal explanation of the intention. Thus, the way the individual makes use of reasons is a cause of her belief, even if she never went through a step-by-step, temporally extended process of reasoning.⁵

When an individual truthfully answers a reasons-seeking why question, we can think of the answer as corresponding to a way of making use of reasons, where that may or may not consist in a temporally extended process of step-by-step reasoning. However, the way one makes use of reasons is a psychological phenomenon that explains why the individual has the intention she has. The way of making use of reasons that corresponds to her answer affects the rationality of her intention partly because that way of making use of reasons psychologically explains why she intends to ϕ . If she had not made use of reasons in that way, or if she intended to ϕ completely independently of making use of reasons in that way, then that way of making use of reasons would not affect the rationality of her intention. Her answer would be an after-the-fact rationalization of her intention, and not a description of

⁵I discuss cases like this in more detail in “Reasons and deliberation in Aristotle: comments on Perin.” Also see Audi (2004).

her actual rational grounds for her intention.

I want, then, to make the following clarification about the claim that reasoning guides intention, action, and belief partly because they are caused by reasoning. The claim is that exercises of one's psychological capacity to make use of reasons guides intention, action, and belief, and affects their rationality, partly in virtue of the fact that they are caused by exercises of that ability. In effect, the clarification is that by 'reasoning' I have in mind exercises of one's ability to make use of reasons, where making use of reasons may or may not consist in a temporally extended, step-by-step process. It is reasoning in this sense that must figure in a causal-psychological explanation of actions, intentions, and beliefs, if it is to affect the rationality of the action, intention, or belief ⁶

Consider an individual playing a game of mastermind, who believes that the third peg is red because she makes use of reasons available to her in a clever way. The quality of the way she makes use of reasons affects the rationality of the belief partly because the belief is due to the way she makes use of reasons. If the individual never actually made use of reasons in that way, then her belief could not be rational in virtue of that way of making use of reasons. If she defended the rationality of her belief by describing a clever pattern of reasoning, but had not actually made use of reasons in way that corresponds to the reasoning she describes, then her defense would be an after-the-fact rationalization. Her defense could not legitimately establish the rationality of her belief.

Or suppose that the individual had made use of reasons in a way corresponding to the description, but then was distracted by something, and completely forgot her conclusion

⁶This terminology is somewhat misleading. The gerundive form of the term 'reasoning' does suggest that reasoning must be a temporally extended process. If one insists on using the term only for such processes, then I would respond to the objection by making a different clarification about the claim. The claim is that, if one's intention, action, or belief is guided by a process of reasoning, it counts as guided by that reasoning partly because it is caused by the reasoning. In cases where an intention, action, or belief is based on reasons, but not on a process of reasoning, those cases are outside the scope of the claim. The claim I am more interested in, however, is the broader one, about exercises of one's ability to make use of reasons. The discussion in this chapter applies *mutatis mutandis* to both this broader claim, and the narrower claim about temporally extended processes of reasoning.

about the color of the third peg. And suppose that, turning back to the game, she then formed a belief that the third peg was red for no reason, as a kind of arbitrary guess. The quality of the initial way of making use of reasons would be irrelevant to the rationality of the second belief about the peg, since that second belief was formed independently of that way of making use of reasons.

Similar points apply to the rationality of intention and action. Consider, for example, an expectant parent's intentional action of buying a bassinet, on the basis of reasoning about the benefits of buying the bassinet. Suppose the agent reasoned that the bassinet is worth the money since with a bassinet the baby can sleep right next to the bed, so feeding at night will be slightly less of an ordeal. His intention to buy the bassinet, and his action of making the purchase, are rationally assessable in virtue of the quality of that reasoning. The quality of the reasoning bears on the rationality of the intention and the action partly because the reasoning is a cause of the action and the intention. If the agent had bought the bassinet independently of that reasoning, the rationality of his action would not be affected by the quality of that reasoning. If he never went through the reasoning, the quality of the reasoning would not bear on the rationality of the action or intention. Or suppose he went through the reasoning, forgot about it, and then formed a new intention to buy the bassinet independently of that earlier reasoning. The quality of that reasoning would not bear on the rationality of his intention, or his action.

In seeking to understand the way that reasoning guides action, we are looking for relations between action and reasoning that help explain why reasoning bears on the rationality of action—where reasoning can bear a similar relation to intention and belief, which helps to explain why the quality reasoning bears on the rationality of intention and belief. Whether a piece of reasoning helps causally determine a belief, intention or action helps to explain why the quality of the reasoning affects the rationality of the belief, intention, or action. Thus, part of what it is for an action to be guided by practical reasoning is for the reasoning to be

a cause of that action.

Those who are suspicious of the view that intentions and reasoning causally explain action sometimes suggest that asserting such a causal relation is motivated by a too-dichotomous view about action and intention.⁷ The dichotomous view is that the action and the intention that causes it are causally related but not constitutively related—like the movement of one billiard ball, and the movement it causes in another ball. The right view, they suggest, is that what action you are carrying out, and whether you are acting at all, is determined partly by your intentions, and reasoning that your intention is based on. They suggest that intentions and reasons are “formal” causes of action, instead of being “efficient” causes. Intentions figure in constitutive explanations of the nature of one’s action, and not in causal psychological explanations of why one acts as one does.

It is very plausible that intentions, and other aspects of practical reasoning, figure in constitutive explanations of the nature of one’s actions. The movement of your arm counts an intentional action, and as a signaling, for example, partly in virtue of the intentions associated with the movement. But that point is compatible with the claim that our actions are caused by intentions, and by reasoning. An action can result from a piece of practical reasoning, and also count as a certain kind of action because it is based on that reasoning. Indeed, it would seem that an intention or a piece of reasoning can help individuate an action *only if* it plays some role in causally determining the action. If an agent came to move her arm as she did independently of her intention to signal, then her arm movement would not count as a signaling. The status of intentions as “formal” causes of action depends on their status as “efficient” causes of action.

A similar point applies to the relation between reasons and actions. An action can count as an act of insulting someone partly because the action was done because the agent believed the action would be hurtful to the other person. The statement that the action was done be-

⁷See Anscombe (1965).

cause the agent believed it would be hurtful is not a formal, constitutive explanation of what kind of action the agent performed. It is an element of the *explanans* in such an explanation. The action counts as an insult partly because the action is causally-psychologically explained by the agent's belief that the action would be hurtful to the other person. Reasons must figure in this causal-psychological way in explaining an action in order to figure in formal, or constitutive explanations of what kind of action the agent is engaged in.⁸

Function

Clearly, not everything that reasoning helps to causally determine is guided by reasoning. Many things that reasoning helps to causally determine are not rationally assessable. Reasoning can figure in causal explanations for headaches, for example. The headache produced by one's reasoning is not guided by the reasoning, and is likewise not open to assessment as rational or irrational in virtue of the reasoning that produces it.

An important difference between a headache produced by reasoning and proper objects of guidance of reasoning is that it is a part of the *function* of reasoning that it causally determines those objects of guidance. Actions and beliefs are rationally assessable partly because they are not "accidental" effects of reasoning, like headaches. Reasoning functions to causally determine its proper objects of guidance, like beliefs and actions.⁹

Epistemic reasoning helps determine what we believe according to its function. It is

⁸I have recently come to think that reasons do not have to have any causal powers in order to provide rational support for beliefs, intentions, or actions. I think that this point is compatible with the central claim here, that reasons play essential roles in causal-psychological explanations of the beliefs, intentions, and actions that they provide rational support for. I intend to address the relation between causation and rational support in detail in future work.

⁹A similar point about function can help to shed light on cases of "deviant causation" that complicate Davidson-inspired causal theories of action. In cases of non-deviant causation, one's intention causes one's behavior through the normal proper functioning of one's agential capacities. In cases of deviant causation, one's intention causes one's behavior, but not through the normal proper functioning of one's agential capacities.

essential to the nature of epistemic reason that it has this function.¹⁰ If epistemic reason did not help determine what we believe according to its function, then it would not be epistemic reason. Just as epistemic reason functions to determine what we believe, practical reason functions to determine what we do, i.e., our actions. If practical reason did not function to determine what we do, it would not be *practical* reason. If practical reason did not issue in action, or only issued in action in a way that was accidental with respect to its proper functioning—in the way that a chair can “accidentally” result in marks in the carpet—then it would not be practical reasoning.

As initial *prima facie* support for this point, it may help to consider a similar point about intention. Intentions, like practical reason itself, function to determine what we do. This fact marks a key point of difference between intentions and mere hopes and wishes. Mere hopes and wishes share a certain “world-to-mind direction of fit” with intentions. When what one wishes for does or does not come about, that is a kind of match or mismatch between the wish and the world. Similarly, when one acts or does not act as intended, that is a kind of match or mismatch between the intention and the world. For both wishes and for intentions, when there is a mismatch, that is not a “descriptive” error to be rectified by changing what one wishes or intends.

However, mere hopes and wishes differ from intentions in that they do not function to determine what we do. In that way, mere wishes are “inert” with respect to action, and are not practical in the way that intentions are. Intentions are practical, and not “inert” with respect to action, in that they function to determine what we actually do. It is essential to a mental state’s counting as an intention that it has this action-determining function. An intention that does not function to determine what we actually do might have a kind of world-to-mind fit, like an idle wish. But what part of what differentiates intentions from idle wishes is that intentions function to determine what we actually do, whereas idle-wishes do not.

¹⁰See Burge (2003) for detailed discussion of this point.

This point is reflected in the fact that an intention is a commitment to act. An intention's being a commitment to actually undertake an action, rather than a commitment to the desirability of some action one might, hypothetically, perform, requires that the intention function to issue in the intended action.

In a similar way, if one's reasoning about what to do did not function to causally determine what one actually does, that reasoning would not be practical. The reasoning would function, at most, as a way of figuring out what it might be nice to do, hypothetically, or as an academic matter. The reasoning would not function as a way of figuring out what *actually* to do. Practical reasoning cannot be hypothetical or academic in this way. It must function to determine what we actually do.¹¹

The core case for the claim that reasoning functions to causally determine its objects of guidance concerns the relation between functions and norms of success and failure. The rationality or irrationality of a belief, intention, or action consists in some kind of success or failure in reason's determining what one believes or intends, or in determining how one acts. This normative vocabulary of success and failure would be in error if the reasoning that makes a belief or action rational or irrational did not *function* to determine what one believes or how one acts.

Standards of success and failure are associated with, and explained by, functions.¹² An outcome is a success or a failure for an individual and her capacities only if, and partly because, that outcome fulfills or violates a function of those capacities. A toaster's failing

¹¹To be clear, I do not mean that an individual incapable of action could not be capable of practical reasoning. An individual who was completely paralyzed, and also incapable of mental action, might still be capable of practical reasoning. In such an individual, practical reason would still *function* to determine what the individual does. Her capacity for practical reasoning would be frustrated, because of defects in other aspects of her action system. In my view, if practical reason cannot issue in action, it must be because of defects in some aspect of the agent's action system. Practical reason functions to causally determine action according to its nature, and the nature of the kind of agent the practical reasoner is. If practical reason cannot issue in action, that incapacity marks a difference between the how the agent actually is, and how she should be according to her nature, and functions associated with her nature.

¹²See Burge (2009).

to brown bread is a failure for the toaster partly because browning bread is a function of the toaster. If strong winds fell a tree, that does not constitute a success or a failure for the winds. That is because the wind does not have a function of felling trees, or keeping them upright. It is not the wind's "job" to determine whether trees stand or fall. In contrast, epistemic and practical reasoning can *succeed* or *fail* in determining what one believes, or how one acts. Reasoning can succeed or fail in this partly because it *is* epistemic reason's "job"—its function—to determine what one believes; and it is practical reason's function to determine what one intends, and how one acts.

I want to emphasize that reasoning functions to causally determine our beliefs, intentions, and actions according to its nature. It is essential to the nature of one's capacity to engage in epistemic reasoning about what is the case that it functions to determine what one believes. It is essential to the nature of one's capacity to engage in practical reasoning about what to do that it functions to determine what one intends, and how one acts.

Our capacities to reason have psychological functions that are not essential to their natures. One kind of example concerns non-voluntary behavior that is "primed" by the conclusion of a piece of reasoning. Consider an agent going for a hike, who believes there are mountain lions in the vicinity on the basis of a piece of reasoning. Suppose that he startles at a quit snapping sound just off the trail. His startle reaction is "primed" by his belief that that there are mountain lions nearby. The reasoning that leads to that belief figures in a causal psychological explanation of why the agent startled. That is, the reasoning helps to causally determine the startle response. Plausibly, one of the psychological functions of reasoning is to help determine certain non-voluntary activity, like startle responses.

However, the capacities for reasoning that function to determine non-voluntary activity do not have that function according to their natures. In the case described above, the reasoning that primes the startle response is epistemic reasoning about whether there are mountain lions nearby. It is not essential to the nature of our capacities for epistemic reasoning that

they function to determine startles, or other non-voluntary activity. Another kind of being could possess capacities for epistemic reasoning that do not function to determine non-voluntary activity. But it is essential to the nature of epistemic reason that it function to determine beliefs. Likewise, it is essential to the nature of practical reason that it functions to determine intentions and actions.

Standards of correctness Reasoning does not function to causally determine beliefs, intentions, and actions just for the sake of proliferating those things. Reasoning determines our beliefs, intentions, and actions in order to get our beliefs, intentions and actions right. Beliefs, intentions, and actions are subject to standards of correctness. Reason's function of causally determining our beliefs, intentions and actions subserves the beliefs', intentions', and actions' functions of satisfying their standards of correctness.

Let me explain what I have in mind in saying that beliefs, intentions, and actions are subject to standards of correctness. For beliefs, the relevant standard of correctness is that the belief be true. If one's belief is false, the belief constitutes a mistake or failure in that respect. If one's belief is true, the belief constitutes a success in that respect. The success or failure in holding a belief is attributable to the individual. She errs or "gets it right" in holding her belief. Epistemic reasoning subserves belief's function of being true.¹³ Arriving at true belief is the point, or one of the points, of one's epistemic reasoning.

For intention the relevant standard of correctness is, roughly, that the intention represent an action that would in fact be good to do in the circumstances. If one intends an action that is in fact not good to do, one's intention constitutes a mistake or a failure in that respect. If one intends an action that is in fact good to do, one's intention constitutes a success in that respect. That success or failure in holding an intention is attributable to the agent. It is she who errs or "gets it right" in holding the intention. Practical reasoning subserves

¹³See Burge (2003).

intention's function of representing action that is good to do. The point, or one of the points, of practical reasoning is to arrive at an intention to carry out an action that is in fact good to do. Epistemic reason's "job" is to get us to correct belief, and practical reason's "job" is to get us to correct intention. Practical reasoning and epistemic reasoning each subserve their respective conclusions' function of satisfying their standards of correctness.¹⁴

I do not assume that the relevant sort of goodness must be moral goodness. An action can be good to do in virtue of its helping to realize a variety of ends, including the avoidance of pain or bodily harm, pleasure, companionship, scientific understanding, artistic achievement, relaxation, light amusement, and so on. I do not assume that these are all ends of morality, though nothing here depends on that issue.¹⁵

These standards on belief and intention are independent of warrant. One's belief can be false and warranted, or true and unwarranted. Similarly, one can intend an action that is in fact good to do, but lack warrant for the intention. For example, one might intend to take medicine that is in fact curative, while having strong misleading reasons to think the medicine is harmful and ineffective. One can also intend an action that is not good to do, but have warrant for the intention—if, for example, the medicine is harmful, but you have good reasons to think that it is curative.

¹⁴If making use of reasons subserves action's function of satisfying its standard of correctness, that might seem to suggest that the goodness of action does not depend on one's reasons for action. If so, then the claim I am making here would conflict with the sensible view that the goodness of one's action can depend on one's reasons for action, at least in some cases. As far as I can see, the claim I am making here actually does not imply that the goodness of action is independent of one's reasons for action. For example, whether one acts justly depends on the reasons one's actions are based on. But nonetheless, practical reason can function to guide one to act justly. The issues here are difficult, and worth discussing in detail. I set them aside here, and hope to address them in future work.

¹⁵In my view, the standard of correctness that serves as the "aim" of practical reasoning is that one's intended action helps, on the whole, to realize legitimate ends, like those listed above. Others claim that the aim of practical reasoning is action that makes sense to the agent (compare Velleman (2000)), or the constitution of our own agency (compare Korsgaard (2009)). Some are rumored to hold that the aim is action that satisfies the agent's desires. For present purposes it does not matter what the relevant standard of correctness is, which practical reasoning subserves. The point is that, if reasoning genuinely guides intention, action, and belief, then there must be standards of correctness on intention, action, and belief that one's reasoning subserves.

Like beliefs and intentions, actions are also subject to standards of correctness. An agent can err or get something right in performing an action. The standard of correctness on action is the same as the standard of correctness on intention. If in fact some intended action is good to do, one would thereby get something right in performing the action, as well as in forming the intention. Satisfaction of the standard of correctness on action is independent of norms of warrant, in the same way as intention. An action of taking medicine can be correct and unwarranted, or incorrect and warranted, just like an intention to take medicine.

The standard of correctness on belief does not have the same kind of double application as the standard of correctness on action and intention. Suppose an individual holds a belief that US policy is to do its best to undermine genuine democracy in Egypt. If that belief is true, that means one gets something right in holding the belief. But, of course, if that is in fact US policy, that does not mean the policy satisfies a standard of correctness. Even if the policy did satisfy some standard of correctness, the correctness of the policy would not constitute the belief-holder's getting something right—whereas the agent does get something right in holding a correct belief or intention, or in carrying out a correct action.

The double application of the standard of correctness on intention and action reflects the constitutive connection between intention and action. Intentions are intentions to act. They are individuated by the actions they represent. If an intention represents an action, it represents that action essentially, according to its nature, and it functions to causally determine the action according to its nature. In the other direction, actions are individuated partly by the intentions they execute. The intention executed by an intentional action is essential to the nature of the action. If one intentionally ϕ s, it is essential to the nature of the action that it is intentional, and the action's relation to the intention is what makes it count as intentional.

Just as practical reasoning subserves intention's function of satisfying its standard of correctness, practical reasoning also subserves action's function of satisfying its standard

of correctness. That is, roughly, the point of reasoning is not just to get one's intentions right, but to get one's actions right. Indeed, intention's function of satisfying its standard of correctness subserves action's function of satisfying its standard of correctness. The point of getting one's intentions right is to get one's actions right.¹⁶

This discussion of reasoning and standards of correctness helps to explain what it is for actions, as well as intentions and beliefs, to be guided by reasoning. The concept of guidance is closely linked to the functions described above. Getting an object of guidance to satisfy something like a standard of correctness is a function of any kind of guidance.

It is common, and appropriate, to use cases of archery to illustrate this kind of point. An archer's hitting her target is a case of satisfying a standard of correctness. Her guidance of her arrow in shooting subserves the satisfaction of that standard of correctness. Every case of guidance has some analogous relation to a standard of correctness. Guidance always subserves some end—what the object of guidance is being guided *to*, as it were. Even in apparently “aimless” cases of guidance, like guiding the direction of a kite, one tugs at the strings in certain way so as to get the kite to move in this or that direction, or to get it to stay in the air. The kite's actually moving in the desired direction, or actually staying in the air, is the success, the satisfaction of a standard of correctness, that one's kite-guiding behavior subserves. It is a conceptual point that guidance always subserves some such standard of correctness. The concept of guidance is, in this respect, teleological. There is no guidance without a “telos” for the object of guidance.

If actions, intentions, and beliefs are objects of guidance of reasoning, that means that the guidance by the reasoning subserves some standard of correctness of the action, intention, or belief. Actions, intentions, and beliefs count as objects of guidance of reasoning partly because reason functions to causally determine action, intention, and belief so that it “hits

¹⁶In other ways, intention is more fundamental. Rationality, irrationality, morality, immorality, and other practical evaluations apply more directly to intention than to action.

some target”—i.e. so that it realizes some standard of correctness.

The management problem

So far, I have suggested that reasoning functions to causally determine beliefs, intentions, and actions, and that this function subserves the belief’s, intention’s, or action’s function of satisfying a standard of correctness. Actions, as well as intentions and beliefs, count as objects of guidance of reasoning partly in virtue of that functional relation to reasoning.

But the discussion so far is only the beginning of an account of what it is for action to be guided by reasoning. To develop a fuller account, it will help to consider the difference between reasoning’s relation to action, and reasoning’s relation to objects of “managerial control,” like the layout of a house one is building, or the trained behavior of a dog.

The term ‘managerial control’ is due to Pamela Hieronymi (2009). In Hieronymi’s work, objects of managerial control are, roughly, things that we affect as indirect results of exercises of agency, and which we do not control directly. For example, the behavior of a trained dog is an object of managerial control, because the trainer brings about changes in the dog’s behavior indirectly through exercises of her (the trainer’s) agency, without having direct control over the dog’s behavior.

Hieronymi contrasts managerial control with what she calls “evaluative control.” Evaluative control is, roughly, the control we exert over our mental states in making up our minds. Objects of evaluative control are committal psychological states, whose adoption constitutes the “settling of a question.” A belief is an object of evaluative control because, in forming a belief that p , one settles for oneself the question of whether it is the case that p . An intention is an object of evaluative control because, in forming an intention to ϕ , one settles for oneself the question of whether to ϕ .¹⁷

¹⁷We have other kinds of direct psychological control over psychological phenomena that are not committal

Our control over our intentional actions—voluntary control—is a central kind of agential control, but it does not fall under either of the two types of control that Hieronymi characterizes. Our control over our actions is not evaluative control, since we only have evaluative control over committal psychological states. Actions are not psychological states. Our control over our intentional actions is also not a case of managerial control. One does not indirectly bring about one’s own behavior in the way that one indirectly brings about the behavior of a dog in training it.¹⁸ To get at the way in which actions are guided by practical reasoning, it will help to distinguish practical reason’s relation to action from its relation to objects of managerial control, like the behavior of a trained dog.

The account so far does not distinguish reason’s guidance of action (or intention, or belief) from reason’s guidance of objects of managerial control. According to the account so far, reasoning guides action in that reason functions to causally determine our actions, and that function subserves our actions’ satisfying their standard of correctness. But reasoning also bears this relation to objects of managerial control, and not just to one’s own actions (and beliefs and intentions). Practical reason’s “job” is not just to get to the agent to intend and perform good actions. An important part of practical reason’s “job” is to effect objects of managerial control, like things in the external world, according to reason’s dictates.

Consider a piece of reasoning along the following lines. “If the dog goes in the street, he might be hit by a car. So the dog should heel. I can get the dog to heel by commanding it to heel. So I shall command it to heel.” The reasoning described above functions to

psychological states. One has a kind of direct control over one’s daydreams which is not evaluative control over committal attitudes. One also has a kind of direct control over whether one makes an inference. Carrying out an inference is not a psychological state, and does not consist in the settling of a question in Hieronymi’s sense. Hieronymi’s notion of evaluative control does not address these kinds of control.

¹⁸However, in my view, there is another kind of indirectness in our control over our intentional actions. We exert control over our intentional actions through forming intentions, and through the exercise of certain non-rational action-guiding capacities involved in executing an action, like capacities for motor control. Our control over our actions is mediated in a way that our control over our beliefs and intentions is not. Our control over our actions is more direct than our control over objects of managerial control, but less direct than our control over our beliefs and intentions.

causally determine one's action of commanding the dog to heal, and that function subserves the action's function of being a good thing to do. But the reasoning also bears that relation to the dog's behavior. The reasoning functions to get the dog to heal. And that function subserves the dog's behavior's satisfying some standard of correctness, like avoiding harm. The "point" of getting the dog to heal is to get the dog to do something it should do, and not just to shape the dog's behavior willy-nilly.

But actions, like intentions and beliefs, figure as objects of guidance of reasoning in a way that objects of managerial control do not. This difference is reflected in the fact that objects of managerial control are not assessed as rational or irrational in virtue of the quality of reasoning that the controlling agent carries out. The dog does not count as acting rationally or irrationally in virtue of the reasoning of the managing agent, who commands the dog to heal. If the dog is open to assessment as rational or irrational for healing, it is in virtue of some success or failure in the exercise of the dog's own rational capacities, involved in the transition from its registration of the command to its act of healing. Similarly, if a house is designed on the basis of good reasoning by its architect, the house itself is not assessable as rational or irrational in virtue of the reasoning of architect—although of course it is assessable as rationally or irrationally *designed*. In contrast, actions, like beliefs and intentions, are assessable as rational or irrational in virtue of the agent's reasoning.

One relevant difference between actions, intentions, and beliefs, and objects of managerial control, is that objects of managerial control are often not exercises of the reasoning agent's abilities. For example, it is the dog that heals, and not its trainer. In contrast, a reasoner's beliefs, intentions, and actions are exercises of the reasoner's own abilities.

However, objects of managerial control can be exercises of the agent's own abilities as well. Just as one can indirectly control a dog's behavior, one can indirectly control what one thinks and does oneself. For example, an individual who thinks her life is worthless might go to therapy in order to come to value her life. If she is successful, she acquires a belief that

her life is worthwhile, which is an object of managerial control with respect to the reasoning that her decision to go to therapy is based on.¹⁹ Unlike a designed building, or the behavior of a trained dog, her belief is an exercise of her abilities.

In a case like that, there are essential aspects of the way the agent “manages” her values that are not exercises of her abilities. The route from her reasoning that it would be good to value her life to her actually coming to value it goes through her therapist. Her therapist’s role in changing her outlook is not an exercise of her abilities. What the therapist says is his own doing, and not his patient’s.

When reason guides action, intention, and belief, it does so as part of the exercise of an overarching ability of the reasoner’s to determine what she believes, intends, or how she acts, on the basis of reasons. If a belief or an intention is based on reasoning, forming the belief or intention under reason’s guidance is, *as a whole*, an exercise of such an ability. Every integral aspect of the determination of one’s belief or intention under reason’s guidance is an aspect of the exercise of these overarching abilities.

This point is nearly trivial as it applies to beliefs and intentions, since beliefs and intentions are elements in one’s reasoning, when they are based on reasoning. Forming the belief or intention is part of the exercise of one’s ability to engage in epistemic or practical reasoning.

The same kind of point also applies to action, even though actions are not elements in one’s reasoning. We have abilities to act on the basis of reasons, as well as abilities to form intentions and beliefs on the basis of reasons. This ability to determine how we act on the basis of reasons is our practical rational agency. When an action is guided by practical reasoning, carrying out the action under reason’s guidance is, as a whole, an exercise of the agent’s abilities. Every aspect of the determination of one’s acting under reason’s guidance is an aspect of this ability. That includes both the reasoning, and the

¹⁹Her belief is also an object of evaluative control, in addition to being an object of managerial control.

performance of the action, as well as the transition from the reasoning to the action. Our practical rational agency comprises more than just capacities to reason, since performing an action constitutively involves non-rational capacities, like capacities to move one's arms and legs. Those non-rational capacities, however, are part of an over-arching ability to act on the basis of reasons, as are our capacities for practical reasoning.

When reasoning guides action, intention, or belief, it does so as part of an overarching ability—where the exercise of that ability comprises the reasoning, its object of guidance (i.e., the belief, intention, or action), and everything in between. That point helps to differentiate the way reasoning guides action, intention, and belief from the way that reasoning influences objects of managerial control.²⁰

One might object to this claim by arguing that objects of guidance of reasoning need not consist just in exercises of the agent's abilities. If an agent raises a glass, the glass seems to have a constitutive role in her action. You cannot raise a glass without a glass. But the glass's role in the action is not an exercise of the agent's abilities. The agent's physical movements, through which the agent raises the glass, are exercises of her abilities. But the motion of the glass itself is not. As a result, it might seem not to be the case that a piece of reasoning, together with its proper object of guidance, consists through and through in exercises of the agent's abilities. Certain constitutive aspects of one's acting under reason's guidance seem not to be part of the exercise of an overarching ability.

But there is an important difference between the role of the glass in the action, and the role of exercises of the agent's physical abilities. We should distinguish between the agent's performance of her action, and factors that help to determine the nature of the action that she is performing. The glass helps to determine the nature of the action, in that it plays a role in individuating the action. Her action would not count as a raising of a glass if not

²⁰At least, that point helps to differentiate reason's guidance of action, intention, and belief from many cases of managerial control. Reason's relation to certain objects of managerial control might not differ in this respect from its relation to action, intention, or belief. See the discussion four paragraphs down.

for the glass and its motion. But the glass is not part of the agent's *performing* her action. The glass's motion is something she makes happen, but not something she does. In contrast, the agent's physical movements are part of her performing her action. The glass and its movements help make her physical movements count as the raising of a glass. But the glass is not part of the phenomenon itself, the performance of the action.

I do not mean to claim that, apart from the action of raising the glass, there is a further action of *performing* the action of raising the glass. The performance of the action is the action. The claim I have in mind is that the glass and its motion play a role in individuating the action, but are not part of the action.

The action can be an object of guidance of reasoning because the reasoning, the performance of the action, and everything in between are aspects of the exercise of an overarching ability. Though the glass helps to individuate the action, it is not part of the performance of the action (i.e., it is not part of the action). So the fact that the glass's movements are not exercises of the agent's abilities is ultimately not threatening to the point here, concerning abilities and objects of guidance of reasoning.²¹ Reasoning functions to determine its objects of guidance as part of the exercise of an overarching ability, which includes the reasoning, its object of guidance, and everything in between.

As the arguments presented above show, this point straightforwardly rules out many objects of managerial control from being objects of guidance of reasoning. Many cases of managerial control (the cases discussed above) involve an external object, like a dog or a therapist. But some cases of managerial control do not involve any external object. An agent's managerial control over herself can be entirely endogenous. If a Kavka-style²² eccentric billionaire offers an individual a large reward if she believes that she is reciting

²¹The issues broached here, concerning the role of external objects in intentional action, are interesting and complex, and deserve more discussion than I give them here. I hope to address this topic in more detail in future work.

²²The allusion is to Kavka (1983).

the alphabet, she can bring about that belief by reciting the alphabet. No external object is involved in her bringing about that belief. The belief, and the reasoning about whether and how to bring about the belief, as well everything in between, are exercises of the agent's abilities. Still, the belief is an object of managerial control, and is not guided by her reasoning in the same way that actions, intentions, and beliefs are guided by the reasoning they are based on.

One might try to deal with such a case of managerial control by arguing that there is no *over-arching* ability in that kind of case. The agent exercises an ability to recite the alphabet on the basis of reasons, and exercises a separate ability to form beliefs about what she is doing. One might argue that the various abilities involved in acting on the basis of reasons are not separate in this way, and have a kind of unity. Perhaps a response like this can be made to work. In my view, a better response is to specify what kind of overarching ability in particular is involved in cases of guidance by reasoning. I develop this response in the next section. The rough idea is that in reasoning one *makes use of reasons* to get one's actions right, in the same way that one makes use of reasons to get a belief or an intention right.

Reasons, reasoning, and guidance

As suggested above, beliefs, intentions and actions are subject to standards of correctness. Reasoning subserves the functions of belief, intention and action of satisfying their standards of correctness. In this way, in reasoning one "attempts" to get one's beliefs, intentions and actions right. Epistemic reasoning subserves belief's function of satisfying its standard of correctness, namely being true. Practical reasoning subserves intention's function of satisfying its standard of correctness, namely representing an action that is good to do. Practical reasoning also subserves action's function of satisfying its standard of correctness, namely

being an action that is good to do.

In reasoning to a belief or an intention, one makes use of reasons that function to indicate that the belief or intention satisfies its standard of correctness.²³ To take the case of belief, in reasoning to a belief, one makes use of reasons to indicate that the belief is true, and satisfies that standard of correctness. For example, in reasoning to a belief about US policy in Egypt, one might make use of one's belief that Egyptian public opinion is hostile to the US to indicate the truth of the belief that the US will do its best to thwart genuine democracy in Egypt. That belief about Egyptian public opinion may or may not in fact be a good indication of the truth of the conclusion. But in either case, in making use of the belief in one's reasoning, one makes use of it to indicate that the conclusion is true, and satisfies its standard of correctness. That is a central function of a reason in one's epistemic reasoning.

In reasoning to an intention, one makes use of reasons that function to indicate that the intended action is good to do, and thus to indicate that the intention satisfies its standard of correctness. For example, in reasoning about whether to vote for Obama, one might make use of a belief that Obama orders illegal attacks in Yemen and Pakistan as a reason to vote for someone else. One makes use of that reason to indicate that voting for someone else would be good to do, and that the intention thus satisfies its standard of correctness. Whether or not that belief actually is a good indication that the intended action would be good to do, it functions in the individual's reasoning to indicate that it is.

I want to ward off a possible misunderstanding here. I do not claim that in reasoning to an intention, one represents the intention and its correctness in the content of elements of the reasoning. Likewise, I do not claim that in reasoning to a belief one represents the belief and its correctness in the content of elements of the reasoning. In reasoning to a belief

²³I use the term 'indicate' instead of 'establish' or 'show' in an effort to avoid suggesting that the relation between reasons and beliefs or intentions must be demonstrative, logical relations, or practical analogs of demonstrative relations (like the relations that correspond to rules of practical inference described by Kenny (1966) and others.

about whether the US supports democracy in Egypt, the subject matter is US policy toward Egypt, and not one's own beliefs. In reasoning to a conclusion about how to vote, the subject matter is the action of voting, and not one's intention.²⁴

In reasoning to a conclusion, like a belief or intention, one makes use of reasons to indicate that the conclusion satisfies its standard of correctness. The reasoning subserves the conclusion's function of satisfying its standard of correctness. Thus, reason functions to get the conclusion right, through making use of reasons that indicate that the conclusion is right. Reasoning can bear on one's action in this same way, even though the action is not an element in one's reasoning. Reasoning functions to get one's actions right, through making use of reasons that show that the action is right.

As stated above, actions are subject to standards of correctness. The standard is the same as the standard on intention. If in fact some intended action is good to do, one would thereby get something right in performing the action, as well as in forming the intention. The double application of the standard of correctness on intention and action reveals that the reasons one makes use of in practical reasoning bear on both action and intention at the same time, in the same way. In practical reasoning one makes use of reasons to indicate that the intended action is good to do. If the action is good to do, that marks a standard of correctness on action as well as intention. In practical reasoning, one makes use of reasons to indicate both that the intention satisfies its standard of correctness, and to indicate that the intended action satisfies its standard of correctness. For example, in the case described three paragraphs up, one makes use of the belief about Obama's attacks in Yemen and Pakistan to indicate both that the intention to vote for someone else is correct, and to indicate that the action of voting for someone else is correct.

²⁴That is not to say that a belief or intention can never be part of the content of one's reasoning. In critical reasoning, where one guides one's belief or intention partly through reasoning about standards of epistemic or practical justification, the belief or an intention can be represented in the content of elements of the reasoning.

Practical reasoning subserves action's function of satisfying its standard of correctness, as well as intention's function of satisfying its standard of correctness. Thus, in practical reasoning one makes use of reasons that function to establish that one's action satisfies its standard of correctness, and the reasoning subserves action's function of satisfying its standard of correctness. In this way, reason functions to get one's action right, through making use of reasons that indicate that the action is right.

Action has this relation to reasoning in common with belief and intention. This relation to reasoning is part of what makes something an object of guidance of reasoning. An object of guidance of reasoning is something that the reasoning functions to get right through making use of reasons that function to indicate that it is right. Epistemic reasoning functions to get beliefs right in this way. Practical reasoning functions to get actions right in this way, as well as intentions.

Accordingly, this relation between reasoning and actions, intentions, and beliefs helps to explain why actions, intentions, and beliefs are rationally assessable in virtue of the quality of one's reasoning. When success or failure in reasoning affects the rationality of a belief, intention, or action, the relevant success or failure is a matter of the way the individual makes use of reasons to indicate that the belief, intention, or action satisfies its standard of correctness. If the individual did not make use of reasons to indicate that her belief was true, or that the intended action was good to do, the quality of the reasoning would not affect the rationality of the belief, intention, or action.

The same point applies again to action. When success or failure in one's reasoning bears on the rationality of one's action, the relevant success is a matter of the way one makes use of reasons to indicate whether the action is good to do. The quality of one's practical reasoning bears on the rationality of one's action partly because in practical reasoning one makes use of reasons to indicate whether the action is good to do. Actions, intentions, and beliefs, are rationally assessable in virtue of reasoning partly because in reasoning one makes use of

reasons to indicate that one's action, intention, or belief satisfies its standard of correctness.

Back to managerial control The way one makes use of reasons in practical reasoning helps shed light on the difference between reason's relation to objects of managerial control and its relation to objects of guidance of reasoning, like beliefs, intentions, and actions.

Let me preface my discussion of this point by saying what the difference is *not*. One might want to suggest that in reasoning, one does not make use of reasons to indicate that objects of managerial control satisfy standards of correctness. But that claim is mistaken, at least in some cases. For example, in reasoning about whether to command a dog to heel, one might make use of reasons that purport to indicate that it would be good for the dog to heel. One might reason that the dog should heel, since otherwise it may run in the street and be hit by a car, and arrive on those grounds at an intention to command the dog to heel. In such a case, the agent seems to make use of reason's that function to indicate that the dog's behavior satisfies its standard of correctness.

While some elements of that reasoning do function to indicate that the dog's behavior satisfies a standard of correctness, the agent's reasoning *as a whole* does not consist in her making use of reasons that function to indicate that the dog's behavior is correct. There are essential elements of the reasoning that play no role in indicating whether the dog should heel. In particular, the transition from the judgment that the dog should heel to the intention to command it to heel plays no such role. There will be some such element in any piece of reasoning involved in cases of managerial control. If some elements of the reasoning concern what change to effect in the world, other elements of the reasoning will concern how to bring that change about. Those latter elements of the reasoning do not function to indicate whether the object of managerial control satisfies a standard of correctness.

Indicating whether an object of guidance of reasoning satisfies a standard of correctness is a *global* function of one's reasoning, and not a function subserved just by particular elements

of the reasoning. When an action, intention, or belief is a proper object of guidance of reasoning, every integral, non-adventitious aspect of the reasoning has a constitutive role in one's making use of reasons to indicate that the object of guidance satisfies its standard of correctness. In epistemic reasoning, every integral aspect of the reasoning has a constitutive role in one's making use of reasons to indicate that the belief is true. In practical reasoning, every integral aspect of the reasoning has a constitutive role in one's making use of reasons to indicate that the intended action is good to do.²⁵

Reasoning is, in all its integral aspects, the exercise of an ability to get something right, through the use of reasons that indicate that it is right. An object of guidance of reasoning is something that reasoning functions globally to get right through the use of such reasons.

A worry about feasibility One might question whether practical reasoning bears this relation to intention and action, because reasoning about the feasibility of an intended action can be part of one's practical reasoning. For example, in reasoning about whether to drive to work or take the bus, an individual might make use of reasons concerning whether or not her car has been repaired, in order to figure out whether driving is a feasible option for her. Reasoning about whether an action is feasible is part of the reasoning that guides one's intention and action. Accordingly, an intention to drive to work, for example, can violate rational standards because the individual errs in the way she makes use of reasons to establish the feasibility of doing so—if, for example, she believes that her car has been repaired on the basis of bad reasons.

It is natural to draw a distinction between reasoning that concerns whether ϕ ing would be good to do, and reasoning that concerns whether ϕ ing is a feasible option. As a result, it might seem that there can be elements in reasoning that guide action and intention, which do not function to indicate that the intended action is good to do, and satisfies its standard of

²⁵Compare Burge (2000).

correctness. If so, that point would undermine my claim that all integral aspects of reasoning have a constitutive role in one's making use of reasons to indicate that the object of guidance of the reasoning satisfies its standard of correctness.

To respond to this worry, it will help to consider a similar issue that arises in cases of epistemic reasoning. When reasoning guides beliefs, reasoning about whether forming the relevant belief is feasible does not enter into the reasoning. But a similar issue does arise for beliefs, since one's reasons for a belief may concern presuppositions for the correctness of the belief or intention.

Suppose an individual believes that her neighbor is back from vacation. Her belief may be guided partly by reasons that function to indicate that the neighbor was away on vacation in the first place. One might be tempted to claim that those reasons do not function to indicate that the neighbor is back, and that they only function to indicate that the neighbor went away. As a result, it may seem that there are elements of the epistemic reasoning that guides a belief that do not function to indicate that the belief satisfies its standard of correctness (i.e., that the belief is true).

In the case of the epistemic reasoning that concludes in a belief, it is clear that the overall function of the reasoning is to make use of reasons to indicate whether the belief is true. Whether one's neighbor is back from vacation depends on whether she was away in the first place. As a result, figuring out whether one's neighbor was away in the first place can be part of figuring out whether the neighbor is back from vacation. The reasons one makes use of to indicate whether that background condition can be part of the case for the primary issue of whether the neighbor is back.

A similar point applies to practical reasoning. Practical reasoning concerns what is *actually* good to do, not what *would* be good to do, hypothetically, if only one could. Whether it is actually good to ϕ depends on whether it is feasible to ϕ . As a result, figuring out whether it is feasible to ϕ can be part of figuring out whether it is actually good to

ϕ .²⁶ Thus, the reasons one makes use of in reasoning about whether an intended action is feasible can play a role in indicating whether the intended action is good to do, and whether it satisfies its standard of correctness.

Conclusion

Epistemic reasoning guides our beliefs, and practical reasoning guides both our intentions and our actions. My aim in this chapter has been to provide an account of what makes actions objects of guidance of reasoning, by looking at parallels between reason's relation to action, and reason's relation to intention and belief. These parallels are supposed to help explain why actions, like intentions and beliefs, are rationally assessable.

When beliefs and intentions are guided by reasoning, they are part of the reasoning that guides them. They are conclusions of one's reasoning.²⁷ In contrast, when actions are guided by reasoning, they are not elements in the reasoning that guides them. The fact that actions are not elements in reasoning does not prevent them from being guided by reasoning, and from being rationally assessable in virtue of one's reasoning. Being a concluding step in a piece of reasoning is not essential to being an object of guidance of reasoning. In order to make progress in understanding what is essential, I have described several relations that reasoning bears to its conclusions, which it also bears to actions.

First, reasoning can rightly be cited in causal psychological explanations of one's actions, just as it can be cited in psychological explanations of intentions and beliefs. Second, rea-

²⁶In contrast, whether it is true that p does not depend on whether it is feasible to believe that p . That point helps to explain why the feasibility of forming a belief does not enter into epistemic reasoning.

²⁷A belief can be guided by reasoning when it is the sub-conclusion of one's reasoning, and not the final conclusion of one's overall reasoning. For example, an individual might reason that his wife is not home from work on the basis of his belief that her car is not in the driveway. And he might reason that his daughter is watching television on the basis of his belief that his wife is not at home, and perceptual belief about the blue light he sees through the living room window. In a case like that, the sub-conclusion—his belief that his wife is not home—is the conclusion of a piece of reasoning itself. That piece of reasoning is an element of the overall reasoning that guides the final conclusion, that his daughter is watching television.

soning *functions* to causally determine our actions, as well as our intentions and beliefs. This function of reasoning subserves our actions', intentions', and beliefs' functions of satisfying their respective standards of correctness. Third, reasoning functions to determine our actions, intentions, and beliefs, as part of an overarching ability. The reasoning, and the action, intention, or belief, and everything in between, are elements of the exercise of one's overarching ability to act, intend, or form a belief on the basis of reasons. Fourth, in practical or epistemic reasoning one makes use of reasons to indicate that one's action, intention, or belief satisfies its standard of correctness. In this way, reasoning functions to get one's actions, intentions, and beliefs right, through making use of reasons that indicate that they are right.

These relations between actions and reasoning help to explain how actions figure as objects of guidance of reasoning in much the same way as conclusions of reasoning, even though actions are not conclusions of reasoning. The core point turns on the double application of the standard of correctness on intention and action. The standard of correctness on intention is also a standard of correctness on the intended action. As a result, the relation between reasoning and action is very similar to the relation between reasoning and intention. Our capacity to get our intentions right through reasoning is also a capacity to get our actions right through reasoning.

In making use of reasons to indicate that one's intention satisfies its standard of correctness, one also makes use of reasons to indicate that one's action satisfies its standard of correctness. And making use of reasons in this way subserves action's function of satisfying its standard of correctness, just as it subserves intention's satisfying its standard of correctness.

Thus, practical reasoning bears on action in much the same way as it bears on intention, even though intentions are concluding elements in practical reasoning, and actions are not. Epistemic reasoning bears on belief in this way as well. The fact that reasoning bears on

action, intention, and belief in this same way helps explain why actions, intentions, and beliefs count as guided by reasoning, and are rationally assessable in virtue of one's reasoning.

Chapter 6

Practical Warrant

Introduction

This chapter presents accounts of practical justification and practical entitlement, using Tyler Burge's distinction between epistemic justification and epistemic entitlement as a template.¹ Epistemic justification and epistemic entitlement are types of epistemic warrant. Roughly, a belief's epistemic warrant is a matter of an individual's (and her belief-guiding sub-systems') doing a good job determining what is the case, instead of arriving at a belief in an arbitrary or misguided way. Burge's discussions of epistemic entitlement reveal two ways in which a belief's epistemic warrant is not just a matter of the way the agent makes use of reasons for the belief. First, a belief's warrant can derive from exercises of non-rational belief-guiding capacities, and not just from exercises of reason. Second, a belief's warrant is not just due to operations of one's belief-guiding capacities (rational or non-rational), but also to epistemically good characteristics of those capacities.

Both these points conflict with certain internalist claims about epistemic warrant. Epistemic internalism includes a variety of views. One internalist claim is that factors that

¹See, for instance, Burge (1993) and (2003).

warrant beliefs are accessible to conscious awareness through some kind of direct reflection. Another internalist claim is that we are responsible for factors that warrant beliefs, and are open to praise or blame on account of those factors.² Another internalist claim is that beliefs are warranted just by exercises of reason.³ Epistemic entitlement is a warrant for a belief, or an element of a warrant for a belief, which is not internalist in these ways. Epistemic entitlements are not exercises of reason; one is not responsible for epistemic entitlements; and epistemic entitlements are not, or not always, accessible to conscious awareness.

This chapter's discussion of practical entitlement supports analogous anti-internalist points about practical warrant. The only discussion of practical analogs to epistemic internalist positions is Audi's "Internalism about actions" (1990).⁴ I argue that practical warrant does not consist exclusively in exercises of one's rational capacity to hold commitments on the basis of reasons. I suggest that there are non-rational elements of practical warrant, which are non-internalist elements of practical warrant. They are not exercises of rational capacities, they are not always accessible to conscious awareness, and one is not responsible for them.

I begin with a stage-setting discussion of epistemic justification and epistemic entitlement. I then present an account of practical justification. Practical justification for an action is a type of practical warrant that conforms to internalist views about practical warrant. Practical justification is roughly a matter of the rational support that actions and intentions get from reasons. That is the type of practical warrant that Audi focuses on. In the rest of

²Though internalists do not claim that we are *morally* responsible for factors that warrant beliefs. They claim we are "epistemically responsible" for factors that contribute to warrant.

³For helpful surveys of the debate between epistemic internalists and epistemic externalists see Alston (1989), and Bonjour and Sosa (2003).

⁴Audi argues that the reasons that contribute to an action's rationality must be reflectively accessible to the agent. I do not argue against that view here. My interest here is in elements of practical warrant other than reasons. Audi's claim about reasons for action seems plausible to me, at least regarding reasons attributable to an agent at the whole-individual level. If there are modular capacities to operate with propositional representational states, those representational states might be counted as reasons for action, and need not be accessible to the agent.

the chapter, sections 4-6, I discuss practical entitlement for actions and intentions. Practical entitlement is a non-internalist type of practical warrant.

Epistemic internalism includes a diverse group of views. Practical analogs of some of these views might be compatible with the points about practical entitlement advanced in this chapter. I do not discuss the bearing of practical entitlements on all possible versions of internalism about practical warrant. My main aim is to support the claim that practical warrant often consists in practical entitlement, and to point out that instances of practical entitlement are counterexamples to practical analogs of a few common commitments of epistemic internalists—in particular, the view that warranting factors are exercises of reason, that they are directly accessible to consciousness, and that we are responsible for them.

Epistemic warrant

The intuitive idea behind the concept of epistemic warrant is, roughly, that a warranted belief is a belief that one forms through doing a good job determining what is the case, instead of forming the belief in an arbitrary or misguided way.

A belief can count as warranted, rather than as misguided, partly in virtue of the way one makes use of reasons to indicate the truth of the belief. We have a capacity to make use of reasons in forming our beliefs. This capacity functions to guide one to true belief. Exercises of this rational capacity contribute to epistemic warrant when they perform well in carrying out this function, given limitations in the individual's perspective and in the rational capacity itself. Epistemic justification is the type of epistemic warrant that is provided by exercises of this capacity to form beliefs on the basis of reasons.

Epistemic justification contrasts with epistemic entitlement. The way one makes use of reasons in forming a belief does not fully explain what makes a belief warranted, at least in most cases. And in many cases a belief is warranted even though one makes no use of

reasons at all. Epistemic entitlement is epistemic warrant that does not consist in the way one makes use of reasons.

We can distinguish between two types of epistemic entitlement. One type consists in exercises of non-rational belief-guiding capacities. Entitlement for a belief can derive from perception and sensation. Like our capacity to make use of reasons, capacities for perception and sensation function to guide our beliefs. I form my belief that it is raining partly under the guidance of perception of the rain. I form my belief that the bread is burning partly under the guidance of my sensory registration of the smell coming from the oven. Beliefs like these are warranted partly in virtue of the guidance of such non-rational capacities. Indeed, only non-observationally warranted beliefs, like beliefs about math and the marital status of bachelors, are warranted independently of guidance by non-rational capacities. When epistemic warrant consists in exercises of non-rational belief-guiding capacities, that warrant is an epistemic entitlement.

But epistemic entitlement is not due exclusively to the exercise of non-rational belief-guiding capacities. For example, consider our epistemic entitlement to rely on memory.⁵ In epistemic reasoning through time, we are entitled to rely on exercises of preservative memory. One relies on preservative memory to accurately preserve earlier conclusions in reasoning, and that reliance is warranted. At least in most cases, when we rely on preservative memory, we are not guided to do so at all. There are no psychological processes that lead the individual to rely on her preservative memory. Such non-guided reliance on memory is not a failure in the functioning of the belief system. The non-guided reliance is warranted. At least in most cases, we have an entitlement to rely on memory that does not derive from any kind of guidance, including guidance by non-rational capacities.

We are entitled to rely on preservative memory because of the essential role that it plays in the effective functioning of epistemic reason. It enables effective reasoning by reliably

⁵See Burge (1993).

preserving the conclusions of earlier reasoning for use in later reasoning. Memory's reliability in carrying out this role in epistemic reasoning helps to explain why we are entitled to rely on memory. Memory's reliability helps to explain why it is epistemically good to rely on preservative memory, and not epistemically misguided or arbitrary. As a result, our entitlement to rely on memory consists partly in memory's reliability, even though memory's reliability is not an exercise of non-rational belief-guiding capacities.

These points about epistemic entitlement conflict with internalist conceptions of epistemic warrant. There are three main commitments commonly associated with epistemic internalism. First, we are responsible for factors that warrant our beliefs. If there is something errant in one's way of forming a belief, but one is not responsible for that error, the mistake does not count against the belief's warrant. A second commitment is that whatever helps warrant a belief is directly accessible to conscious awareness, either because it is a conscious state, or because it is available for use in conscious reasoning. This commitment is often taken to follow from the first one. If warrant were to depend on some factor outside of one's ken, inaccessible to conscious awareness, then one would not be culpable if that factor failed to obtain.⁶ As a result, one could lack warrant without making mistakes for which one is responsible. A third commitment is that epistemic warrant is a matter of the way one makes use of reasons. This commitment dovetails with the other two. Internalists claim that reasons for belief are accessible to conscious awareness. And they claim that one is responsible for errors in the way one makes use of reasons.⁷

⁶If something is accessible to consciousness, that does not imply that one is responsible for it. For example, a sensation might be consciously accessible to the agent even if she is not responsible for having that sensation. One lacks direct control over the sensation, even though it is consciously accessible. Some types of accessibility might guarantee responsibility, or at least come close to guaranteeing it. For example, if a belief is accessible for use in conscious reasoning, that suggests that one has direct control over the belief. If so, then accessibility for use in conscious reasoning might imply a kind of responsibility for the belief. On the other hand, perhaps an individual could have this kind of access to a belief, but not be responsible for it because she has severe cognitive impairments.

⁷More specifically, internalists claim that one is responsible for how one makes use of reasons, and has access to one's reasons, when the reasons warrant or justify one's belief. Internalists can allow that one

The claim that epistemic warrant consists partly in guidance by non-rational capacities like perception and sensation conflicts with these internalist commitments. It conflicts with the third because perceptions and sensations are not reasons for belief.⁸ The claim also conflicts with some versions of the second internalist commitment. It conflicts with the claim that factors that contribute to epistemic warrant are accessible for use in conscious reasoning. Perceptions and sensations cannot be elements in reasoning, so *a fortiori* cannot be elements in conscious reasoning. They are also sometimes unconscious and inaccessible. Paradigm cases of perceptual and sensory states are accessible to conscious awareness in that they are conscious states. Many perceptions and sensations also can be accessed consciously in episodic memory. But some perceptual and sensory states can contribute to epistemic warrant without being conscious, or accessible in episodic memory. Non-conscious proprioceptive sensory input is an example. The role of perception and sensation in warranting belief also conflicts with the first internalist commitment as well. One is not responsible for failures or successes in the exercise of one's perceptual and sensory capacities. It is also likely that, at least in many cases, one is not responsible for failures or successes in transitioning from perceptions and sensations to beliefs.

The claim that epistemic warrant does not consist exclusively in the exercise of belief-guiding capacities conflicts with these internalist commitments in similar ways. First, taking memory's reliability as an example, the reliability of memory in preserving conclusions of reasoning is not an exercise of one's capacity to make use of reasons. Second, memory's reliability cannot be an element in reasoning, and *a fortiori* cannot be an element in conscious reasoning. It can be *represented* in a belief that can be an element in reasoning (though it almost never is), but it is never itself an element in reasoning. The reliability of memory

might not be responsible for one's reasons, or have access to them, if one has certain cognitive impairments. But, in that case, one's reasons would not warrant or justify one's belief, according to internalists.

⁸I discuss the claim that perceptions and sensations are not reasons in the beginning of section 4.

also cannot be a conscious mental state, since the reliability is not a mental state at all, but a characteristic of a mental state. Third, an individual's memory could be extremely unreliable, without the individual's having any indication of the problems with her memory. That unreliability could undermine her epistemic warrant for relying on memory and for holding a belief, even if she makes no culpable error in the reliance on memory, or in holding the belief.⁹

I want to defend similar anti-internalist points about practical warrant. I begin by discussing practical justification in the next section. I discuss practical entitlement in the subsequent sections.

Practical justification

As I said earlier, the intuitive idea behind the concept of epistemic warrant is that warranted belief is roughly a matter of doing a good job in arriving at a belief, instead of forming a belief in an arbitrary or misguided way. Similarly, the rough, intuitive idea behind practical warrant is that warranted intention or action is a matter of doing a good job in arriving an intention or action, instead of arriving at the intention or action in an arbitrary or misguided way.

Often our intentions and actions are warranted in virtue of the way they are based on reasons. If one intends to ϕ on the basis of reasons that indicate that it would be good to ϕ , the way the intention is based on reasons helps to show why the intention is not formed in a

⁹It is possible for one's memory to be reliable in its own functioning, while not being reliably causally connected to other elements of one's belief system. In a case like that, one would lack warrant for relying on one's memory, even though it is reliable. Warrant for relying on memory depends on its being reliably causally connected with other elements of one's belief system, as well as on the reliability of memory itself. Analogous points can be made about the role of reliability various action guiding capacities in providing for practical warrant. I focus here on the reliability of the capacities themselves, and not on the importance of their reliable causal connection with other aspects of one's agency. There is some discussion of the latter issue in my "Entitlement for reliance on physical abilities."

misguided or arbitrary way. The intention is not formed in a misguided way partly because the agent arrived at her intention through doing a good job in making use of reasons to figure out what to do. Practical justification is the type of practical warrant that consists in exercises of the agent's capacity to make use of reasons to figure out what to do.

Actions and intentions are subject to a standard of correctness analogous to truth for belief. In practical reasoning one makes use of reasons that function to indicate that one's action and intention satisfy their standards of correctness. Roughly the standard of correctness on action is that it be good to do, on the whole, given the context. The standard of correctness on intention is, again roughly, that it represent action that is good to do. Exercises of practical reason function to indicate whether an intended action is good to do. Exercises of practical reason also function to causally determine our intentions and actions (or to help causally determine, in the case of actions—exercises of practical reason are not sufficient to produce action). These functions of practical reason subserve action's and intention's satisfaction of their standards of correctness. In this way, practical reason functions to guide one to correct intentions and actions.¹⁰ Standards of practical justification mark successes in the exercise of practical reason with respect to this guidance function.

Roughly, an intention is justified when practical reason does a good job in guiding the agent to that intention. Somewhat more carefully, an intention is justified when it is causally determined through the normal proper functioning of practical reason, wherein practical reason performs well in making use of reasons to establish that the intended action satisfies its standard of correctness.¹¹ Any action that constitutes an attempt to execute a justified

¹⁰I discuss these points in more detail in "The rationality of actions."

¹¹One might object to this account already, by suggesting that an intention can be justified even if the agent does not form the intention on the basis of any reasons. Hursthouse (1991) and Anscombe (1965) suggest that an action need not be misguided just because it is not based on reasons. (Though that claim is contested, by Quinn (1993), for example.) I reserve the word "justification" for cases in which intentions and actions are based on justifying reasons. But it seems to me plausible that an action or intention can have practical warrant independently of any justifying reasons, and even independently of any guidance by non-rational capacities, like sensory capacities and motor capacities. I suspect that we have a default

intention is itself justified. The action is justified whether or not it is well-executed—as long as any failure in execution is not due to a failure in the way the intention is based on reasons. The justification one has for a justified intention or action consists in the totality of the reasoning that satisfies norms of practical justification, including both the mental states that figure in that reasoning, and the transitions between mental states.

This notion of practical justification is teleological, in that norms of practical justification concern success in guidance to some telos. Exercises of practical reason that satisfy those norms subserve action's and intention's satisfaction of their standards of correctness. The satisfaction of those standards of correctness is a telos of those exercises of practical reason. One satisfies norms of practical justification in virtue of practical reason's success in carrying out a guidance function that subserves that telos.

Norms of epistemic justification are teleological in an analogous way. Epistemic reasoning subserves a standard of correctness for belief. That standard of correctness is a telos for epistemic reasoning. One satisfies standards of epistemic justification in virtue of epistemic reason's success in carrying out a guidance function that subserves that telos.

The relevant telos for belief is truth. Truth is a standard of correctness on belief that one satisfies or violates independently of the quality of the reasoning (or the quality of guidance by non-rational capacities) that the belief is based on. Epistemic reason's guidance of belief functions to guide one to belief that satisfies this standard of correctness.¹²

entitlement to certain types of intentions and actions, where our warrant for such intentions and actions need not depend on any guidance by exercises of the agent's action guiding capacities. It seems to me plausible that we have a default entitlement to examine objects that attract our attention, to perform basic physical actions like walking and manipulating objects, to talk with others, and to deliberate. If we have a default entitlement to perform these actions, and to corresponding intentions, it is because of the essential role of those types of actions in agency, and the development of agency.

¹²Some claim that the "aim" of belief, and of epistemic reasoning, is knowledge, and not truth. Advocates for this view are sometimes misled by the fact that knowledge is a fuller success for belief, and for epistemic reasoning, than is truth. To see the error, it may help to take up Sosa's (2002) analogy between epistemic success and success in archery. For an archer, hitting the target because of the skillfulness of the shot is a fuller success than hitting the target through luck. Nonetheless, the "aim" of the archer's shot is to hit the target, and not to hit the target through an exercise of one's skill. In the same way, the "aim" of belief

As suggested above, the relevant telos for intention and action, their standard of correctness, is roughly that the intended action be good. Somewhat more specifically, my view is that the relevant telos for action and intention is that intended action helps, on the whole, to realize the agent's legitimate ends. Legitimate ends include ends like nutrition, being well-rested, the avoidance of illness and bodily harm, and physical fitness. Legitimate ends also include ends like scientific or artistic accomplishment, physical or intellectual pleasure, and companionship. They also include moral ends, like the fair distribution of goods, telling the truth, helping those in need, treating others kindly, and so on. I doubt that all legitimate ends are moral ends. Pre-moral practical reasoners, like young children and higher animals, can have legitimate ends. Non-rational animals, like mice, also have legitimate ends. Reproduction, survival, eating, and the avoidance of bodily harm and pain (for pre-rational animals that feel pain), are legitimate ends for such animals. Legitimate ends can also include the fulfillment of relatively unimportant goals and desires, like making a shot in a casual basketball game.¹³

Legitimate ends contrast with deficient ends, like ends that are harmful, wicked, pointless, or shallow. Ends like these can only be cited in criticizable answers to questions like “why should she ϕ ?” If in answer to the question “why should she drive to the pet store?” one says, “so as to acquire kittens for drowning,” the answer is criticizable because there is something wrong with the end specified in the answer. The end specified is deficient as a ground for the action. Unlike deficient ends, citing a legitimate end in specifying the end of an action is not criticizable in this way. If in answer to the question “why should she go to the doctor?” one says, “so as to be cured of her illness,” in normal circumstances the specification of the

is to get things right, to arrive at true belief, and not to arrive at true belief in a way that constitutes knowledge—even though arriving at true belief in a way that constitutes knowledge is a fuller epistemic success than arriving at true belief through luck.

¹³I assume that there is something good about ends like these which makes them worth pursuing, beyond their not detracting from more important ends like health, companionship and intellectual accomplishment. Compare the discussion of “pointless” ends two paragraphs down.

action's purpose in the answer would not be criticizable as giving deficient grounds for the action.

In my view, for an end to be cited in a non-criticizable answer to such a question, there must be something positive about realizing that end, given the details of the context of action, and the agent's attitudes toward the end. If there is nothing positive about realizing some end, then citing that end would not really specify a purpose, or point to the action. There must be, so to speak, a genuine "point" to an action, on pain of the action's being "pointless." That there is nothing positive about realizing the cited end of an action is a kind of criticism of the explanation of why the action is to be done. A non-criticizable end of an action cannot be bad, but also cannot be evaluatively neutral, in that there is nothing positive about performing the action. It is natural to think there is nothing positive or negative about a light breeze sweeping a little sand on an isolated beach. In my view, a legitimate end cannot be neutral in this way.

Even an action like idly drumming one's fingers can be positive in certain respects, in that it helps release nervous energy, or in that the sound it makes is source of mild enjoyment to the agent, for example. If an action is in no way oriented toward any such positive end, the action is criticizable as arbitrary or pointless, and lacks practical justification (and also lacks other kinds of practical warrant). The claim here is similar to Quinn's (1993) claim that an agent acts incoherently if she acts on a pure orectic impulse, sees nothing good in her action of, for example, turning on a radio. However, the main points in this chapter do not depend on this view about legitimate ends. The account of practical warrant I develop below is consistent with the view that legitimate, non-criticizable ends can be entirely neutral, and need not be positive in any way.

Practical justification is a matter of the agent's rational capacities performing well with respect to their function of guiding the agent to action that helps realize the agent's legitimate ends. Fulfilling standards of practical justification constitutes a practical good, a practical

success, insofar as those standards mark success with respect to that function. In that way, the telos of realizing legitimate ends grounds standards of practical justification, just as the telos of truth grounds standards of epistemic justification.

Standards of practical justification mark successes in the functioning of our capacity to make use of reasons to determine what to do, *given* limitations in one's perspective, and in one's cognitive capacities. Standards of practical justification do not prescribe that we reason like an ideally informed and ideally rational being. Whether an individual's intention satisfies standards of practical justification depends, in part, on what sort of reasoning she is capable of, given limitations in her perspective and her cognitive abilities. Practical reasoning that is in certain respects misguided can satisfy standards of practical justification, if the error is due to limitations in one's perspective or in one's cognitive capacities. Likewise, actions and intentions that are based on such reasoning, and which are likewise misguided, can also satisfy standards of practical justification.

As a result, whether one's intentional action does in fact help to realize one's legitimate ends does not always covary with the practical justification of an action. For example, a thirsty individual might have strong-but-misleading reasons to believe that a particular cup of water is safe to drink, when in fact there is poison in the water. The individual might make good use of reasons in forming an intention to drink, given the limitations in her perspective. Her intention, and her corresponding action of drinking the water, can be justified even though the action falls far short of helping to realize the agent's legitimate ends.¹⁴

¹⁴Limitations in cognitive capacities can bear on practical justification too. Consider two chess players, one advanced player, and one beginner. The disadvantageous move of an advanced chess player may lack practical justification because, with her level of expertise, she could have easily seen that the move would eventually lead to disadvantage. The same move, made by a beginner, on the basis of the same reasons, can be justified because, given her lack of expertise, it would have been very difficult for her to see that the move would lead to disadvantage. The beginner's move is justified because she does a good job making use of reasons to determine what move to make, given limitations in certain cognitive capacities involved in playing chess at a high level.

In the other direction, an intention and action might violate standards of practical justification, even though the intended action helps to realize the agent's legitimate ends. For example, one might take medicine that is in fact curative, even though one is aware of strong-but-misleading evidence suggesting that the medicine is harmful. In that case, one's rational capacities perform poorly in guiding the agent to action, and the intention and action lack practical justification. But the action does help realize the agent's legitimate ends. Whether a belief is true is independent of epistemic justification in these same ways. A belief can be epistemically justified but false, or it can be true while violating standards of epistemic justification. Analogously, actions and intentions can be justified even if the intended action does not help to realize the agent's legitimate ends; and an action and intention can violate standards of practical justification even if the intended action happens, by luck, to help to realize the agent's legitimate ends.

Perception, sensation, and practical entitlement

Practical justification is a type of practical warrant which consists in exercises of one's rational capacities. But exercises of non-rational capacities also contribute constitutively to practical warrant. Actions usually count as warranted, instead of being misguided or arbitrary, partly in virtue of guidance by non-rational capacities like capacities for perception and sensation.¹⁵ The practical warrant provided by these nonrational capacities is a type of practical entitlement, and conflicts with internalist claims about practical warrant.

Perceptual capacities are non-rational capacities I am not aware of anyone who argues that sensations are exercises of rational capacities. But many philosophers do claim

¹⁵Capacities for basic desiderative states like hunger and thirst are also non-rational capacities that can contribute to practical warrant. So are capacities for basic kinds of pains and pleasures. In this section I focus on the role of perception and sensation in warranting intentions and actions. I mean to discuss the warranting role of other non-rational capacities in future work.

that our capacity for perception is a rational capacity. I will not attempt a detailed criticism of that view here.¹⁶ But I do want to briefly note some grounds for claiming that perception is a non-rational capacity. First, animals who lack rational capacities for thought and understanding have perceptual capacities. Bees, shrimp, mice and other animals have perceptual capacities but lack rational capacities.¹⁷ Perceptual psychology finds no difference in kind between our perceptual capacities and those of non-rational animals. The capacities seem to operate in largely the same ways. This point strongly suggests that our perceptual capacities, like those of these other animals, are non-rational capacities.

Second, rational capacities are widely thought to be capacities to operate with representations having sentence-like propositional form. Perceptions do not have propositional form. Their structure is more similar to the structure of a complex noun phrase, than to the structure of a sentence. The word ‘white’ in the phrase ‘the white whale’ is part of a term that representationally picks out the object that is being described as white. In the sentence ‘the whale is white,’ the word ‘white’ is not part of such a term. The structure of perception is more like that of ‘the white whale’ than ‘the whale is white.’ When attributives are applied to objects in perception, the attributive is part of a representation that picks out that object. In perceiving an object as white, for example, the perceptual attribution of whiteness to the object is part of a representation that picks out that object. The fact that perceptions do not have propositional form suggests that perceptions are not exercises of rational capacities.¹⁸

Third, our abilities to conceptually represent observational attributes like *red* and *moving* in thought develop from our abilities to represent such attributes in perception. If perceptions are exercises of rational capacities, then their content is conceptual. If perceptual content

¹⁶Compare, for instance, McDowell (1996), Brewer (1999), Martin (1993), Siegel and Silins (forthcoming).

¹⁷Compare Burge (2003) and (2010a), and Peacocke (2001).

¹⁸See Burge (2010b) and (2010c).

is conceptual, then our perceptual capacities presuppose abilities to conceptually represent observational attributes in thought, and cannot developmentally explain the acquisition of those conceptual capacities.¹⁹

It is sometimes argued that perceptions must have conceptual content, and must be exercises of rational capacities, if they are to provide rational epistemic support beliefs.²⁰ Presumably, “rational epistemic support” is epistemic support provided by exercises of one’s rational capacities to make use of reasons in forming beliefs. *reasons* for one’s belief. The claim that perceptions provide rational epistemic support for beliefs begs the question at issue.

In fact, the relation between perceptions and standards of rationality suggests that perceptions are not exercises of rational capacities, and do not support beliefs as reasons for belief. Perceptions epistemically support beliefs even though they do not satisfy or violate standards of rationality. It is a mistake to evaluate perceptions as rational or irrational, or as intelligent or foolish. If perceptions were exercises of rational capacities, they would be subject to assessments like these.

A related point is that one is never directly culpable for a failure in the exercise of one’s perceptual capacities. Rational capacities are usually conceived of as active capacities, where one can be directly culpable for failures in the exercises of such capacities. This point dovetails with the point about perception and standards of rationality. One is culpable for violations of norms of rationality. The fact that one is not culpable for failures in exercises of perceptual capacities implies that perceptions cannot violate standards of rationality.

The point that perceptions cannot violate rational standards is also borne out by the fact that perceptions cannot be based on reasons. It is a mistake to ask for an individual’s reasons for her perceptions. One cannot perceive that an object is moving on the basis of reasons

¹⁹Compare Ayers (2002), Peacocke (2001).

²⁰See McDowell (1996).

indicating that it is moving, for example. Since perceptions cannot be based on reasons, they are not under our rational control. Since they are not under our rational control, they are not assessable as rational or irrational. However, exercises of rational capacities which purport to represent the world accurately are under our rational control, and are assessable as rational or irrational.

It is also worth noting that perceptions can provide epistemic support for beliefs in a way that reasons for belief cannot. Perceptions can provide epistemic support for beliefs whose content does not differ in a substantive way from the content of the perception. For example, my belief that the object in front of me is moving is based on my perception of that object as moving. That perception cannot epistemically support that belief as a reason for the belief. On pain of circularity, the content of a reason for a belief must differ in a substantive way from the content of the belief it is a reason for. Perceptions do not support such beliefs as reasons for the beliefs. But, if perceptions were exercises of rational capacities, one would expect that they would epistemically support beliefs as reasons for beliefs.

The question of whether perceptions have conceptual content, and are rational capacities, has been well worked over. The remarks here need more detailed development to fully address the arguments of those who claim that perception is a rational capacity. I do not want to focus here on the question of whether capacities for perception and sensation are rational capacities. Having briefly sketched some grounds for doubting that they are, I want to focus on the way that exercises of these capacities help provide practical warrant. It is worth pointing out though, that even if one insists that perceptual and sensory capacities are rational capacities, there is still another type of practical entitlement. Practical warrant for actions and intentions consists partly in the reliability of the agent's action guiding capacities, and their causal connection to action. That reliability and causal connection are not exercises of the agent's rational capacities. I discuss the role of reliability in warranting

action in section 3.²¹ The claim there, about the role of reliability of one's action guiding capacities in warranting intentions and actions, is the most important anti-internalist point I want to make in this chapter.

The role of perception and sensation in warranting intentions and actions One way to see that perception and sensation help to warrant intentions and actions is by noting that our justification for our intentions and actions is usually partly observationally grounded, and not apriori. In contrast to one's justification for believing that $2+2=4$, or for believing that bachelors are unmarried, the justification for intentions usually does not consist solely in support from reasoning or understanding. Usually, our justifying basis for action and intention is not solely apriori reflection about what to do, or recognition of self-evident truths about what to do. For example, my justification for going to the grocery store depends partly on observational evidence that indicates that there is not much food in the house.

When our intentions are observationally warranted, and not justified apriori, they depend for their justification partly on perceptions or sensations. Apriori justification is justification that depends only on guidance by reason, and on facts about the nature of reason.

A belief is epistemically justified apriori when it does not depend for its warrant on perception or sensation, or on the nature of perception or sensation. Similarly, intentions and actions are practically justified apriori when they do not depend for their warrant on perception and sensation, or other 'passive' non-rational states.²²

If intentions and actions did not depend for their warrant on guidance by such non-

²¹Although it is an important issue, I do not address the warranting role of causal connection to action in this chapter.

²²When an intention or action is warranted partly by the guidance of hunger, fatigue, or basic pains, that also renders the practical warrant non-apriori. The role of these non-rational capacities in guiding action is different from the role of perception, and non-orectic sensation like smell. They do not help guide action just by causally registering information relevant to the action. They play motivational roles, and end-setting roles, in the guidance of action. Again, I mean to discuss their role in warranting intention and action in future work.

rational capacities, then the warrant for intentions and actions would always be apriori. Clearly, the justification of intentions and actions is not always apriori. Guidance by non-rational capacities can contribute constitutively to warrant for intentions and actions.

The claim that perception and sensation help warrant intentions and actions also finds support in the relation between the practical warrant of intentions and actions, and the epistemic warrant of beliefs that they are based on.

Practical warrant for actions and intentions depends on epistemic warrant for beliefs that they are based on. If an individual acts or intends to act on the basis of an epistemically unwarranted belief, that lack of epistemic warrant undermines the practical warrant for the action or intention. For example, if an individual has an epistemically groundless belief that her favorite team is sure to win, and places a bet on the basis of that belief, then her action and intention to place the bet are practically unwarranted. Actions and intentions are practically warranted partly in virtue of the epistemic warrant of the beliefs that they are based on.

Beliefs that intentions and actions are based on often depend for their epistemic warrant partly on perception and sensation. Intentions and actions can depend for their practical warrant partly on the epistemic warrant of such beliefs. So intentions and actions can depend for their practical warrant partly on perception and sensation.

One might deny that practical norms can include epistemic norms in the way I am suggesting. Parfit (2011) claims that the practical rationality of actions cannot depend on the epistemic rationality of beliefs that actions are based on. He claims, for example, that smoking can be rational even if the action is based on a wildly irrational belief that smoking is healthy. Parfit has a dichotomous thinks in a dichotomous way about practical success and epistemic success. Epistemic success and practical success are different kinds of success, and the one has little to do with the other.

This dichotomous view is mistaken. As Parfit himself claims, whether an action is rational

is roughly a matter of whether it is stupid, foolish, or crazy. But actions can count as stupid, foolish, or crazy in virtue of epistemic mistakes one makes in holding beliefs that actions are based on. If one smokes on the basis of a wildly irrational belief that smoking is healthy, the action is foolish and practically irrational. And the action counts as foolish because of the egregious epistemic failure of believing that smoking is healthy. In the other direction, actions can count as clever, intelligent, or even skillful in virtue of epistemic successes one achieves in arriving at beliefs that the action is based on. For example, a chess move can count as particularly clever in virtue of an epistemic success of noticing some subtle vulnerability in an opponent's position.

Norms that apply to actions and intentions are practical norms. Thus, negative and positive evaluative concepts, like 'foolish' and 'clever' apply to actions and intentions in virtue of the satisfaction and violation of practical norms. Those practical norms can be satisfied and violated in partly in virtue of epistemic successes and failures. Fulfilling practical norms is, roughly, a matter of doing well in determining what to do. Fulfilling epistemic norms is, again roughly, a matter of doing well in determining what is the case. But our ability to determine what to do consists partly in our ability to determine what is the case. Doing well in determining what is the case is an important part of doing well in determining what to do.

We would be quite helpless in determining what to do if we lacked epistemic capacities to discern matters of fact, like whether there are dangerous animals nearby, whether the room is on fire, or whether smoking is healthy. Our reliance on our belief-guiding epistemic capacities in determining what to do shows the practical significance of our epistemic capacities. Epistemic capacities to determine what is the case have an epistemic function of guiding one to true belief. But our capacities to determine what is the case also have a practical function. They function to help guide our actions. Determining what is the case is an important part of determining what to do. As a result, epistemic success in determining what is the case

can be a part of practical success in determining what to do.

The practical function of our action guiding capacities is to guide us to action that will help realize our legitimate ends. Whether an action will help realize the agent's legitimate ends depends on facts relevant to the action. For example, whether smoking will help, on the whole, to realize the agent's legitimate ends, depends on whether smoking is healthy. The epistemic function of our belief guiding capacities is to guide us to belief that accurately represent such facts. Success with respect to the practical function of guiding action depends on success with respect to this epistemic function. As a result, satisfaction of norms of practical warrant can depend on the satisfaction of norms of epistemic warrant.

An intention or action based on an epistemically groundless belief thereby lacks practical warrant. Perception and sensation can contribute constitutively to practical warrant by providing epistemic grounds for beliefs that actions and intentions are based on. If a belief is epistemically warranted by a perception or sensation, an intention or action based on that belief can be warranted partly in virtue of that perception or sensation. Consider an individual who crosses the street on the basis of a groundless belief that the light is green. That action lacks practical warrant. If instead of being epistemically ungrounded, the belief were based on a perception of the light as green, the action would be warranted. The action is warranted partly because of the perception's role in epistemically grounding the belief that the action is based on.

I suggested earlier that practical reason functions to indicate whether an intention and an intended action satisfy their standards of correctness. Satisfying these standard of correctness is a matter of whether the intended action helps, on the whole, to realize one's legitimate ends. In practical reasoning one makes use of reasons that function to indicate that one's action satisfies that standard of correctness. Making use of reasons in this way is a core intention guiding, and action guiding function of practical reason. Standards of practical justification mark successes in the exercise of practical reason with respect to these functions.

Like practical reason, non-rational capacities like perception also have a practical function of indicating whether an intention and intended action satisfies its standard of correctness. These capacities are not involved in making use of reasons to establish that an action or intention satisfies its standard of correctness. Non-rational capacities are not capacities to operate with reasons. But they have a similar function. These non-rational capacities function to register information. In practical reasoners like us, registrations of information in perception and sensation can function to help indicate whether an intended action satisfies its standard of correctness. Thus, like practical reason, these non-rational capacities also have an intention guiding and action guiding function. When these non-rational capacities contribute to practical warrant, they do so in virtue of successes with respect to this practical guidance function.

When non-rational capacities contribute to practical warrant, they do so by contributing to practical entitlement, and not to practical justification. Standards of practical justification mark successes in the functioning of one's rational capacities. Those standards are not fulfilled in virtue of operations of non-rational capacities. As suggested above, capacities for perception and sensation are non-rational capacities.

The fact that these non-rational action guiding capacities help warrant intentions and actions conflicts with some internalist claims about practical warrant. One such internalist claim is that whatever helps to warrant intentions and actions is accessible for use as an element in reasoning, or in conscious reasoning. But our capacity to reason is a rational capacity. Accordingly, elements of reasoning are exercises of rational capacities. Likewise, elements of reasoning are open to assessment as rational or irrational. Perceptions and sensations are not open to that kind of assessment, and are not exercises of rational capacities. Thus, they are not available for use as elements in reasoning, including conscious reasoning.

Another internalist claim about practical warrant is that we are responsible for whatever helps to warrant intentions and actions. But we are not responsible for perceptions and

sensations, even though they help to warrant our intentions and actions. If one's perceptual capacities or sensory capacities fail to function properly—if one is myopic, for example—one is not responsible for that failure. Nor is one responsible for successes of those capacities when they function well. I get no more credit for my accurate perceptual discrimination of the color red than I do for my digestive system's effective functioning.

Perception and sensation are non-rational capacities. As a result, they do not practically justify intentions and actions. Only exercises of rational capacities justify intentions and actions. However, the fulfillment of norms of practical justification does depend on perception and sensation in another noteworthy way.

I want to introduce the point by focusing first on an analogous point about epistemic warrant. Consider an epistemically warranted perceptual belief that an object is white. The perception of the object as white does not justify the belief, since justification is due to guidance by one's rational capacities. The warrant provided by the perception is a kind of entitlement, not justification. But now "subtract" the perception from the formation of the belief, so that it is formed without being based on perception, or on any other mental state. Such a groundless belief would violate norms of epistemic justification, and epistemic rationality. It would be an irrational, unjustified belief. It would seem, then, that whether a belief violates standards of epistemic rationality can depend on perception.

This point might seem somewhat puzzling. Rational norms, and norms of justification, apply non-derivatively only to the functioning of one's rational capacities. Perception is a non-rational capacity. If so, it seems that standards of rationality and justification cannot be satisfied or violated in virtue of perception. How, then, can epistemic justification and rationality depend on perception?

The irrationality in the second case is due, in the first instance, to a failure in the exercise of one's rational capacities. Forming the belief constitutes a mistake in exercise of one's rational capacities, since the belief is groundless. If the belief is based on perception,

as in the first case, one's rational capacities perform well in forming the belief. Forming the belief does not constitute a mistake in the functioning of one's rational capacities in that case, so the belief does not violate standards of rationality or justification.

What the individual's rational capacities do is in important respects the same way in both cases. In both cases, one's rational capacities form the relevant belief, and in neither case does one recognize supporting considerations through the exercise of reason.²³ In one case forming the belief without rational grounds violates standards of rationality and justification, and in the other case doing so does not violate those standards. Whether doing so violates those standards depends on exercises of one's non-rational perceptual capacities.

Standards of rationality and justification are non-derivatively fulfilled only by successes and failures in the exercise of one's rational capacities. However, whether an exercise of one's rational capacities constitutes a rational failure can depend partly on exercises of non-rational belief-guiding capacities like perception. Whether a belief violates standards of rationality depends in this way on exercises of non-rational capacities.

Reliability

I have been discussing the contribution to practical entitlement made by non-rational action guiding capacities. I want to turn now to a second type of practical entitlement.

I suggested earlier that practical reason functions to guide one to action that helps realize one's legitimate ends. Non-rational capacities like perception have that action guiding function as well. Practical reason and non-rational capacities contribute to practical warrant

²³The content of the beliefs will differ in certain respects, since perceptual beliefs inherit their contents from the associated perceptions, and the perceptions have demonstrative elements. (Compare Burge (2009a).) The point here does not depend on the exercises of rational capacities being the same in every respect. The point depends on the observation that in both cases one's rational capacities form the belief, but do not "recognize" grounds for the belief. In one case, that constitutes an error in the functioning of one's rational capacities, and in the other case it does not.

in virtue of their success with respect to this function. Action guiding capacities contribute to practical warrant partly in virtue of their reliable effectiveness with respect to this function. This reliability is a second kind of constitutive contributor to practical entitlement for intentions and actions.

There is an analogous source of epistemic entitlement. Belief guiding capacities contribute to epistemic warrant partly in virtue of their being good routes to true belief. Being a good route to truth requires that the capacities reliably perform in a way that helps to lead one to true belief.

There are other epistemic standards that do not require that one's belief guiding capacities be reliable in this way. One satisfies certain epistemic standards in doing as epistemically well as one can, even if one's belief-guiding capacities are utterly ineffective. But those are not standards of epistemic warrant. Epistemic warrant is a matter of the successful exercise of epistemic competence. It is not attained by the least bad exercise of epistemic incompetence.

Standards of practical warrant are associated with being a good route to action that helps realize legitimate ends, in the same way that standards of epistemic warrant are associated with being a good route to true belief. Satisfying standards of practical warrant is a matter of the successful exercise of a practical competence in determining how one acts. The reliable effectiveness of one's action guiding capacities is essential to their being aspects of practical competence, rather than incompetence. As a result, one's actions and intentions are warranted partly in virtue of the reliable effectiveness of one's action guiding capacities in guiding the agent to action that realizes legitimate ends.

This point applies both to practical reason, and to non-rational action guiding capacities. Let me start by explaining how the point applies to perception, as an example of a non-rational action guiding capacity.²⁴

²⁴The point also applies to non-representational sensory representations, and to basic desiderative states

A warranted action that is based on a perceptual belief is warranted partly because of the reliable role that perception plays in the reliably effective guidance of one's actions. Many of our perceptual capacities are reliably accurate. We depend on the accuracy of perception in determining what we do. We are able to determine how to act in ways that help to realize our legitimate ends partly because of the accuracy of perception. For example, an endangered pedestrian would not know to step out of the road if not for the accuracy of her perception of oncoming traffic. The reliable accuracy of perception is part of what makes our overall action guiding system reliably effective in guiding us to action that realizes our legitimate ends. If perception were wildly and systematically inaccurate, we would succeed in acting in ways that realize our legitimate ends only rarely, and by luck. We would be incompetent agents.

The reliable accuracy of perception helps to make one's actions and intentions count as guided by the successful performance of reliably effective action guiding capacities. Since practical warrant depends constitutively on one's intentions and actions being guided by the successful performance of such capacities, the reliability of those capacities contributes constitutively to practical warrant. Perception's reliability helps to make our actions count as guided by the successful performance of an over-all *competence* in determining what we do. In this way, practical warrant is due in part to the reliable accuracy of our perceptual capacities. The reliable accuracy of our perceptual capacities is not an exercise of our rational action guiding capacities. Thus, it contributes to the practical warrant of an action or intention as a constitutive contributor to practical entitlement, and not practical justification.

The reliable accuracy of perception contributes constitutively to both practical warrant and epistemic warrant. However, it is worth emphasizing that the reliable accuracy of perception contributes to practical warrant because of the *practical* significance of that reliable

that do not have conceptual content, like hunger and thirst. I focus on perception here in order to exploit parallels with its contribution to epistemic warrant.

accuracy. The reliable accuracy of perception contributes to practical warrant because that accuracy helps to make one's action guiding capacities effective with respect to their practical function. The reliable accuracy of perception helps to make one's action guiding capacities effective in guiding one to action that realizes one's legitimate ends. That feature of perception's reliable accuracy is not relevant to epistemic warrant. Perception's reliable accuracy contributes to epistemic warrant because of the way that accuracy helps guide one to true belief, and not because of its connection to action.

Practical warrant can depend constitutively on the reliability of practical reason, just as it can depend constitutively on the reliability of non-rational action guiding capacities like perception. If a warranted action or intention is based on practical reasoning, the action or intention has practical warrant partly in virtue of the reliability of practical reason in guiding the agent to action that realizes the agent's legitimate ends.

Some aspects of practical reason contribute to practical warrant because of the practical significance of their epistemic reliability. As suggested earlier, epistemic warrant can contribute constitutively to practical warrant. Accordingly, if a warranted intention is based on a belief, and that belief in turn is based on epistemic reasoning, that reasoning contributes to the practical warrant of the intention. Such epistemic reasoning is an element in the agent's practical reasoning, even though the same pattern of reasoning could occur without being an element in any practical reasoning.

Epistemic reasoning contributes to practical warrant partly in virtue its being a reliable guide to true belief. This reliability helps to make one's over-all action system an effective guide to action that realizes legitimate ends. We are frequently successful in settling on action that realizes legitimate ends in part because of the reliability of epistemic reasoning in guiding us to true beliefs that are relevant to the action. Thus, practical warrant can be due in part to this epistemic reliability of epistemic reason.

Some elements of practical reasoning are not elements or instances of epistemic reason-

ing. Transitions in reasoning that have intentions as conclusions are elements of practical reasoning, but not elements of epistemic reasoning.²⁵ The rational capacities involved in those transitions do not function to guide one to true belief. Thus, they do not contribute to the effectiveness of one's action guiding capacities as reliable guides to true belief. They contribute to the effectiveness of one's action guiding capacities as reliable guides to correct intention. That is, they contribute to the effectiveness of one's action guiding capacities as reliable guides to intention that represents action that helps to realize one's legitimate ends.

This reliability is not epistemic, since it does not concern the guidance of belief. But it is similar to epistemic reliability. It is a kind of reliability in making use of reasons to guide mental states (specifically intentions) so that they “accord” with certain facts—facts about whether the intended action helps to realize one's legitimate ends. In non-epistemic reasoning one does not “get the facts right” through arriving at a belief that accurately represents the facts. One “gets the facts right” through arriving at an intention, a commitment to act, which counts as errant or correct in virtue of whether the intended action in fact helps realize one's legitimate ends. Intentions and actions based on practical reasoning are warranted partly in virtue of the reliability of practical reason in getting things right in this way.

This non-epistemic reliability of non-epistemic practical reasoning contributes constitutively to practical warrant. So does the epistemic reliability of epistemic aspects of practical reasoning. They both help make it the case that our actions and intentions are guided by successful performances of a genuine competence in determining what we do. These types of reliability of one's capacities for practical reasoning are not exercise of capacities for prac-

²⁵Some deny that forming an intention is an element in practical reasoning. Joseph Raz and Robert Audi, for example, claim that the concluding element of practical reasoning is a judgment about what one ought to do. The view that intentions are not elements of practical reasoning does not imply that practical reasoning is a special case of epistemic reasoning. Even if intentions are not elements in practical reasoning, transitions to desires with conceptual content can be elements in practical reasoning. These would be non-epistemic elements of practical reasoning. Transitions from desires with conceptual content to judgments about what one ought to could also be elements in practical reasoning, even if intentions are not elements in practical reasoning. Such a transition would not be a transition in epistemic reasoning.

tical reasoning. As a result, the warrant provided by these types of reliability is practical entitlement, and not practical justification.²⁶

Entitlement provided by reliability figures in apriori practical warrant. Even if a warranted intention is guided just by practical reason, and not by non-rational action guiding capacities, the intention is not warranted exclusively by exercises of one's action-guiding rational capacities. When an intention is warranted by exercises of practical reason, the reliability of practical reason in carrying out its action guiding function contributes constitutively to the intention's warrant. The kind of warrant provided by that reliability is an entitlement, not a justification, since the reliability is not an exercise of one's rational capacities. Thus, even an apriori warranted intention would have an element of entitlement, and would not consist exclusively in exercises of rational capacities.

The practical warrant provided by the reliability of our action guiding capacities, including both practical reason and non-rational action guiding capacities, conflicts with internalist claims about practical warrant. First, one is not directly responsible for whether one's rational and non-rational action guiding capacities are reliable. Whether one's action guiding capacities are reliable is not under one's direct control.²⁷ The reliability of these capacities nonetheless contributes constitutively to practical warrant.

Second, the reliability is not an exercise of our capacity to make use of reasons (nor is it an exercise of non-rational action guiding capacities). It is a characteristic of our action guiding capacities, and not an exercise of those capacities. Accordingly, the reliability of these capacities is not accessible for use in reasoning. The reliability can be represented in beliefs that are elements in reasoning.²⁸ But reliability itself cannot be an element in one's

²⁶The reliability of capacities that operate "downstream" of intention, like capacities for motor control, also contribute to practical warrant. I discuss the capacities, and their reliability, in section 6.

²⁷We can, however, sometimes affect their reliability indirectly, through training or medical intervention, for example.

²⁸However, representing reliability in belief requires that one possess the concepts of reliability, and perception or reason. An individual need not possess these concepts to be capable of warranted action. As

reasoning, any more than objects in the external world can be.

Third, the reliability of these action guiding capacities is not accessible to consciousness, except, in some cases, as something represented by conscious mental states. Reliability is not a mental state, and so not a conscious mental state. And, since reliability is not an element of reasoning of any kind, it is not an element of conscious reasoning.

Post-intention action guiding capacities

In this final section, I want to discuss the connection between practical warrant and non-rational action guiding capacities whose operation is posterior to intention. Some non-rational action guiding capacities carry out their action guiding function by guiding our intentions. The capacities for practical reasoning we make use of in forming our intentions guide actions in this way, as do capacities for perception and sensation that help ground our intentions. For example, reasoning about whether the car is in good condition that helps guide one to buy the car does so by helping to guide one to intend to buy the car. Perceptual registration of the sound of the engine can also help to guide the agent to buy the car by helping to guide her intention to buy the car. Non-rational action guiding capacities can also operate “downstream” of intention. Capacities for motor control are an example. If one intends to grasp an object, success in doing so depends on exercises of non-rational motor capacities that determine details of how one moves one’s body in carrying out the action. Many of these details are not represented in intentions with conceptual content. The details are determined by motor capacities that operate downstream of intention. Such non-rational action guiding capacities are, as it were, intermediaries between intention and action.

I want to briefly note some grounds for the claim the action guiding capacities we make

a result, reliability of one’s action guiding capacities can contribute to practical warrant even if the agent cannot represent reliability in belief.

use of to execute our intentions are non-rational capacities.²⁹ First the details of physical action, and the psychological processes that determine them, are often not accessible to conscious reflection by the agent. Second, capacities that determine these details often have to be acquired through practice, and not through an intellectual grasp of the details of skillful action. Third, mistakes in the details of attempts at skilled action often do not open the agent to rational criticism. Mistakes in the details of skillful action often open the agent to criticism as clumsy or inept, without opening her to criticism as irrational, unreasonable, foolish, or stupid. Fourth, intentions are conclusions of practical reasoning. Thus, any action guiding capacities that operate posterior to intention operate posterior to one's reasoning. Rational capacities involved in guiding action are exercised in practical reasoning. Since post-intention exercises of action guiding capacities are posterior to practical reasoning, they are not rational capacities.³⁰

The claim finds further support in the distinction drawn by psychologists between procedural knowledge and declarative knowledge. Psychologists claim that the procedural knowledge that we make use of in carrying out our intended actions consists largely in information that is inaccessible to consciousness, that is available for use only in a circumscribed range of contexts, and which is not encoded in truth-assessable representations.

Exercises of these non-rational action guiding capacities, which operate downstream of intention, do not help to warrant intentions. Exercises of action guiding capacities that help warrant an intention must play a role in guiding the intention.³¹ However, exercises of action

²⁹I discuss this claim in detail in "Nonrational aspects of agency."

³⁰By "practical reasoning" I have in mind exercises of one's ability to make use of reasons to guide commitments about what to do. I do not assume that exercises of this ability must be temporally extended processes that unfold in steps.

³¹Let me underscore that the claim here is that *exercises* of such action guiding capacities do not help to warrant intentions. In my view, the *reliability* of action guiding capacities that operate downstream of intention can play a constitutive role in warranting our intentions. Warrant for an intention depends on whether one is warranted in relying on one's action system to carry out the intended action. If one's action system is systematically ineffective in carrying out the agent's intended actions, one is not warranted in relying on the action system to carry out intended actions. Accordingly, one's intentions also lack warrant

guiding capacities that operate downstream of intention can help warrant actions. Action guiding capacities that operate downstream of intention function to guide us to action that helps realize our legitimate ends. These capacities carry out this function by enabling us to successfully perform the actions we intend.

If we manage to successfully act in ways that realize legitimate ends while lacking such capacities, or while failing in the exercise of such capacities, the success in action is due to luck, and accidental with respect to the normal proper functioning of our action guiding capacities. For example, suppose an individual playing tennis intends to hit a hard winner, and swings her racket in a clumsy way, but happens to hit a hard winner because of the way the ball glances off the side of the racket. Her success in action is due to luck, and is accidental with respect to the normal proper functioning of her action guiding capacities. In cases like this, one's action is successful by luck in a way that is similar to epistemically unwarranted beliefs that happen to be true by luck. The agent succeeds in doing what she intends even though her action guiding capacities perform poorly in determining how she acts.

Whether an action is warranted depends on whether it is well guided by one's action guiding capacities. Whether an action is well guided consists partly in the exercise of non-rational capacities that operate downstream of intention, like capacities for motor control. As a result, warrant for actions consists partly in the exercise of these kinds of capacities.

Since these are non-rational capacities, the practical warrant they provide is a kind of entitlement, and not a kind of justification. Practical reasoning is an exercise of one's rational capacities. So exercises of non-rational capacities are not elements in practical reasoning. Practical justification only consists in practical reasoning, and its elements. Accordingly, the

in that case. In this way, warrant for an intention depends partly on whether aspects of one's action system that operate downstream of intention are reliable or systematically ineffective in carrying out one's intended actions. I regard this claim as more tenuous than other claims about practical warrant in this chapter. As a result, the claim requires more detailed discussion and defense, which I leave for another occasion.

warrant provided by exercises of post-intention action guiding capacities conflicts with the internalist claim that whatever contributes to practical warrant must be accessible for use in reasoning, or in conscious reasoning.

I do not claim, however, that we cannot be responsible for exercises of these non-rational post-intention action guiding capacities, at least given a sufficiently broad notion of responsibility. One can be at fault for an error in performing a physical action when the error is due to failures in these post-intention action guiding capacities, and when there is no failure in the exercise of the agent's intention guiding capacities. For example, Bill Buckner is famously at fault for his failure to handle a routine ground ball that could have won the 1986 world series for the Boston Red Sox. (Though of course he is not morally at fault.) He is at fault for his error even though the error was due to failures of non-rational post-intention action guiding capacities involved in carrying out his intended action of collecting the ground ball.³² It seems to me that one can be culpable for a very wide range of failures in the performance of intentional actions, and a very wide range of post-intention action guiding capacities can constitute the locus of those failures. One can be at fault for losing one's balance, for stuttering, for failing to find something in scanning a scene, and so on—provided that success it is within the agent's abilities, and she is not helpless with respect to her failure.

But there are some non-rational post-intention action guiding capacities for which one cannot be culpable or creditable. Perception and sensation have a post-intention action guiding role, in addition to their role in guiding intention. For example, the post-intention guidance of details of how one moves one's hand in grasping a cup are keyed to perception of the cup and one's hand, proprioceptive sensory input concerning the position of one's hand, and sensory information concerning the degree and direction of pressure of the cup on one's

³²I discuss fault for errors of this kind in more detail in "Non-rational aspects of agency." One might wonder why we should be culpable for failures in these capacities, but not culpable for failures of perception or sensation. In my view, the reason is that these post-intention action guiding capacities are part of our capacity for agential control, and are active in a way that perception and sensation are not. I set aside that issue here.

fingers. We are not culpable for failures in the post-intention operation of these perceptual and sensory capacities.

The role of post-intention non-rational action guiding capacities in warranting action is quite different from the role of non-rational capacities that guide action by guiding intention. This difference is reflected in the fact that exercises of the latter capacities bear on standards of justification in a way that exercises of the former capacities do not. I claimed earlier that whether an action violates norms of practical rationality and practical justification can depend on exercises of non-rational capacities, even though such norms apply non-derivatively only to exercises of rational capacities. That claim only applies to non-rational capacities that guide action by guiding intention. Exercises of capacities that operate downstream of intention do not bear on the violation of standards of practical irrationality and practical justification.³³

If one's intention is well guided by one's rational and non-rational capacities, an action based on the intention does not violate norms of rationality or justification. Success or failure in the functioning of post-intention action guiding capacities, like capacities for motor control, are irrelevant to rationality and justification. For example, suppose an individual intends to throw a ball to her friend, and has good evidence that she can throw the ball accurately. Suppose that she intends to throw the ball on the basis of good reasons, and good guidance from non-rational capacities like perception. But suppose she flubs the throw through a failure in her motor capacities, and breaks a window. If there really is no error

³³However, in my view, the reliability of capacities that operate downstream of intention can contribute constitutively to the satisfaction of standards of practical rationality and practical justification—though the reliability is not an element in a practical justification, since reliability is not an exercise of one's rational capacities. This claim is closely related to the claim that an intention's warrant depends partly on warrant for relying on one's action system to carry out the intended action. (See footnote 32 above.) If an intention is unwarranted, then it lacks practical justification. As a result, practical justification depends partly on warrant for relying on the action system to carry out the intended action. That warrant for relying on the action system depends on the reliability of action guiding capacities that operate downstream of intention. As a result, so does the practical justification of an intention. One's practical reasoning is in a certain respect defective if it issues in an intention that one cannot rely on oneself to carry out—even if there is no mistake endogenous to the practical reasoning itself.

in the guidance of the intention that her errant action is based on, then she is not open to charges of irrationality, or charges of acting without adequate justification. She is open to charges of clumsiness, or ineptness, but not charges of rational failure.

Exercises of post-intention action guiding capacities help warrant actions even though they do not affect practical justification. As a result, a justified action can violate norms of practical warrant because of failures in the exercise of post-intention action guiding capacities. Practical justification is independent of practical entitlement provided by post-intention action guiding capacities.

In addition to exercises of post-intention action guiding capacities, the reliability of such capacities also helps to warrant our actions.³⁴ Warrant for our actions depends on whether our action guiding capacities constitute a practical competence, and are a good route to action that helps to realize our legitimate ends. Whether our action guiding capacities are a good route to action depends on action guiding capacities that operate downstream of intention. Our action guiding capacities constitute a good route to action partly in virtue of the reliable effectiveness of these post-intention capacities in enabling us to perform the actions we intend. Without the reliable effectiveness of capacities like these, if we managed to act in ways that realize legitimate ends, it would be due to luck, and accidental with respect to the normal proper functioning of our agency. Thus, the reliable effectiveness of those actions contributes constitutively to warrant for our actions.

The reliability of action guiding capacities that operate downstream of intention contributes constitutively to practical entitlement, and is not a practical justification, or an element in a practical justification. Practical justification consists in the agent's practical reasoning, and elements of the reasoning. Like the reliability of practical reason, and the reliability of non-rational intention guiding capacities, the reliability of these post-intention

³⁴As described in footnote 32, I think the reliability of these capacities also plays a constitutive role in warranting our intentions. But I do not attempt to defend that claim here.

action guiding capacities is not an element of reasoning. Reliability can be represented in elements of reasoning, like beliefs, but cannot itself be an element of reasoning. Since the reliability of post-intention action guiding capacities is not a practical justification, or an element of a practical justification, the kind of warrant it provides is a practical entitlement.

Again like the reliability of other action guiding capacities, the warrant provided by the reliability of post-intention action guiding capacities conflicts with internalist claims about practical warrant. The reliability is not available for use in reasoning, or in conscious reasoning. Nor is the reliability a conscious mental state. Moreover, reliability of these capacities can contribute to practical warrant even if one lacks the ability to represent reliability. Furthermore, one is not directly creditable or culpable for the reliability or unreliability of these action guiding capacities. We do not have direct control over the reliability of these action guiding capacities.

Chapter 7

Entitlement for relying on physical abilities

Introduction

Many philosophers are drawn to Anscombe's claim that we can have non-observational knowledge of action. But action, at least physical action, depends on what we are actually doing with our physical bodies. One might object to Anscombe's claim on the grounds that we cannot have knowledge concerning the vicissitudes of the physical world—including our physical bodies—independently of observational evidence. In this chapter I present an objection to Anscombe's claim along these lines, and defend a variant of her claim in a way that is responsive to the objection.

According to the claim that we can have non-observational knowledge of our actions, the warrant¹ associated with knowledge of action need not depend at all for its force on obser-

¹It is more common to use the term 'justification' than 'warrant' in discussions of knowledge of action. Following Burge (2003), I use the word "warrant" to avoid internalist connotations of the word "justification." That is, I avoid the word 'justification' here because it is often thought that justification must consist in rational support, or in fulfilling some kind of duty. Very roughly, I take warrant to be a matter of a kind of doing well in determining what is the case, or in determining what to do. I do not assume that the relevant

vational evidentiary grounds like perceptions, sensations, and proprioceptive registrations.²

Many discussions of the role of observation in knowledge of action focus on the question of whether knowledge of action is based on observation of our actions as they unfold. That question is different from the question of whether knowledge of action is observational *at all*. Knowing my wife's daily routine, I have observational knowledge that she is working right now, even though that knowledge is not based on perception of her action as it unfolds.³ Some deny that knowledge of one's own physical actions must depend on observation of the action itself, but still claim that knowledge of action has to depend on observational grounds indicating that one is able to perform the action, where the grounds are given by past experience.⁴ The view that knowledge of action is non-observational is the view that knowledge of action depends on *no* observational evidentiary grounds, including those given by past experience.

There is a reasonable *prima facie* case to be made against the claim that we can have such knowledge of our physical actions. Whether you are successfully carrying out a physical action, like extending your arm, or juggling, or wiggling your ears, is a contingent, extra-mental matter of fact. It depends on what body parts you have, on whether you have the relevant motor skills, whether you are tied up or not, and so on. Except in special cases, one does not have apriori knowledge of contingent, extra-mental matters.⁵ Knowledge

sort of "doing well" must consist in having good rational support, or in fulfilling duties. The usage of the term 'warrant' here differs from Plantinga's (1993). According to Plantinga's usage, epistemic warrant is whatever makes the difference between knowledge and true belief. There is no epistemic warrant in Gettier cases, according to Plantinga's usage of the term, even though there is a kind of doing well in determining what is the case. I take epistemic warrant to consist in that kind of doing well in determining what is the case. I take practical warrant to consist in an analogous kind of doing well in determining what to do.

²It goes against ordinary English usage to call sensations proprioceptive registrations "observations." I call these mental states "observational" because they register, or function to register, information about the world through causal interaction with the world.

³O'Brien (2003, p. 8.) marks this difference by claiming that knowledge of action is only "relatively" apriori.

⁴Compare Paul (2009, p. 15), who follows Grice (1971, p. 276).

⁵Exceptions include non-observational knowledge expressed by the sentence "I am here now" (see Kaplan

of whether it is raining or sunny, for example, must be based on observational grounds. Knowledge about such matters depends on observational registrations of the world, which provide information about what is going on in the world. It is not clear why we should have license to believe we are successfully acting as intended without grounding by observational information that indicates what is actually going on with our physical bodies.

Indeed, capacities for observational registration of what the body is doing, like perception and proprioception, play a very large role in facilitating action, and in guiding our beliefs about our actions. At least in most cases, beliefs about what one is doing are formed and sustained under the guidance of proprioception and perception of one's bodily movements,⁶ or past experience indicating that one is likely to be successful in carrying out an intended action. For example, one would not normally believe that one is successfully carrying out an intention to wiggle one's ears without observational grounds, like past perception or sensory registration of doing so. It would be surprising if such observational capacities did not play an important role in providing for our knowledge of our physical actions. In this chapter I defend a variant of Anscombe's view in a way that aims to address these concerns.

The variant I defend differs from the view most often associated with Anscombe's claim. Most defenses of Anscombe's view defend a variant that claims we have direct, non-observational epistemic warrant for *beliefs* about action. I call this a "variant" because Anscombe's own view seems not to be about epistemic warrant for beliefs. She denies that our non-observational knowledge of action ("practical knowledge") is a kind of "contemplative knowledge" or "speculative knowledge" involved in action. Mocking that view, she compares it to thinking there is "a very queer and special sort of seeing eye in the middle of the acting" (1965 p. 57). Moreover, she says of contemplative knowledge that it is "judged as such by

(1977)). And perhaps the "baptizer" of the name 'Neptune' can have non-observational knowledge that, if Neptune exists, it perturbs the orbit of Uranus (see Kripke (1980)).

⁶See Pickard (2004) for an account of knowledge of action that emphasizes the role of observational awareness of the body as one acts.

being in accordance with the facts,” and implies that practical knowledge does not purport to accord with any facts. It is not easy to see what sort of thing Anscombe has in mind when she talks about practical knowledge of what one is doing. But it seems unlikely that she has in mind epistemic knowledge, and epistemic warrant, since she does not have in mind a success in getting one’s beliefs to accord with facts.⁷

Like Anscombe’s own view, the view I defend does not concern epistemic warrant for beliefs about action. The Anscombe-inspired view that I defend concerns warrant for our *reliance* on physical abilities in intentional action. Normally, when one intends to be performing some physical action, one relies on the physical abilities through which one executes one’s intentions.⁸ When one grasps a cup, or plays Bach’s first cello suite, one relies on exercises of the physical abilities involved in performing those actions.

In this chapter, I defend the view that, in core cases, when one performs an action, one has a default entitlement to rely on exercises of one’s physical abilities to play their role in carrying out the action. The entitlement applies by default, and does not depend on whether and how the agent makes use of any kind of evidence concerning what physical abilities one possesses, and whether they are functioning properly in the particular case. Thus, in core cases, one has a warrant for relying on exercises of physical abilities that does not depend for its force on any evidence that the abilities are effective and working properly. So, for example, warrant for relying on exercises of one’s ability to walk need not depend for its force on evidence that one is able to walk, or that one is successfully making use of one’s

⁷Another interpretation of Anscombe is that, for her, practical knowledge can be regarded as a special case of epistemic knowledge, which is the cause of what it understands. One worry about this interpretation is that a piece of putative knowledge can be judged according to whether it accords with the facts, even if the knowledge is the cause of what it understands. For example, Descartes *cogito* is supposed to count as knowledge partly because it accords with the facts, even though the relevant fact does not obtain independently of the knowledge itself. The aim of this chapter is not Anscombe interpretation, so these issues can be set aside here.

⁸The physical abilities I have in mind, like the ability to walk, consist partly in certain psychological capacities, like capacities for motor control. But the abilities do not consist exclusively in such psychological capacities. Having legs is part of having the ability to walk.

ability to walk. In core cases, we have non-observational warrant for relying on exercises of our physical abilities.

9

When we are entitled to rely on exercises of physical abilities in particular cases, the entitlement in those particular cases are instances of a more general entitlement to rely on our physical abilities. In general, we have a default entitlement to rely on exercises of our physical abilities in carrying out our various actions. That default entitlement is defeasible. If you have evidence that your physical abilities are inadequate for a carrying out particular action, that evidence can “override” your default entitlement. So, for example, if you know that juggling is difficult and that you never learned how to do it, then you would lack warrant for relying on your physical abilities to carry out an action of juggling. But, according to the view I argue for, if you had no such discouraging evidence, you would have a default, non-observational entitlement to rely on your physical abilities to carry out an action of juggling.¹⁰

The account I offer is indebted to Burge’s accounts of various epistemic entitlements,

⁹It would have been possible to formulate the view I defend as being about warrant for exercising one’s abilities, instead of being about warrant for relying on exercises of abilities. I formulate my claim in terms of reliance because warrant for exercising abilities to carry out action may be interpreted as warrant for action, where that warrant can depend on evidence of whether the action is good to do. Talking about reliance emphasizes that the warrant concerns the matter of whether one’s abilities are working properly, and not whether they are being put to good use. I am not claiming in this chapter that our actions themselves have non-observational warrant (though I do think in certain cases we can have non-observational warrant for actions).

¹⁰I have distinguished here between specific entitlement on specific exercises of physical abilities, and a general entitlement to rely on physical abilities in general. There is another kind of entitlement concerning physical abilities, which concerns a *prospective* kind of reliance on physical abilities. In intending to perform a physical action, antecedently to actually carrying out the action, one relies on one’s physical abilities to carry out the intended action in a prospective way. Since the prospective reliance is antecedent to the exercise of one’s physical abilities, one is not relying on specific exercises of one’s abilities in a case like that. This reliance is generic in one way, and specific in another way. The reliance is specific in that one is relying on one’s physical abilities in a specific case, to carry out a specific action. The reliance is generic in that one is not relying on a specific *exercise* of physical abilities, since the reliance is antecedent to a specific exercise of physical abilities. One is relying on one’s physical abilities to carry out an action of the intended type prospectively, and antecedently to relying on specific exercises of one’s physical abilities to carry out a particular action that is underway.

especially his accounts of epistemic entitlements to rely on memory (1993) and perception (2003).¹¹ Suppose you arrive at a conclusion in reasoning, store that conclusion in memory, and later on you make use of that conclusion in further reasoning. There is a question as to whether you are entitled to rely on your memory to accurately store your conclusion, or whether you need some evidence that your memory can be relied on. Burge's view is that you are entitled to rely on memory partly because you are entitled to rely on reason, and memory plays an essential, constitutive role in reason's functioning.¹²

A similar issue arises if you form a perceptual belief on the basis of perception. There is a question as to whether you are entitled to rely on perception to represent accurately, or whether you need some evidence that your perception can be relied on. Burge's view is that we are entitled to rely on perception.¹³ The entitlement is not due to our entitlement to rely on reason, since perceptual capacities are not rational capacities. When our perceptual capacities are reliable, we are entitled to rely on them because they are reliably accurate representational capacities. The reliable accuracy of perception is relevant to epistemic warrant because that reliability enables perception to facilitate our belief system's being a good route to truth, and epistemic warrant is associated with the epistemic success of true belief.

The main idea of the account I present is that we are entitled to rely on exercises of

¹¹Burge's discussions of warrant for reliance on testimony (1993) and computer proof (1998) also inform the discussion in this presentation. Burge does not now hold the view that warrant for relying on computer proof or on testimony can be non-observational.

¹²According to Burge, beings who have perceptual capacities but lack reason have an analog of epistemic warrant to rely on perceptual memory, because of the role that perceptual memory plays in the functioning of their representational capacities.

¹³Burge's view is that we have a non-observational entitlement to rely on perception in general. Entitlement for relying on particular perceptions is trivially observational, inasmuch as it is an entitlement to rely on observational evidence registered by perception. There is another sense in which entitlement to rely on particular perceptions is non-observational. The entitlement is not grounded in observational evidence in the individual's psychology, indicating that the particular perception can be relied on. Nevertheless, the entitlement to rely on particular perceptions is grounded partly in the reliably successful exercise of perceptual capacities in the individual's ancestors.

physical abilities because the physical abilities being exercised are reliably effective aspects of agency. We are entitled to rely on our agential capacities, just as we are entitled to rely on representational capacities like memory and perception. I develop this account in sections 3 and 4. I emphasize the analogy to perception, since neither perceptions nor exercises of physical abilities are involved in reasoning.

In section 5, I provide more support for the view that we are entitled to rely on exercises of physical abilities by discussing the role that physical abilities play in the full well-functioning of practical reason. We have a default entitlement to rely on practical reason to carry out its full function of guiding action, including physical action. As a result, we must be entitled to rely on our abilities to carry out our actions, including physical abilities. I do not claim, however, that our entitlement to rely on our physical abilities *derives* from our entitlement to rely on practical reason.

I conclude the chapter by suggesting that our entitlement to rely on practical reason to carry out its full function derives in part from the role of our physical abilities in our agency. As a result, our entitlement to rely on practical reason to carry out its full function is not grounded just in the nature of reason itself, and so is not apriori.¹⁴

Our entitlement to rely on practical reason to carry out its full action guiding function, and our entitlement to rely on exercises of physical abilities, are non-observational in particular cases, without being apriori. An individual's warrant is non-observational when the warrant does not depend for its force on exercises of any of the individual's capacities for sensory or perceptual registration. To put the point in a rougher and more intuitive way, a warrant is observational when it depends on the way the individual makes use of perception and sensation to indicate that she is right to be and do as she is and does. But a warrant's

¹⁴I distinguish between practical reason's full action guiding function, and its narrower function of guiding our intentions. In this chapter, I do not address the question of whether our entitlement to rely on practical reason to guide intention is apriori. My view is that this warrant is also non-apriori. I defend that view in other work.

being non-observational does not mean it is apriori. To be apriori, the warrant must be due only to exercises of rational capacities, and to the nature reason. Warrant consists in something practically or epistemically good about how things are with an individual. Even if an individual's warrant is non-observational, the explanation of what is practically or epistemically good about how things are with her can refer to more than just exercises of her rational capacities, and the nature of reason. The explanation can refer to the nature of non-rational capacities, and to exercises of those capacities in the history of the species.

We can have a non-observational entitlement to rely on exercises physical abilities. But a full explanation of the relevant practical goodness of the reliance will not refer just to exercises of rational capacities and the nature of reason. Part of the explanation will refer to the reliable effectiveness of the physical abilities being exercised. The explanation might also refer to the role of sensory and perceptual registration of one's actions in the history of the species. The role of these sorts of considerations in the explanation of the warrant prevents the warrant from being apriori, but is consistent with the warrant's being non-observational. The warrant for relying on exercises of physical abilities can still be independent of any perceptual or sensory grounds possessed by the individual. In a similar way, our entitlement to rely on practical reason to carry out its full action-guiding function can be non-apriori, while still being non-observational.

An objection to Anscombe's view about knowledge of action

Before presenting my variant of Anscombe's view, I want to say more about the problem for Anscombe's view that motivates the account. I also want to briefly criticize one way of responding to the problem, which bears some similarity to my own. In characterizing the problem, I focus on the more widespread variant of Anscombe's view, which concerns

epistemic warrant for beliefs about action. I then go on to suggest that we can deal with this sort of problem if we think of Anscombe's view as a view about *practical* warrant for reliance on exercises of physical abilities, instead of being about epistemic warrant for beliefs about actions. Defending an epistemic variant of Anscombe's view would take further argument. I am unsure about the prospects of the epistemic variant. In any case, I do not defend that view here.

Let me turn to the problem for Anscombe's view. It is not clear why our knowledge of what we are doing should be different from other kinds of knowledge about contingent, extra-mental matters of fact, like knowledge about the weather. Clearly, knowledge about one's physical actions cannot derive non-observationally from one's understanding of concepts in the way that one's knowledge that red is a color can. The fact that one is walking rather than standing, sitting, or jumping, cannot be known in that way. Moreover, though physical actions have deep relations to aspects of one's mental life, like one's intentions and one's values, knowledge about one's physical actions cannot be acquired as a straightforward extension of one's knowledge of one's intentions, or of other aspects of one's mental life. Whether one is actually carrying out a physical action depends not just on how things are with one's mind, but also on how things are with one's physical body, and on the effects of one's physical body on the larger world.¹⁵

Like the weather, one's physical actions are contingent, extra-mental matters. Usually knowledge about contingent, extra-mental matters of fact requires some kind of observational tracking of how things happen to be with the world. Without grounding by observation, one's view about whether it is *raining* could only be, at best, accidentally correct, and could not

¹⁵Discussions of Anscombe's view often focus on another objection—that knowledge of action must be based on one's awareness of one's intention to act. (See, for example, Grice (1971), Velleman (1989), Setiya (2008).) Whether one is actually carrying out some intentional action depends on whether one intends to carry out that action. As a result, it might seem to be a kind of "leap of faith" to believe that one is ϕ ing if the belief is not inferentially based on an intention to ϕ , or a belief about such an intention. I do not focus on this criticism here, because it does not target the claim that warrant concerning action is non-observational.

count as knowledge. It is initially quite plausible that knowledge of one's physical actions requires observational evidence of how things are with one's physical body, for the same reason that knowledge of the weather requires such observational grounding. It would seem that, to count as knowledge, one's view about physical action must be based on observational evidence, which indicates how things are with one's physical body, and its effects on the larger world. If not, one's view about one's physical action would seem to be like an arbitrary guess.¹⁶

This problem is pronounced in cases of action that involve objects in the external world. For example, whether you are successfully carrying out an intention to be grasping a cup depends in part on whether the object in your hand is a cup. Whether the object in your hand is a cup is a contingent matter of fact, that seems to be knowable only empirically. As a result, it would seem that warrant for a belief that you are actually grasping a cup depends on observational grounds for thinking the object in your hand is a cup. If you could know without observation that you are grasping a cup, then, by reasoning off of your non-observational knowledge of what you are doing, you could know without observation that there are cups in the world, and that one of them is nearby. It would be very surprising if you could gain non-observational knowledge of those contingent facts in that way.

There is a similar problem for the Anscombian view even when it comes to actions that only concern the agent's own body, like clenching one's fist. Whether you are actually clenching your fist depends on whether you have hands at all, on whether you possess the required physical abilities, whether circumstances are hospitable to their exercise, and so on. Those are contingent matters of fact, which seem to be knowable only empirically. If you

¹⁶One strategy for responding to this sort of worry is to claim that whether one is carrying out an action does not depend on what one's body is actually doing, and its effects on the world. Perhaps it can be correct to say of an individual that she is writing, even if her pen is out of ink, or if she is unable to produce letters on the page because her hand is trembling so violently. In my view the truth conditions of physical action descriptions require that the body succeed in at least some aspect of the intended action. In any case, in this chapter I am interested in the question of whether one can have non-observational warrant concerning *successful* action.

believe you are actually clenching your fist without observational grounds indicating that the required empirical matters of fact actually obtain, then your belief would seem to be like an arbitrary guess, and not to be warranted.

There is a tradition of claiming that one can have non-observational knowledge of elements of one's mental life, like beliefs and intentions. Though it is difficult to explain the source of warrant for beliefs about one's mental life, it makes sense that the warrant would be non-observational. Crudely, observational capacities are "input interfaces" between the mind and the extra-mental world, which register information about the world for use in one's thinking. It makes sense that the mind would not require that sort of input for knowledge about itself. But knowledge about our physical actions is knowledge about what our bodies are actually accomplishing. It is not evident how we can have knowledge about our own bodies without observation, any more than we can have knowledge about the weather, or about others people's bodies, without observation.

We can avoid the unwelcome proliferation of non-observational knowledge if we focus on warrant for reliance on our physical abilities, instead of focusing on knowledge about our actions. There is no straightforward way of gleaning non-observational knowledge about the existence and location of cups, or about the make-up of one's body, from non-observationally warranted *reliance* on the abilities involved in grasping a cup. But the main worry sketched here still seems to apply to that variant of Anscombe's view. Whether you are right to rely on your physical abilities depends on how things are with your physical body. Without observational evidence indicating how things are with your physical body, relying on your physical abilities would seem to be like an arbitrary, unwarranted guess. The account I present here is meant to address that issue. The main idea is that warrant for relying on physical abilities does not require observational evidence because our physical abilities are integral to the functioning of our agency.

Setiya on knowledge of action and know how

Let me turn to Setiya's Anscombian account (2008, (2009) of the source of our knowledge of physical action, which is meant to answer an objection along these lines. Setiya's account is meant to answer the objection that, without observational warranting grounds, beliefs about action would be at best accidentally true, and would lack warrant. His account tries to explain the possibility of non-observationally warranted beliefs¹⁷ about action in terms of know how. As Setiya uses the term 'know how,' know how consists in abilities to carry out intentions, like the ability to walk, to dance a tango, or to extend one's arm. What Setiya calls 'know how' is roughly the same as what I call 'physical abilities.'¹⁸

According to Setiya, know how provides warrant for beliefs about action because know how makes the agent's beliefs non-accidentally true (2009, pp. 136-7). When an individual is intentionally extending her arm, for example, and she believes that she is extending her arm, Setiya claims that her knowing how to extend her arm "ensures non-accidentally true belief" that she is doing so (2009, p. 136). The idea seems to be that, because the action is an exercise of know how, the belief is different from a guess that happens by chance to be true.

I want to claim that Setiya's account is inadequate, because the sort of non-accidentalness that exercises of know how help give to the truth of beliefs about action seems not to be relevant to whether a belief is warranted, or amounts to an unwarranted guess. The truth of a belief can be accidental in two ways. It can be accidental in that the fact that the belief is about is a fluke, an accidental result of chance. Or the truth of the belief can be accidental in that the agent's belief amounts to a guess, which happens by chance to be true. Only this second way is directly relevant to warrant. It is easy to see that know how helps to make

¹⁷Or non-observationally "justified" beliefs, in Setiya's terminology. (See note 1 above.)

¹⁸Setiya's term is broader though, since it includes abilities to carry out mental actions as well as abilities to carry out physical actions.

the truth of an agent's belief about her physical actions non-accidental in the first way. But it is not clear how an agent's know how helps to make the truth of her belief non-accidental in the second way.

When an agent is in fact performing an intentional action that she believes she is performing, her know how helps to make it “non-accidentally” the case that she is actually doing the thing that she believes she is doing. If she did not know how to juggle, for example, she would not be juggling. Or else, if she were juggling, it would be an extraordinary fluke that she somehow managed to juggle despite not knowing how.

But that point does not show that the know how helps to make the agent's belief non-accidentally true in a way that is relevant to warrant. Consider my belief that my eastward transatlantic flight will be faster than the westward return trip. The jet stream helps to explain the phenomenon that my belief is about, but it does not help to explain why my belief is better than a lucky guess. Even if my belief were a baseless, unwarranted lucky guess, the jet stream would still explain the truth of the belief in that way. Similarly, the juggler's know how helps to explain the phenomenon that her belief is about—it helps to explain how it came to pass that she succeeding in her attempt to juggle. But the fact that know how makes the truth of an agent's belief about her action non-accidental in this way is not evidently relevant to the belief's warrant. At any rate, Setiya's discussion does not explain how it might be relevant.

To explain how a non-observational belief about physical action can be warranted, it does not help to show that the occurrence of the action was no fluke. One has to explain why the accuracy of the belief was no accident, instead of being a lucky guess. It is difficult to see how Setiya's invocation of know how helps to provide such an explanation. At least, Setiya has not shown how it helps to provide such an explanation.¹⁹

¹⁹One might suggest that know how can explain warrant for beliefs about action if possessing know how consists partly in the ability to discern what one is doing. It is plausible that many of our aptitudes do consist partly in such abilities. The ability to lift a cup to one's mouth consists, partly, in sensory “feedback”

Nonetheless, the link Setiya draws between action and know how, or physical ability, is useful for explaining how we can have non-observational warrant concerning physical action (practical warrant, as I will argue). But, whereas Setiya makes use of such abilities as *explanans*, I think it is more fruitful to use them to look at the *explanandum* in a different way—to think in a different way about what our non-observational warrant concerning physical action is warrant *for*. The importance of physical abilities is not that they explain the occurrence of intentional actions in the way that the jet stream explains the speed of my trans-atlantic flight. The importance of physical abilities is that they allow us to understand warrant concerning physical action as warrant that concerns reliance on agency itself. I want to argue that we have a default, non-observational practical entitlement to rely on exercises of our physical abilities, because of the role of physical abilities in our agency. Because the entitlement is due to the role of physical abilities in agency, it is a kind of *practical* warrant, and not epistemic warrant.

In intentional action, we rely on exercises of our physical abilities to play their role in carrying out the intended action. That reliance is subject to norms of warrant. If an individual has good reason to doubt that she can carry out some action, but nonetheless relies on exercises of her physical abilities, that reliance is a mistake, and is unwarranted. For example, if an agent relies on her physical abilities to carry out an intention to juggle, but knows that juggling is difficult and that she never learned how, then her reliance is unwarranted. If instead her reliance is supported by warranting grounds, like observational evidence that she has successfully juggled in the past, then her reliance can be warranted. I want to argue that warrant for reliance on exercises of physical abilities does not have to

capacities that monitor how the action is developing, so one can make responsive adjustments as the action continues. If these “feedback” capacities are partly constitutive of know how, then at least in many cases we can tell what we are doing through the exercise of know how. But that point is no help in explaining how warrant for beliefs about action could be non-observational. These capacities involved in discerning what one is doing, which are partly constitutive of know how, are observational capacities. As a result, knowledge of what one is doing that is acquired through their exercise would not be non-observational knowledge.

depend for its force on observational evidence.

Perceptual entitlement

Our entitlement to rely on perception helps reveal why non-observational reliance on exercises of physical abilities need not be like an unwarranted arbitrary guess. I want to lay the groundwork for my account of our entitlement to rely on physical abilities by discussing our entitlement to rely on perception.²⁰

We have an entitlement to rely on our perceptual capacities, without evidence establishing that they are effective in general, or functioning effectively in the present situation. If one is not warranted in relying on perception, then it is an epistemic mistake to form beliefs on the basis of perceptions, which undermines the warrant of those beliefs. But one can have warranted beliefs based on perception without making use of evidence that one's perceptual capacities operate effectively. That is the norm for children, higher animals capable of belief, and probably many adults. Thus, we must have an entitlement to rely on perception, which does not depend on evidence establishing the effectiveness of perception.

But it would seem that whether our perceptual capacities are generally effective, and whether they are functioning effectively in any particular case, are questions to be settled with empirical evidence. One cannot discover the effectiveness of one's perceptual capacities through conceptual analysis or introspection. If one wanted to know whether some other animal's perceptual capacities (or putatively perceptual capacities) were effective, one would have to make use of empirical evidence. It is tempting to think that the effectiveness of our own perceptual capacities would likewise have to be determined through empirical investigation. How, then, can we be entitled to rely on our perceptual capacities to function effectively without the sort of evidence one would make use of in such an investigation?

²⁰The discussion of entitlement to rely on perception is indebted throughout to Burge (2003).

Lacking such evidence, our reliance on our perceptual capacities would seem to be like an arbitrary guess, and to be unwarranted.

To begin to respond to this challenge, one might point out that our perceptual capacities are in fact reliably accurate. The reliable accuracy of perceptual capacities helps to make them part of a good route to truth, when they are integrated with epistemic reason. Since perception is reliably accurate, that makes relying on perception epistemically good in a certain respect. One might suggest that we are epistemically entitled to rely on perceptual capacities because of the way in which their reliable accuracy makes it epistemically good to rely on them.^{21, 22}

It seems to me correct that our entitlement to rely on perception is due in large part to the fact that the reliable accuracy of perception makes it part of a good route to truth. If our perceptual capacities were too far out of sync with reality, that would undermine our entitlement to rely on perception.

But these points do not yet answer the question about why reliance on perception is different from an arbitrary guess. In relying on our perceptual capacities, we are “presuming” that they are effective. If our perceptual capacities are reliably accurate, and part of a good route to truth, that shows that the “presumption” is in a certain respect correct, and the reliance is rightly placed. But one might still ask why this rightly placed non-observational reliance on perceptual capacities is not a lucky, unwarranted accident. Even if our perceptual

²¹According to Burge (2003), it is the reliability of our perceptual capacities in their normal environment that matters to entitlement. The normal environment is the environment that plays a role in determining the content of perceptual states. This is the environment in which the being’s perceptual capacities evolved.

²²Pre-rational perceivers have license to rely on their perceptual capacities that is analogous to our epistemic entitlement. It is representationally good for them to rely on perception. This representational goodness is analogous to the epistemic goodness of our reliance on perception. It is representationally good for them to rely on perception partly because perception is a good route to accurate representation. In those beings perception is not part of a good route to truth, because truth is accuracy for propositional representational capacities, and pre-rational beings lack capacities for propositional representation. That is why their reliance on perception is not an *epistemic* good, and so not an entitlement, even though it is a representational good.

capacities are *in fact* reliably accurate, one may still ask what licenses us to “presume” that they are reliably accurate. In general, the fact that something actually is reliable does not show that one has license to rely on it. The fact that your car is reliable does not make me warranted in relying on it to get me where I need to go. Warrant for relying on the car depends not on the fact of its reliability, but on one’s tracking evidence of its reliability, like whether the car is rickety-looking, how the engine sounds, and so on. Without making use of such evidence, reliance on something that is in fact reliable would seem to be like an unwarranted guess that happens to be right by accident.

We can respond to this worry by considering what sort of accidental-ness undermines epistemic warrant. The relevant accidental-ness cannot consist in making a commitment, or relying on capacities, without evidence. That proposal precludes the possibility of epistemic entitlements. Precluding the possibility of epistemic entitlements makes knowledge and epistemic warrant a remote ideal, nearly impossible to attain, except perhaps in very special cases of self-evident truths.

Here is an alternative. Roughly, epistemic warrant is a matter of one’s representational capacities functioning as a good route to truth. Accordingly, the relevant sort of accidental-ness does not concern, in the first instance, the good use of evidence, but the good functioning of one’s representational capacities. If you get something accidentally right in a way that undermines warrant, your getting it right is accidental with respect to the good functioning of your representational capacities. In the other direction, when you get something right in a way that is non-accidental with respect to the good functioning of your representational capacities—that is explicable in a systematic way in terms of the good functioning of your representational capacities—your getting it right is not accidental in a way that undermines warrant.²³

²³The phrasing here echoes Burge’s characterization of brute truth (2003, p. 507). Cases of brute truth are Gettier cases. They are cases of justified true belief that falls short of knowledge. In such cases, your representational capacities are functioning well, but there is some way in which the truth of your belief is

The use of evidence has a role in providing for this non-accidental-ness in some cases, but not in all cases. Suppose your belief is formed through a good exercise of your representational capacities to make good use of evidence in beliefs. The belief's being formed in that way shows the truth of your belief not to be accidental with respect to the good functioning of your representational capacities.

But what about your exercise of those representational capacities? You get something right in making use of those capacities (i.e., in relying on them) , in a way that is non-accidental with respect to the good functioning of your representational capacities. But making use of representational capacities does not need to be based on evidence in order for the use to be non-accidentally right in that way. It is part of the normal good functioning of one's representational capacities that their possessors make use of them and in doing so rely on a good route to accurate representation.

It is worth noting that our representational capacities are not just capacities to engage in demonstrative reasoning. As a result, an understanding of the validity of canons of demonstrative reasoning falls far short of an understanding of how representational capacities function as a good route to truth. Representations are often formed according to complex patterns of induction. We are usually unaware of the inductive procedures we make use of, and our reliance on these procedures is not grounded in evidence of their reliability. Moreover, if warrant for reliance on representational capacities depended on such evidence, that would lead to a circle. We could not gain this kind of evidence for the reliability of our representational capacities without relying on the capacities.

When you rely on your perceptual capacities, you are in some way getting it right (rightly “presuming”) that your perceptual capacities are reliably accurate. It is part of the normal

accidental with respect to that well-functioning.

Here I have in mind that, if your representational capacities are performing poorly, but your belief is true anyway, then there is also a way in which the truth of your belief is accidental with respect to the good functioning of your representational capacities. The truth of your belief cannot be explained in terms of the good functioning of your representational capacities.

good functioning of our representational capacities that we rely on reliably accurate perceptual capacities to operate as a good route to accurate representation, without our reliance being guided by observational evidence of the effectiveness of perception. So non-observational reliance perceptual capacities can be non-accidentally right with respect to the good use and functioning of one's representational capacities. As a result, relying on perceptual capacities without observational capacities need not be accidentally right in a way that undermines the warrant of the reliance.

The reliable accuracy of perception is relevant to *epistemic* warrant because that reliability enables perception to facilitate our belief system's being a good route to truth, and epistemic warrant is associated with the epistemic success of true belief. Our perceptual capacities play a role in guiding belief. Our reliance on perception is part of the normal good functioning of the representational capacities involved in guiding belief to truth (our "epistemic capacities," as I call them). It is the functioning of these epistemic capacities that is relevant to epistemic warrant.

Pre-rational, pre-epistemic beings with perceptual capacities have a non-epistemic license to rely on their perceptual capacities. This license is a matter of meeting certain norms on good representation. If an animal gets something accidentally right, in a way that undermines the pre-rational analog to warrant, the right-ness is accidental with respect to the good functioning of its representational capacities. It is part of the normal good functioning of the animal's representational capacities that it rely on its perceptual capacities as a good route to accurate representation, and that it be right to do so. Thus, in relying on its perceptual capacities without anything like evidence of their effectiveness, the reliance is not accidentally right in a way that undermines the pre-rational analog to warrant.

Let me emphasize that the core element of the explanation of our entitlement to rely on perception is the same as the core element in the explanation of this pre-rational analog of entitlement. It is the fact that, in relying on perception, one relies on representational

capacities that are reliably accurate, and are thus a good route to accurate representation. The connection of perception to belief in human beings is what makes that representational goodness epistemically relevant.

To clarify the line of thought here about our entitlement to rely on perception, it will help to consider an objection. One might suggest that a parallel line of thought leads to the conclusion that we have an entitlement to any belief, no matter how it is formed. It is part of the normal good functioning of your epistemic capacities that you form true beliefs. So you should be able to “presume” that any belief of yours is true.

The problem with this suggestion is that it is not part of the normal good functioning of your epistemic capacities that you form beliefs willy-nilly. Doing so constitutes a failure in the exercise of those capacities. So, if you form a belief in that way, and happen to be right, your being right is accidental with respect to the good functioning of your epistemic capacities. Norms on the functioning of epistemic capacities mark aspects of the normal good functioning of those capacities. Satisfying these norms consists in forming beliefs through certain “routes” that constitute the normal good functioning of one’s epistemic capacities. Part of what makes those “routes” to belief epistemically good is that they are reliably successful in their role in guiding the individual to true belief. If one forms a belief willy-nilly, the route to the belief is unreliable, and thus not epistemically good. So it is not part of the normal good functioning of one’s epistemic capacities, and the belief lacks warrant.

Being non-accidentally right in the relevant way does not mean being right about something when the good functioning of your epistemic capacities usually enough issue in your being right about such things. It means being right when your epistemic capacities actually are functioning well, and that good functioning explains your being right in the particular case at hand.

Unlike forming beliefs willy-nilly, non-observational reliance on perception is part of the normal good functioning of one’s epistemic capacities. As a result, non-observational reliance

on perceptual capacities can be non-accidentally right with respect to the good functioning of one's epistemic capacities. Thus, we can be entitled to rely on our perceptual capacities.

Physical abilities and agency

I want to suggest that we can have a similar entitlement to rely on exercises of physical abilities. We are entitled to rely on exercises of physical abilities because rightly relying on exercises of physical abilities is non-accidental with respect to the normal good functioning of our agency.

The warrant we have for relying on exercises of our physical abilities is not epistemic warrant, unlike warrant for relying on perceptual capacities. Epistemic warrant is a success in the operation of our epistemic capacities, the capacities constitutively involved in the guidance of belief. Physical abilities, like the ability to walk, or raise one's arm, are not epistemic capacities.²⁴

The warrant that applies to our reliance on exercises of our physical abilities is not epistemic warrant but *practical* warrant. Whereas epistemic warrant is associated with being a good route to true belief, practical warrant is associated with being a good route to good action. Practical warrant is a success in the operation of the capacities constitutively involved in guiding one to good action. I refer to these capacities collectively as 'agency.' Roughly, guiding action consists in determining how one acts through the exercise of capacities that function to determine how one acts, where that function subserves action's function of being good. Our agency guides our action in carrying out this function. Carrying out the function comprises exercises of capacities to carry out actions, as well as capacities to determine what

²⁴It is true that physical abilities can have an important role in providing for epistemic success. You use your hands to turn the pages of a book that you are learning from. But the role of your hands in arriving at true belief is very different from the role of your capacities for reasoning, understanding, perception, and sensation in guiding your beliefs. A belief cannot be based or grounded in the exercise of your physical ability to use your hands.

is to be done.²⁵

By “good action,” I mean action that realizes some standard of correctness analogous to truth for belief. I do not assume that the relevant standard is *moral* goodness. In my view the standard is that an action help, on the whole, to realize the agent’s legitimate ends. Non-moral agents have legitimate ends, like survival, mating, pleasure, and avoiding bodily harm. One might disagree about what the relevant standard of goodness is. That disagreement does not matter for present purposes. What matters is that we have agential capacities that function to determine how we act, so that our actions satisfy some standard of goodness. Practical warrant is associated with that practical function of our agency, just as epistemic warrant is associated with the epistemic function of our epistemic capacities.

Getting something accidentally right in a way that undermines epistemic warrant is getting it right in a way that is accidental with respect to the good representational functioning of one’s epistemic capacities. Similarly, getting something accidentally right in a way that undermines practical warrant is getting it right in a way that is accidental with respect to the good practical functioning of one’s agency.

For example, suppose one buys a good car in spite of strong evidence that it is shoddy. The action is right, in a way, in that the purchase worked out, and serves the legitimate transportation needs it is supposed to serve. But the action is accidentally right in a way that undermines its warrant, since the individual’s agential capacities performed poorly in determining how she acts. The action’s being right is thus accidental with respect to the good functioning of her agency.

Now, if non-observational reliance on exercises of physical abilities can be non-accidentally right with respect to the good functioning of our agency, that should help to explain how we can be entitled to rely on our physical abilities. I want to suggest that one can have an entitlement to rely on physical abilities for that reason.

²⁵I discuss the concept of action guidance in more detail in “Reason’s guidance of action.”

Physical abilities play a central constitutive role in our agency. Most actions of most agents are physical actions. Our abilities to carry out actions, including our physical abilities, are a part of our agency. An agent, or at least a non-defective agent, is a being capable of acting, and not just capable of determining what is to be done. Agency comprises the capacities constitutively involved in guiding our actions, where carrying out an action is part of successful guidance of action—in roughly the same way that forming a belief is part of successful guidance of belief. Agency comprises both our abilities to actually carry out actions, as well as our abilities to determine what is to be done. Our physical abilities are the capacities through which we carry out most of our actions. In this way, our physical abilities have a central constitutive role in agency.

Practical warrant is associated with our agency's functioning as a good route to correct action. Being practically warranted is not just a matter of doing as well as possible as an agent, given the limitations of one's agency. It is a matter of "actually" doing well, where that involves the good functioning of good equipment, as it were. Doing as well as possible as an agent given the limitations of one's agency is also a genuine practical success, but it is not the one I am interested in here. Practical warrant is a success analogous to a good toaster that is functioning properly, rather than a bad toaster that burns toast even though it is functioning as well as it can (i.e., is not malfunctioning). Or, to make a closer comparison, it is a success analogous to the intelligent thinking of an intelligent person, not the hopelessly poor thinking of a person who is thinking as well as possible given his severe intellectual limitations.

For agency to be a good route to action in a way that is relevant to warrant, it must be reliably successful in fulfilling its action-guiding function when it functions normally, in its normal environment.²⁶ If one's agency fails to meet this condition, it is like a bad toaster,

²⁶In my view, the normal environment is the environment that plays a role in individuating the representational mental states involved exercises of agency. That view is similar to Burge's (2003) view that the normal environment relevant to perceptual entitlement is the environment that plays a role in the individuation of

or a person with severe intellectual limitations. The reliable effectiveness consists in one's agency reliably guiding one to correct action. We carry out our actions through the exercise of our physical abilities, at least in large part. As a result, the reliable success of our agency consists partly in one's physical abilities being reliably effective in carrying out one's actions. The reliable effectiveness of our physical abilities is part of what makes our overall agency a good route to action.

It may be possible in principle for a being's agency to be a good route to action, without a role for physical abilities. If angels are a coherent possibility, then their agency would be an example. But I assume that it is a central part of the nature of our agency, and the agency of all existing agents, that it functions to guide physical actions, even if it also functions to guide mental actions.²⁷ Moreover, I assume that our kind of agency is embodied agency according to its nature, and not contingently. It is part of the nature of our kind of agency that it functions to guide physical actions. Thus, it is given by the nature of our agency that its being a good route to action depends constitutively on the reliable effectiveness of our physical abilities.²⁸

In us, the reliable effectiveness of physical abilities in our agency consists largely in their role in executing intentions with conceptual content. In some cases, physical abilities can operate as part of our agency without being guided by concept-involving intentions. For

the contents of perceptual states.

²⁷In saying agency functions to guide action, I mean that the various capacities that constitute our agency, including practical reason, motor skills, for example, function to guide our actions. These capacities function to causally determine how we act, where that function subserves action's satisfaction of its "standard of correctness"—namely, being something that is good to do. I discuss the guidance of action by such agential capacities in chapters one, three, and four.

²⁸I am inclined to think that the reliability of our physical abilities is *essential* to our agency being a good route to action. But the point here does not require that claim. Maybe it is in principle possible that agency's being a good route to action could be overdetermined. Angels with physical arms and legs added might be an example. Physical abilities could contribute constitutively to the systematic effectiveness of their agency, even though the effectiveness of our agency would not depend essentially on the effectiveness of their physical abilities. The agency of the unlikely embodied angels would still be a good route to action, even if they had lacked physical abilities.

example, reflexively withdrawing one's hand from a hot surface seems to be an exercise of agency that is not guided by a conceptualized intention to move one's hand. In pre-rational representational animals, the reliable effectiveness of physical abilities consists in their role in executing non-conceptualized representations of actions, or in carrying out actions that function to realize representations of practical goals. In both cases, the reliable effectiveness of physical abilities helps to make agency a good route to action.

When we rely on exercises of reliably effective physical abilities, we are practically right to do so, in that the physical abilities being exercised are a good part of a good route to action. This point echoes the point that we are epistemically right to rely on exercises of reliably accurate perceptual capacities because their exercise is a good part of a good route to true belief. In the perception case, the question arose of whether it is a warrant-undermining lucky accident that we rightly rely on our perceptual capacities. The same question arises about our reliance on exercises of our physical abilities. Even if we are right to rely on exercises of our physical abilities, it would seem to be an unwarranted lucky accident, unless the reliance is based on evidence of the effectiveness of our physical abilities.

It is part of the good functioning of our agency that we rely on exercises of our physical abilities to carry out our actions. We cannot physically act without relying on exercises of our physical abilities.²⁹ Physically acting is a central part of the normal good functioning of our agency, and doing so involves relying on exercises of one's physical abilities. Moreover, it is part of the normal good functioning of our agency that we are *right* to rely on exercises of our physical abilities. In the normal good functioning of agency, one's agency operates as a good route to correct action, reliably guiding one to correct action. The reliably effective exercise of physical abilities is a constitutive aspect of this good functioning, at least in

²⁹There might be exceptions. Perhaps if you carry out some action when you are convinced that it is beyond your capabilities, you act without relying on your abilities. It seems to me more plausible to say that in such a case you do rely on exercises of your physical abilities, while judging that it is a mistake to do so.

embodied agents like us. Thus, being right in relying on exercises of physical abilities is not accidental with respect to the normal good functioning of agency. It is part of the normal good functioning of our agency that we rely on exercises of our physical abilities when in fact they are reliably effective in carrying out our actions.

An objection

One might object that this line of thought does not suffice to show how we can have non-observational warrant for relying on exercises of our physical abilities. In the earlier discussion of perceptual entitlement, I suggested that non-observational reliance on reliably accurate perception is part of the normal good functioning of our epistemic capacities. As a result, non-observational reliance can be non-accidentally right with respect to the normal good functioning of one's epistemic capacities. Thus, we can have non-observational warrant for relying on perception. It is true that reliance on exercises of reliably effective physical abilities is part of the normal good functioning of our agency. But one might doubt that *non-observational* reliance on exercises of reliably effective physical abilities is part of the normal good functioning of our agency.

It is plausible that, when our agency is functioning well, we rely on our physical abilities on the basis of observational evidence about the effectiveness of our physical abilities. In the normal course of development, as one develops physical abilities to walk and to grasp, for example, one "picks up on" the fact that one possesses these abilities through perception, proprioception and sensation. One might suggest that, when our agency is functioning well, we do not rely on one's physical abilities willy-nilly, but under the influence of sensitivities like these. Rightly relying on exercises of physical abilities might be part of the good functioning of our agency only when the reliance is guided by observational evidence.

If so, then it would seem that non-observational reliance on exercises of reliably effective

physical abilities would be accidentally right with respect to the normal good functioning of our agency. Non-observational reliance would seem to be unwarranted. To show that we can have non-observational warrant for relying on exercises of physical abilities, one would have to show that non-observational reliance on physical abilities is part of the normal good functioning of our agency.

As a preliminary point, note that even if the objection is correct, and the warrant is observational, the warrant is still likely to be an entitlement rather than a justification through reasoning. Suppose that when our agency is functioning well, our reliance on exercises of our physical abilities is informed by some kind of observational evidence about what physical abilities we possess. Even if this were true, it is very likely that much of the observational guidance of our reliance on our physical abilities is not accomplished through reasoning about one's physical abilities. When an infant "picks up on" her developing ability to grasp objects, to control the movements of her neck, arms and legs, she likely does not do so through reasoning. It is not clear that an infant even possesses the capacity to reason. It is likely that her disposition to rely on exercises of her developing physical abilities is guided by non-rational sensitivities, like proprioception, perception, and sensation, and not by reasons she makes use of to establish that she possesses those physical abilities.

If so, it is part of the normal good functioning of our agency that we rely on exercises of certain physical abilities without using observational evidence of their effectiveness in reasoning. We can be warranted in relying on exercises of those physical abilities without such rational support. Our warrant for our reliance on exercises of those physical abilities would be a kind of entitlement, since it does not depend for its force on reasons. But the warrant would still be observational, like warrant for perceptual beliefs. Though the warrant would not depend on reasons, it would depend for its force on guidance by observational capacities that register information about the world.³⁰

³⁰Our entitlement to rely on these sensitivities would be due, in large part, to their reliable accuracy in

But I do not want to concede that warrant for relying on exercises of physical abilities must depend for its force on observation. Whether the objection sketched above is sound depends on empirical facts about how agency actually works. Dispositions to rely on physical abilities can be innately “wired in,” where those physical abilities contribute to one’s agency being a good route to action. In such an individual, non-observational reliance on exercises of physical abilities could be non-accidentally right with respect to the normal good functioning of its agency.

Indeed, it seems very likely that many agents work this way. It seems likely that insects, for example, do not rely on their abilities to walk, fly, or eat under the guidance of sensitivities that provide evidence about the effectiveness of their abilities.³¹ If not, these creatures rightly rely on such physical abilities as part of the normal good functioning of their agency, without observational evidence about their abilities. I do not want to say these creatures have an “entitlement” to rely on their physical abilities. I reserve the words ‘entitlement,’ ‘justification,’ and ‘warrant,’ for successes of reasoners—in guiding belief, or guiding intentional action. But the creatures have an analog of entitlement for their reliance on their physical abilities. Furthermore, it is likely that some human physical abilities involved in intentional action are non-observationally relied on in the normal good functioning of our agency. It is likely to be innately “wired in” that we rely on basic abilities to open and close our eyes and mouth, to swallow, and to move our limbs and neck. Non-observational reliance on basic abilities like these is very likely to be part of the normal good functioning of our agency. If so, we would have a non-observational entitlement to rely on those basic physical abilities.

registering information about our bodies—or, to their being reliably correlated with the states of affairs that they function to be correlated with, in the case of non-representational sensory capacities involved in guiding reliance. But it is worth pointing out that their reliable accuracy (or reliable correlated-ness) would be relevant to *practical* warrant only insofar as it contributes to one’s agency being a good route to action.

³¹This is of course an empirical question. I have not found science that speaks to the question.

In any case, reliance on physical abilities need not be “wired in” for the reliance to be non-observationally warranted. Even if, in the normal good functioning of one’s agency, one’s reliance on certain reliably effective physical abilities is due to guidance by observational capacities, the reliance can still be non-observationally warranted.

Like warrant for relying on perception, memory, and other capacities, warrant for relying on exercises of physical abilities is not essentially observational. The case of warrant for “wired in” reliance on physical abilities brings out this point in a clear way. If non-observational reliance on reliably effective physical abilities is an integral part of the good functioning of one’s agency, then one *thereby* has a default entitlement to rely on those abilities. The entitlement is due to the role of the physical abilities in the good functioning of agency.

The case of “wired in” reliance on physical abilities shows that the role of reliably effective physical abilities in the good functioning of agency gives one a default entitlement to rely on exercises of the abilities. Generally, one has a default entitlement to rely on reliably effective physical abilities when, and because, their exercise is an integral part of the good functioning of one’s agency. When the exercise of physical abilities is an integral part of the good functioning of agency, one’s reliance on those abilities is therein non-accidentally right with respect to the good functioning of one’s agency. The role of reliably effective physical abilities gives us a default entitlement to rely on their exercise, whether or not the reliance is guided by observational capacities like perception and proprioception.

When reliance on exercises of reliably effective physical abilities is guided by observational capacities, warrant for relying on the abilities can be overdetermined. There can be two warrants for the reliance. One warrant for the reliance depends for its force on observational capacities that register evidence that helps establish that the physical abilities are effective and can be relied on. The other warrant is due to the role of the physical abilities in the good functioning of one’s agency. That second warrant is non-observational. The second warrant

depends on the facts about one's physical abilities—whether they are in fact a part of the good functioning of one's agency—and not on whether the agent makes use of observational evidence of the effectiveness of the abilities.

Consider the physical ability involved in grasping a cup and bringing it to one's mouth. As an individual acquires the relevant abilities, she “picks up” on the fact that one possesses those abilities through perception, sensation, and so on. She may have two warrants for relying on exercises of these physical abilities. One warrant is due to the observational capacities that register evidence that one possesses those abilities. Another warrant is a non-observational entitlement, due to the role that those abilities actually play in the good functioning of the individual's agency. The latter warrant would stand without the observational evidence of the abilities' effectiveness. As long as the physical abilities play an integral role in the good functioning of her agency, she is thereby warranted in relying on them.

Being right to rely on exercises of your physical abilities is a matter of being right to make use of your physical abilities to carry out their role in the functioning of your agency. If the exercise of your physical abilities *is* an aspect of the good functioning of your agency, your reliance on your physical abilities is *therein* non-accidentally right with respect to the good functioning of your agency. If your reliance on exercises of your physical abilities is guided by perception and sensation, those observational capacities can also be an integral aspect of the good functioning of your agency. As a result, one has a default practical entitlement to rely on those observational capacities. But the physical abilities themselves are also an integral aspect of the good functioning of one's agency. Thus, one has a default non-observational entitlement to rely on the physical abilities themselves.

This account is meant to be responsive to the problem for Anscombe's view that I sketched earlier. Whether one possesses and is successfully making use of effective physical abilities is a contingent, extra-mental matter of fact. So without observational grounds for relying on exercises of one's physical abilities, the reliance would seem to be like an arbitrary,

unwarranted guess, that is at best accidentally correct. Non-observational reliance would seem to be no more warranted than a non-observational belief about the weather.

The important difference between intentional physical actions and events in the external world, like the weather, is that the former are exercises of agency. Warrant concerning physical action is different from warrant concerning the external world because our physical actions are exercises of agency. We are entitled to rely on aspects of the good functioning of our agency. Focusing on reliance on physical abilities, instead of beliefs about physical action, allows us to see how non-observational warrant concerning physical action is due to the fact that our physically acting is an integral aspect of the exercise of agency.

A non-observational belief about the weather is unwarranted because, if such a belief were true, its truth would be accidental with respect to the good functioning of one's epistemic capacities. Non-observational reliance on exercises of one's physical abilities is non-accidentally right with respect to the good functioning of one's agency because it *is* an integral aspect of the good functioning of one's agency. Doing well as an agent does not begin with assessment of one's agency. It begins with reliance on one's agency, which includes the physical abilities we make use of to execute our actions. Relying on exercises of one's physical abilities constitutes doing well as an agent when and because our physical abilities are an integral part of the good functioning of our agency. Thus, we have a non-observational default entitlement to rely on exercises of our physical abilities because in relying on our physical abilities we rely on integral aspects of our agency.

Physical abilities and practical reason

The view that we have a defeasible non-observational entitlement to rely on exercises of physical abilities finds some support in the relation between reliance on exercises of physical abilities and reliance on exercises of practical reason.

To get at the relation between practical reason and physical abilities, it will help to begin with the relation between practical reason and action. Practical reason functions to guide both intentions and actions. This point is evinced by the fact that we both act and intend for reasons, and by the fact that practical reasoning affects the rationality of both actions and intentions.

Practical reason's relation to action differs from its relation to intention. I assume that forming an intention is the conclusion in a piece of practical reasoning, just as forming a belief is the conclusion in a piece of epistemic reasoning. Forming a belief is a part of one's epistemic reasoning, and forming an intention is part of one's practical reasoning. In contrast, an action is not a part of the practical reasoning it is based on. Raising your arm, or walking around, is not a part of the reasoning your action is based on. In guiding action, practical reason guides something not contained in the guiding exercise of practical reason.

This point marks a significant difference between epistemic reason and practical reason. What epistemic reason functions to guide, namely belief, is contained in the exercise of epistemic reason. In contrast, something that practical reason functions to guide is not contained in the exercise of practical reason.

Because belief is contained in the exercise of epistemic reason, epistemic reason can carry out its belief guiding function on its own, without contributions from capacities beyond those involved in epistemic reasoning. In a similar way, practical reason can carry out its function of guiding intention on its own. But for practical reason to carry out its action guiding function, capacities involved in practical reasoning are not enough. For practical reason to successfully carry out its full action guiding function, one must *act* on the basis of one's reasoning. That requires the exercise of capacities through which one actually carries out one's actions. Without such capacities, practical reason could guide intention, but it could not guide action. Practical reason's guidance of action requires not just practical reason itself, but also other aspects of the larger action system that it functions as a part of.

Rational agency is a capacity to act on the basis of reasons, and not just an ability to form intentions on the basis of reasons. Our rational agency comprises both our capacities for practical reasoning, and capacities through which we carry out our actions. It is essential to the nature of practical reason that it functions as part of rational agency.

At least for embodied rational agents like us, a central part of practical reason's action guiding function is the guidance of *physical* action. Practical reason can successfully carry out this function only in concert with the physical abilities involved in executing physical actions—like the physical ability to put one foot in front of another and walk.

When we make use of practical reason, we rely on it not just to carry out its intention guiding function, but to carry its full action guiding function. We rely on exercises of practical reason to guide action even though practical reason guides action only as part of a larger action system. Analogously, when one relies on a steering wheel in driving, one relies on the wheel to guide the direction of the car—even though a steering wheel only guides the direction of the car as part of a larger system that includes the steering shaft, the wheels of the car, and so on. One does not rely on the steering wheel just to determine the movement of the steering shaft, or just to be turnable. One relies on the steering wheel to do its job of guiding the direction of the car, even though it cannot do that job on its own. Likewise, one relies on exercises practical reason to guide action, even though practical cannot guide action on its own.

When we rely on exercises of practical reason to carry out its full action guiding function, we must also rely on the larger action system that practical reason is a part of—including the physical abilities comprised by the larger actions system. If we do not rely on that larger action system, we can only rely on exercises of practical reason to guide intention, and cannot rely on it to guide action. Since physical abilities are integral to our larger action system, reliance on practical reason to guide action requires reliance on physical abilities.

There is a corresponding connection between *entitlement* to rely on practical reason, and

entitlement to rely on physical abilities. We have an entitlement to rely on practical reason to carry out its full action guiding function, which implies that we have an entitlement to rely on exercises of our physical abilities.

It is *prime facie* plausible that we have an entitlement to rely on practical reason to carry out its guidance functions. We do not need to have observational evidence of practical reason's effectiveness in guiding intentions to be warranted in relying on it to guide our intentions. Nor do we start out with an entitlement to rely on practical reason to guide intentions, and then have to "piece together" that it is also effective in guiding action.

I assume that we do indeed have a defeasible default entitlement to rely on practical reason to carry out its full action guiding function. If you are indeed entitled to rely on practical reason to carry out its full action-guiding function, then you are entitled to rely on your physical abilities.

To be clear, I do not mean to claim that our entitlement to rely on our physical abilities derives from the nature of practical reason. I mean to claim that our entitlement to rely on practical reason to carry out its full function implies that we have an entitlement to rely on our physical abilities. I want to briefly discuss two considerations that help to motivate this claim. The considerations do not constitute a knock-down argument, but I think they help strengthen the case for the claim.

First, the claim that we have a default entitlement to rely on practical reason is closely related to the claim that standards of practical reason have a kind of non-derivative normative force. Standards of practical reasoning concerning action have genuine normative force. I say "concerning action" to mark the fact that I have in mind rational standards that apply to action, and not just rational standards that apply to intentions. We genuinely ought to conform to them, and act on the basis of good reasons.³² Standards of practical rationality concerning action do not have this normative force because of vicissitudes of the agent's

³²Some would deny this claim. See, for example, Kolodny (2005).

personal history. For example, the explanation for why you should act on the basis of good reasons is not that you recognize that doing so is advisable as a means to your ends. Standards of practical reason concerning action apply to individuals in virtue of their being rational agents.

Now, if you are not warranted in relying on practical reason to guide action, then standards of practical reason concerning action do not have genuine normative force for you. If you lack warrant for relying on practical reason to carry out its action guiding function, it is a mistake to act under its guidance. It is not the case that you genuinely ought to act under practical reason's guidance, in conformity with standards of practical reasoning. As a result, if warrant for relying on practical reason to guide action were observational, then the normative force of standards of practical rationality concerning action would be derivative in a certain respect. Whether you are warranted in relying on practical reason would depend on whether you have made use of information indicating that practical reason is an effective guide to action. And whether standards of practical reason concerning action have genuine normative force would likewise depend on whether one has made use of this sort of information. So, if one accepts that standards of practical rationality concerning action have non-derivative normative force, one should also accept that warrant for relying on practical reason is non-observational.

Reflecting on practical reasoners with limited cognitive capacities provides a second source of support for the view that we have a non-observational entitlement to rely on practical reason. An agent can be capable of determining how she acts through practical reasoning, without being capable of meta-reasoning about whether practical reason is effective in guiding action. Nor must a practical reasoner's reliance on practical reason be susceptible to guidance by non-rational sensitivities to the effectiveness of her practical reasoning in guiding action. It seems likely that young children and some higher animals are practical reasoners of this sort. Such practical reasoners rely on practical reason "blindly."

But their reliance on practical reason to guide action is not thereby mistaken. Their reliance on practical reason is warranted even though it is not based on warranting grounds that provide information about practical reason's effectiveness. They are entitled to rely on practical reason. If basic practical reasoners are entitled to rely on practical reason, it is natural to suppose that sophisticated practical reasoners like us would have such an entitlement too. If you are a practical reasoner who is capable of effectively determining what you do through reasoning, then you are entitled to rely on practical reason to determine what you do. You are entitled independently of whether you have evidence concerning practical reason's effectiveness in guiding action.

If you are entitled to rely on practical reason to carry out its full action-guiding function, then you must be entitled to rely on exercises of the larger action system that practical reason functions as a part of.³³ Without an entitlement to rely on exercises of the larger system that practical reason functions as a part of, we would not be entitled to rely on practical reason to carry out its action guiding function. Practical reason only guides action as part of that larger system. A default entitlement to rely on exercises of practical reason means that practical reasoners do not need to examine the credentials of practical reason to be warranted in relying on it. Since practical reason only guides action as part of a larger action system, a default entitlement to rely on practical reason to carry out its full function means that we do not need to examine the credentials of the larger action system to rely on it either. Our physical abilities are a central part of that larger action system. As a result, our default entitlement to rely on practical reason to guide action means that we have a default entitlement to rely on exercises of the physical abilities involved in practical reason's guidance of action.

³³This claim about entitlements is less vulnerable to some objections than the earlier claim that reliance on practical reason to guide action requires reliance on the larger action system. To that earlier claim, one might object that one can rely on practical reason to guide action without relying on physical abilities, even though it would be a mistake. I do not believe that is the case, but the claim in the objection is compatible with the claim here about entitlement.

If there were no default entitlement to rely on the abilities through which we execute our actions, there would be a default requirement for evidence of those abilities' effectiveness to rely on them. A default requirement for such evidence would mean that it would not be reasonable to "presume" that those abilities are effective in carrying out their role in practical reason's guidance of action. We would need evidence to establish that those "executive" abilities are effective and can be relied on. That demand for evidence would mean that we do not have a default entitlement to rely on practical reason to guide action. The pressing doubt about one's executive abilities would mean there is pressing doubt about the effectiveness of practical reason in guiding action. If there were such a requirement for evidence, then perhaps one could be entitled to rely on practical reason to guide one to an appropriate *intention* about whether, for example, to go for a walk, but one would not be entitled to rely on practical reason to guide one to the appropriate course of *action*. In that case, warrant for relying on practical reason to carry out its full action-guiding function would depend on evidence that one's larger action system is effective in executing one's intentions.

Our entitlement to rely on practical reason to carry out its full function reveals that we are entitled to rely on other aspects of the action system. Our physical abilities are an important part of that larger action system. Practical reason can only do what we are entitled to rely on it to do in concert with other aspects of the action system, including our physical abilities. As a result, if we are non-observationally entitled to rely on exercises of practical reason to guide action, we must also be non-observationally entitled to rely on exercises of other aspects of our action system, including our physical abilities.

This line of thought coheres with my earlier claim that one is entitled to rely on aspects of the normal good functioning of one's agency. In rational agents, exercises of practical reason are aspects of the normal good functioning of one's agency. So are exercises of the other aspects of the larger action system, including exercises of physical abilities. Both entitlements can be thought of as entitlements to rely on aspects of the normal good functioning of one's

agency. Our entitlement to rely on practical reason to carry out its full action guiding function depends on practical reason's being part of a good route to action. As a result, the entitlement depends on the reliable effectiveness of the agent's abilities to carry out actions, including her physical abilities. Without such physical abilities, practical reason would not function as part of a good route to action.

Entitlement to rely on practical reason

I want to conclude with a discussion of the way physical abilities figure in the explanation of our entitlement to rely on practical reason. The line of thought developed just above might seem to suggest the following account of our entitlement to rely on physical abilities. We are entitled to rely on practical reason to carry out its full function because of the nature of practical reason. In that sense, our entitlement to rely on practical reason is an apriori entitlement. Because of our entitlement to rely on practical reason to carry out its full function, we are entitled to rely on the larger action system that practical reason depends on to carry out its full action guiding function, including our physical abilities. Thus, our entitlement to rely on our physical abilities is an apriori entitlement that derives from the nature of practical reason.

I do not think that our entitlement to rely on our physical abilities is apriori in this way. Our entitlement to rely on our physical abilities derives in part from the systematic effectiveness of those abilities in carrying out our actions. That effectiveness is part of what makes those abilities part of a good route to action. Whether our physical abilities are effective in carrying out our actions is not determined by the nature of practical reason.

The discussion in this chapter suggests an opposite view. Our entitlement to rely on practical reason to carry out its full action guiding function is not an apriori entitlement. We are entitled to rely on practical reason to carry out its full action guiding function partly because practical reason is a good route to action. That is, the entitlement depends on

whether practical reason is systematically effective in guiding the agent to correct action, when it is functioning normally in normal circumstances. But whether practical reason is a good route to action in that way is not determined just by the nature of practical reason. It depends on other aspects of the larger action system within which practical reason functions. If there were no larger action system, or the action system were massively inept in carrying out actions, then practical reason would not be a good route to action. Since practical reason's effectiveness in carrying out its full function is hostage to the larger action system, entitlement to rely on practical reason to carry out its full function is likewise dependent on the larger action system. Entitlement to rely on practical reason to carry out its full function depends not just on the nature of practical reason itself, but also on the effectiveness of other aspects of the action system. Thus, entitlement to rely on practical reason to carry out its full action guiding function is not apriori.

At least from a rationalist perspective, it may seem somewhat surprising that entitlement to rely on practical reason is not apriori. One strand in Kantian rationalist views about ethics and agency is that the authority of practical reason is due to the nature of reason itself. In referring to 'the authority of practical reason,' I have in mind the fact that we genuinely ought to conform to standards of practical rationality, including rational standards that concern physical action. The rationalist claim is that we ought to conform to standards of practical rationality because of the nature of practical reason.³⁴

The claim that rational norms apply to actions apriori seems to be in tension with the claim that our entitlement to rely on practical reason to carry out its full action guiding function is not apriori. Our entitlement to rely on practical reason to carry out its full function is

³⁴Kant-inspired rationalists claim that rational norms apply in the first instance to willings, and not to actions. That is certainly Kant's view. Kantians also claim that rational norms apply to physical actions, even if derivatively. Some Kant-inspired rationalists claim that rational norms apply to physical actions apriori. Barbara Herman's 1993 interpretation of Kant suggests that rational norms do not apply to any physical actions apriori. Apriori reflection can ground formal conclusions about the contours of any system of moral or rational requirements on action. But to ground substantive rational norms prescribing actions, one must make use of grounds concerning our contingent empirical natures.

a condition on the authority of practical reason. As suggested earlier, we ought to conform to standards of practical rationality that concern action only if we are warranted in relying on practical reason to guide action. My claim is that our entitlement to rely on practical reason to carry out its full action guiding function is due in part to the effectiveness of the larger action system, including the effectiveness of our physical abilities. This claim seems to suggest that the effectiveness of the larger action system is a condition on the normative force of standards of practical rationality—at least standards of practical rationality that concern action. Thus, the normative force of standards of practical rationality do not derive just from the nature of practical reason.

This line of thought seems to me more or less sound as it applies to standards of rationality that concern action. Our entitlement to rely on practical reason to carry out its full action guiding function is a condition on the normative force of standards of rational action that concern action. The entitlement is not apriori. Thus, such rational norms are not apriori.

However, it can still be the case that rational standards that concern intention are apriori. Our entitlement to rely on practical reason to carry out its full action guiding function can be factored into separate elements. We have an entitlement to rely on practical reason to guide intention. We also have an entitlement to rely on the larger action system to execute our intentions. Our entitlement to rely on practical reason to carry out its full function of guiding action comprises both of these entitlements.

The aspect of that full entitlement that concerns the larger action system is not apriori. It depends on the effectiveness of the larger action system. However, our entitlement to rely on practical reason to guide intention might still be apriori. Forming an intention is a concluding step in one's practical reasoning. So no capacities beyond practical reason itself are required for practical reason to carry out that function. The effectiveness of practical reason in guiding intention does not depend on the effectiveness of the larger action system. As a result, our entitlement to rely on practical reason to guide intention can be due just to

the nature of practical reason, and so can be apriori.

Bibliography

William Alston. *Epistemic Justification*. Cornell University Press, 1989.

Michael Anderson. Embodied cognition: a field guide. *Artificial Intelligence*, 149(1), 2003.

Julia Annas. Ancient ethics and modern morality. *Philosophical Perspectives*, 6, 1992.

Elizabeth Anscombe. *Intention*. Harvard University Press, Cambridge, MA, 1965.

Robert Audi. *Practical Reasoning*. Routledge, New York, 1989.

Robert Audi. An internalist conception of rational action. *Philosophical Perspectives*, 4, 1990.

Robert Audi. Reasons, practical reason, and practical reasoning. *Ratio*, 42(2), 2004.

Michael Ayers. Is perceptual content ever conceptual? *Philosophical Books*, 43(1), 2002.

Lawrence Bonjour and Ernest Sosa. *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues*. Blackwell Publishing, Oxford, 2003.

Michael Bratman. *Structures of Agency*. Oxford University Press, Oxford, 2007.

Bill Brewer. *Perception and Reason*. Oxford University Press, Oxford, 1999.

John Broome. Practical reasoning. In José Luis Bermúdez and Alan Millar, editors, *Reasons and Nature*. Oxford University Press, Oxford, 2002.

- Tyler Burge. Content preservation. *The Philosophical Review*, 102, 1993.
- Tyler Burge. Computer proof, apriori knowledge, and other minds. *Philosophical perspectives*, 12, 1998.
- Tyler Burge. Reason and the first person. In Crispin Wright, Barry Smith, and Cynthia Macdonald, editors, *Knowing Our Own Minds*. Oxford University Press, Oxford, 2000.
- Tyler Burge. Perceptual entitlement. *Philosophy and Phenomenological Research*, 67, 2003.
- Tyler Burge. Five theses on de re states and attitudes. In Joseph Almog and Paul Leonardi, editors, *The Philosophy of David Kaplan*. Oxford University Press, Oxford, 2009a.
- Tyler Burge. Primitive agency and natural norms. *Philosophy and Phenomenological Research*, 79(2), 2009b.
- Tyler Burge. *Origins of Objectivity*. Oxford University Press, Oxford, 2010a.
- Tyler Burge. Origins of perception. *Disputatio*, 4(29), 2010b.
- Tyler Burge. Steps towards origins of propositional thought. *Disputatio*, 4(29), 2010c.
- Judee K. Burgoon, Lesa A. Stern, and Leesa Dillman. *Interpersonal Adaptation: Dyadic Interaction Patterns*. The Press Syndicate of the University of Cambridge, Cambridge, 1995.
- David Charles. *Aristotle's Philosophy of Action*. Cornell University Press, Ithica, 1984.
- Andy Clark and David Chalmers. The extended mind. *Analysis*, 58, 1998.
- Philip Clark. The action as conclusion. *Canadian Journal of Philosophy*, 31(4), 2001.
- Randolph Clarke. Skilled activity and the causal theory of action. *Philosophy and Phenomenological Research*, 80(3), 2010.

- Antonio Damasio. *The Feeling of What Happens*. Harcourt, New York, 1999.
- Jonathan Dancy. *Ethics without Principles*. Oxford University Press, Oxford, 2004.
- Donald Davidson. Actions, reasons and causes. In *Essays on Actions and Events*. Clarendon Press, Oxford, 1980.
- Ronald de Sousa. The rationality of emotion. *Dialogue*, 18, 1979.
- Hubert Dreyfus. Intelligence without representation. *Phenomenology and the Cognitive Sciences*, 2002.
- Hubert Dreyfus. Overcoming the myth of the mental. *Topoi*, 25, 2006.
- Hubert Dreyfus. Detachment, involvement, and rationality: Are we essentially rational animals? *Human Affairs*, 17, 2007.
- Hubert L. Dreyfus and Stuart E. Dreyfus. What is morality: A phenomenological account of the development of ethical expertise. In David Rasmussen, editor, *Universalism vs. Communitarianism: Contemporary Debates in Ethics*. MIT University Press, Cambridge, MA, 1990.
- Elizabeth Ennen. Phenomological coping skills and he striatal memory system. *Phenomenology and the Cognitive Sciences*, 2, 2003.
- Philippa Foot. *Natural goodness*. Oxford University Press, Oxford, 2003.
- Harry G. Frankfurt. The problem of action. *American Philosophical Quarterly*, 15(2), 1978.
- Harry G. Frankfurt. Identification and externality. In *The Importance of What We Care About*. Cambridge University Press, Cambridge, 1988.
- David W. Franklin and Daniel M. Wolpert. Computational mechanisms of sensorimotor control. *Neuron*, 72, 2011.

- Gerd Gigerenzer. *Gut Feelings: The Intelligence of the Unconscious*. Viking, New York, 2007.
- Carol Gilligan. In a different voice: Women's conceptions of self and of morality. In Hester Eisenstein and Alice Jardine, editors, *The Future of Difference*. Rutgers University Press, New Brunswick, NJ, 1985.
- Erving Goffman. *The Presentation of Self in Everyday Life*. Doubleday, New York, 1959.
- Peter Goldie. *The Emotions: A Philosophical Exploration*. Oxford University Press, Oxford, 2000.
- Paul Grice. Intention and uncertainty. *The Proceedings of the British Academy*, 57, 1971.
- Barbara Herman. *The Practice of Moral Judgment*. Harvard University Press, Cambridge, MA, 1993.
- Barbara Herman. *Moral Literacy*. Harvard University Press, Cambridge, MA, 2007.
- Pamela Hieronymi. Controlling attitudes. *Pacific Philosophical Quarterly*, 87(1), 2006.
- Pamela Hieronymi. Two kinds of agency. In Lucy O'Brien and Matthew Soterious, editors, *Mental Actions*. Oxford University Press, Oxford, 2009.
- Rosalind Hursthouse. Arational actions. *The Journal of Philosophy*, 88(2), 1991.
- Julie T. Irish and Judith Hall. Interruptive patterns in medical visits: the effects of role, status and gender. *Social Science and Medicine*, 41(6), 1995.
- Richard L. Street Jr. and David B. Buller. Nonverbal response patterning in physician-patient interactions: a functional analysis. *Journal of Nonverbal Behavior*, 1987.
- David Kaplan. Demonstratives. In Joseph Almog and Howard Wettstein, editors, *Themes From Kaplan*. Oxford University Press, Oxford, 1977.

- Gregory Kavka. The toxin puzzle. *Analysis*, 43(1), 1983.
- Sean D. Kelly. Grasping at straws: Motor intentionality and the cognitive science of skilled behavior. In Mark Wrathall and Jeff Malpas, editors, *Heidegger, Coping, and Cognitive Science: Essays in Honor of Hubert L. Dreyfus*. MIT University Press, Cambridge, MA, 2000.
- Anthony Kenny. Practical inference. *Analysis*, 26(3), 1966.
- Niko Kolodny. Why be rational? *Mind*, 114, 2005.
- Christine Korsgaard. *Sources of Normativity*. Cambridge University Press, Cambridge, 1996.
- Christine Korsgaard. *Self-Constitution: Agency, Identity, and Integrity*. Oxford University Press, Oxford, 2009.
- Saul Kripke. *Naming and Necessity*. Harvard University Press, Cambridge, MA, 1980.
- George Lakoff and Mark Johnson. *Philosophy in the Flesh: The Embodied Mind and its Challenge to Western Thought*. Basic Books, New York, 1999.
- Pawel Lewicki, Maria Czyzewska, and Hunter Hoffman. Unconscious acquisition of complex procedural knowledge. *Journal of experimental psychology*, 13(4), 1987.
- Pawel Lewicki, Thomas Hill, and Elizabeth Bizot. Acquisition of procedural knowledge about a pattern of stimuli that cannot be articulated. *Cognitive Psychology*, 20, 1988.
- Michael Martin. The rational role of experience. *Proceeding of the Aristotelian Society*, 93, 1993.
- John McDowell. Virtue and reason. *Monist*, 62(3), 1979.
- John McDowell. *Mind and World*. Harvard University Press, Cambridge, MA, 1996.

- John McDowell. Response to dreyfus. *Inquiry*, 50(4), 2007a.
- John McDowell. What myth? *Inquiry*, 50(4), 2007b.
- Martha Craven Nussbaum. *Love's Knowledge: Essays on Philosophy and Literature*. Oxford University Press, Oxford, 1992.
- Lucy O'Brien. On knowing one's own actions. In *Agency and Self-Awareness*. Clarendon Press, Oxford, 2003.
- Derek Parfit. *On What Matters*, volume 1. Oxford University Press, Oxford, 2011.
- Sarah Paul. How we know what we're doing. *Philosopher's Imprint*, 9, 2009.
- Christopher Peacocke. Does perception have a nonconceptual content? *The Journal of Philosophy*, 98, 2001.
- Hanna Pickard. Knowledge of action without observation. *Proceedings of the Aristotelian Society*, 104, 2004.
- Alvin Plantinga. *Warrant and Proper Function*. Oxford University Press, Oxford, 1993.
- Warren Quinn. Putting rationality in its place. In Warren Quinn, editor, *Morality and Action*. Cambridge University Press, Cambridge, 1993.
- Joseph Raz. The active and the passive. *Proceedings of the Aristotelian Society*, 71:221–227, 1997.
- Joseph Raz. *Practical Reason and Norms*. Oxford University Press, Oxford, 1999.
- Joseph Raz. Reason, reasons, and normativity. In Russ Shafer-Landau, editor, *Oxford Studies in Metaethics*, volume 5. Oxford University Press, Oxford, 2010.

- Paul J. Reber and Larry R. Squire. Parallel brain systems for learning with and without awareness. *Learning and Memory*, 1, 1994.
- Erik Rietveld. Situated normativity: The normative aspect of embodied cognition in unreflective action. *Mind*, 117(468), October 2008.
- Jeffrey Robinson. Getting down to business: Talk, gaze, and body orientation during openings of doctor-patient consultations. *Human Communication Research*, 25(1), 1998.
- David-Hillel Ruben. *Action and its Explanation*. Oxford University Press, Oxford, 2003.
- T. M. Scanlon. *What We Owe to Each Other*. Harvard university press, Cambridge, 1998.
- Kieran Setiya. Practical knowledge. *Ethics*, 118, 2008.
- Kieran Setiya. Practical knowledge revisited. *Ethics*, 120, 2009.
- Susanna Siegel and Nicholas Silins. The epistemology of perception. In Mohan Matthen, editor, *The Oxford Handbook on the Philosophy of Perception*. Oxford University Press, Oxford, forthcoming.
- Ernest Sosa. The place of truth in epistemology. In M. DePaul and L. Zagzebski, editors, *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Oxford University Press, Oxford, 2002.
- Larry R. Squire. Declarative and non-declarative memory: Multiple brain systems supporting learning and memory. *Journal of Cognitive Neuroscience*, 4, 1992.
- Larry R. Squire. Memory systems of the brain: a brief history and current perspective. *Neurobiology of learning and memory*, 2004.
- Larry R. Squire. Memory and brain systems: 1969-2009. *The journal of neuroscience*, 2009.

- Jason Stanley and Timothy Williamson. Knowing how. *The Journal of Philosophy*, 98(8), 2001.
- Peter Strawson. Freedom and resentment. In Gary Watson, editor, *Free Will*. Oxford University Press, Oxford, 2003.
- Bart Streumer. Practical reasoning. In Timothy O'Connor and Constantine Sandis, editors, *A Companion to the Philosophy of Action*. Blackwell, Oxford, 2010.
- Sergio Tenenbaum. The conclusion of practical reasoning. In Sergio Tenenbaum, editor, *New Trends in Philosophy: Moral Psychology*. Rodopi, Amsterdam, 2007.
- J. David Velleman. *Practical Reflection*. Princeton University Press, Princeton, 1989.
- J. David Velleman. *The Possibility of Practical Reason*. Oxford University Press, Oxford, 2000.
- Candace Vogler. Anscombe on practical inference. In Elijah Millgram, editor, *Varieties of Practical Reasoning*. MIT University Press, Cambridge, MA, 2001.
- David Wiggins. Deliberation and practical reason. *Proceedings of the Aristotelian Society*, 76, 1975 - 1976.
- Bernard Williams. Morality and the emotions. In *Problems of the Self*. Cambridge University Press, Cambridge, 1973.
- Daniel B. Willingham. A neuropsychological theory of motor skill learning. *Psychological Review*, 105(3), 1998.
- Daniel B. Willingham, Mary J. Nissen, and Peter Bullemer. Development of procedural knowledge. *Journal of Experimental Psychology*, 15(6), 1989.

Daniel M. Wolpert, Kenji Doya, and Mitsuo Kawato. A unifying computational framework for motor control and social interaction. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 358(1431), 2003.