

Oxford Studies in Metaethics

Volume 6

Edited By
RUSS SHAFER-LANDAU

OXFORD
UNIVERSITY PRESS

OXFORD
UNIVERSITY PRESS

Great Clarendon Street, Oxford OX2 6DP

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide in

Oxford New York

Auckland Cape Town Dar es Salaam Hong Kong Karachi
Kuala Lumpur Madrid Melbourne Mexico City Nairobi
New Delhi Shanghai Taipei Toronto

With offices in

Argentina Austria Brazil Chile Czech Republic France Greece
Guatemala Hungary Italy Japan Poland Portugal Singapore
South Korea Switzerland Thailand Turkey Ukraine Vietnam

Oxford is a registered trademark of Oxford University Press
in the UK and in certain other countries

Published in the United States
by Oxford University Press Inc., New York

© The several contributors 2011

The moral rights of the authors have been asserted
Database right Oxford University Press (maker)

Crown copyright material is reproduced under Class Licence
Number C01P0000148 with the permission of OPSI
and the Queen's Printer for Scotland

First published 2011

All rights reserved. No part of this publication may be reproduced,
stored in a retrieval system, or transmitted, in any form or by any means,
without the prior permission in writing of Oxford University Press,
or as expressly permitted by law, or under terms agreed with the appropriate
reprographics rights organization. Enquiries concerning reproduction
outside the scope of the above should be sent to the Rights Department,
Oxford University Press, at the address above

You must not circulate this book in any other binding or cover
and you must impose the same condition on any acquirer

British Library Cataloguing in Publication Data
Data available

Library of Congress Cataloging in Publication Data
Data available

Typeset by SPI Publisher Services, Pondicherry, India
Printed in Great Britain
on acid-free paper by

MPG Books Group, Bodmin and King's Lynn

ISBN 978-0-19-960637-5 (Hbk.)
978-0-19-960638-2 (Pbk.)

1 3 5 7 9 10 8 6 4 2

Contents

<i>Notes on Contributors</i>	vi
<i>Introduction</i>	vii
1. Mind-Independence without the Mystery: Why Quasi-Realists Can't Have it Both Ways <i>Sharon Street</i>	1
2. How Much Realism? Evolved Thinkers and Normative Concepts <i>Allan Gibbard</i>	33
3. Parfit's Case against Subjectivism <i>David Sobel</i>	52
4. Desire-Based Theories of Reasons, Pleasure, and Welfare <i>Chris Heathwood</i>	79
5. Moral Knowledge and Experience <i>Sarah McGrath</i>	107
6. Passing the Deontic Buck. <i>Matt Bedke</i>	128
7. The Accidental Error Theorist <i>Richard Joyce</i>	153
8. Getting Real about Moral Fictionalism <i>Jonas Olson</i>	181
9. A New and Improved Supervenience Argument for Ethical Descriptivism <i>Campbell Brown</i>	205
10. Activity and Passivity in Reflective Agency <i>Paul Katsafanas</i>	219
11. Why be an Internalist about Reasons? <i>Julia Markovits</i>	255
12. Instrumental Rationality <i>Ralph Wedgwood</i>	280
<i>Index</i>	311

10

Activity and Passivity in Reflective Agency¹

Paul Katsafanas

Lately, a pair of ideas has become increasingly widespread in the literature on action theory. The first is that agents can be more and less active in the production of their own actions. For example, Michael Bratman writes,

When a person acts because of what she desires, or intends, or the like, we sometimes do not want to say simply that the pro-attitude leads to the action. In some cases, we suppose, further, that the agent is the source of, determines, directs, governs the action and is not merely the locus of a series of happenings; of causal pushes and pulls. (Bratman 2007, 91)

Bratman here claims that we need to distinguish actions that the agent actively produces from actions that spring from the agent in a more passive fashion. David Velleman concurs, claiming that “full-fledged” or “paradigm” actions cannot simply be behaviors that are caused by a belief and a desire, for the process so described

fails to include an agent. In this story, reasons cause an intention, and an intention causes bodily movements, but nobody—that is, no person—does anything. Psychological and physical events take place inside a person, but the person serves merely as an arena for these events: he takes no active part. (Velleman 2000, 123)

¹ I presented versions of this chapter at Texas Tech University, the Second Annual Rocky Mountain Ethics Congress at the University of Colorado, Boulder, and the Sixth Annual Metaethics Workshop at the University of Wisconsin, Madison. I thank the audiences for their extremely helpful comments and questions. Thanks also to David Velleman, Christine Korsgaard, David Enoch, and Lisa Damm for their insightful written comments.

These philosophers invoke the idea of an agent or self, who in some cases actively brings about her own action, rather than serving merely as a passive conduit for various desires and affects that cause acts.²

The second idea is that *reflective* or *deliberative* actions are paradigmatic cases of an agent's actively bringing about her own actions. For example, Velleman argues that actions preceded by deliberation are "full-blooded" actions or actions "par excellence" (Velleman 2000, 124 and 189), and Jay Wallace claims that full-fledged agency requires "reflective self-control," which is manifest in "choices" and "decisions" (2006, 149). These philosophers agree that reflective, deliberative actions are paradigmatic cases of agential activity.³

In this chapter, I will ask whether philosophers are correct in associating agential activity with reflective or deliberative activity. In other words, assuming that there is a distinction between activity and passivity in action, is it true that reflective actions are active and unreflective ones passive? I will argue that, on the contrary, there is no reason for maintaining that reflective or deliberative actions are paradigmatically active. In particular, I will argue that reflective or deliberative actions will seem to be active only if we work with an impoverished conception of motivation. I close by defending a new model of agential activity, which is based upon a more realistic notion of motivation and which avoids the problems faced by the current accounts.

The chapter comprises four sections. Section 1 explicates the notions of activity and passivity and provides some illustrations of the ways in which philosophers typically associate agential activity with reflection or deliberation. Section 2 asks why one would think that there is an essential connection between agential activity and reflective activity. I begin by teasing apart three distinct claims about reflective agency, which have not been clearly distinguished in the literature: that choice causes action, that motives do not determine choice, and that reflection suspends the effects of motives. Sections 2 and 3 assess these three claims, arguing that while the first and second claims are true, there are philosophical arguments and results from empirical psychology indicating that the third claim is false. Moreover, I argue that the third claim is the crucial one; its truth is necessary in order to support the idea that reflective agency is paradigmatically active. Thus, I maintain that once we distinguish these three claims, and see that the third

² Velleman continues, "What makes us agents rather than mere subjects of behavior—in our conception of ourselves, at least, if not in reality—is our perceived capacity to interpose ourselves into the course of events in such a way that the behavioral outcome is traceable directly to us" (2000, 128).

³ When an agent plays an active role in producing her own action, I will say that the agent manifests *agential activity*.

claim is false, we discover that our assumptions about reflective agency require modification and perhaps even rejection. Reflective acts are not necessarily more active than unreflective acts.

However, in Section 4, I argue that there is something correct in the idea that reflection plays an important role in agential activity. I develop a new model, which maintains that agential activity requires a certain kind of unity between one's reflective and unreflective motivations. I show that this model avoids the problems to which the traditional model succumbs. My proposed model thus provides a way of distinguishing cases of agential activity without relying on a deficient conception of motivation.

1. THE ALLEGED CONNECTION BETWEEN AGENTIAL ACTIVITY AND DELIBERATION

1.1. Activity and passivity in action

Action theory can be envisioned in two quite different ways: On the one hand, some action theorists seek to explain the distinction between *actions* and *undergoings*. Consider the following movements: my muscle spasms; my hair grows; my heart beats in order to distribute blood; my pupils dilate in order that I may see. These are all changes that, in some sense, I effect. In addition, some of these events are purposive movements. But there seems to be an important distinction between these events and genuine actions, such as walking into the kitchen in order to get a cup of coffee, typing in order to write a paper, or deciding where to take my next vacation. The events in this latter set seem attributable to me, to my own activity, in a quite different way than do dilations of my pupils or the growth of my hair. A central task in action theory is explaining this distinction between actions and mere undergoings.

However, some action theorists undertake a second and more ambitious task. They seek not only to distinguish actions from other events, but also to grade actions as more and less exemplary members of their kind. On this view, it makes sense to speak of "paradigm" or "full-fledged" actions, for there is some essential feature that is fully manifest in some actions, and less completely manifest in other actions.^{4,5}

⁴ Philosophers endorsing this second approach include Michael Bratman (2007), Harry Frankfurt (2004), Christine Korsgaard (1999 and 2009), Joseph Raz (2002), David Velleman (2000), and Jay Wallace (2006).

⁵ The two approaches to action theory may coincide, in the end. There may be some feature of action that is realizable to different degrees, so that at the highest degrees of

The feature in question is a subject of debate: it has been characterized by turns as *autonomy*, *reflective self-control*, *guidance by the agent*, *direction by the agent*, *rational control*, *agential authority*, and so forth. All of these terms have different connotations, and are analyzed in different ways. But many of the theorists who employ these terms agree on one point: these terms are meant to characterize the idea that agents can be more and less *active* in the production of their own actions.

The underlying idea can be brought out by considering a few cases:

Abe is walking along a dark path at night. Out of the corner of his eye, he suddenly sees a shape lurching through the trees, and he recoils.

Bob has been at the party for a while, and he is trying to decide whether to have another drink. He reflects on his reasons for drinking, as well as his reasons for declining to drink. He recognizes that if he has another drink, he won't be able to drive home in a safe and responsible fashion. Accordingly, he decides that he will not drink. Nevertheless, when his friend hands him a drink a moment later, he ends up drinking after all.

Claire is torn between her need to study for the exam that she has tomorrow, and her strong desire to go out to the party tonight. She reflects on her motives for each course of action, and decides that she has most reason to stay at home and study. She stays at home and studies.

These three agents seem to play progressively more active roles in the production of their movements. Abe doesn't seem to play much of a role in the production of his movement; his recoiling seems to be something that he undergoes, something that happens to him, rather than something that he performs. (On most accounts, this behavior will not even qualify as an intentional action.) Bob and Claire, on the other hand, certainly play some role in the production of their actions, but there seem to be important differences between them. Bob endeavors to act in a certain way, but finds himself yielding to or being overpowered by his desires. Claire also experiences strong desires to perform a certain action, but she resists these desires and acts in accordance with her decision. Accordingly, there is a sense in which Claire plays *more* of a role in the production of her action than does Bob. These examples illustrate one way in which we might wish to distinguish between actions that the agent plays more and less active roles in producing.⁶

fulfillment we have the paradigmatic, full-fledged actions, whereas at the lowest degrees, acting shades off into mere undergoing.

⁶ Of course, I do not intend to suggest that invoking a notion of agential activity is the *only* way in which we can mark the distinctions between these actions. There are alternatives. For example, one relevant distinction is that Bob's action is akratic, whereas

We can characterize these cases by introducing a bit of terminology. Let "agential activity" serve as the most general term for notions that are meant to pick out the agent's contribution to the production of action. Agential activity is a genus whose species are notions such as *agential control*, *agent causation*, *guidance*, *direction by the agent*, *self-government*, *self-control*, *choice*, *decision*, and so on. Then we can say that the philosophers who endorse the second conception of action theory, the philosophers who grade actions as more and less exemplary members of their kind, are claiming that different actions manifest different degrees of agential activity. Claire is presumably the most active member of the lot; Bob is somewhat less active; Abe may not be active at all.

In this chapter, I am going to assume that there is a coherent notion of agential activity. This is tendentious; some may wish to argue that the idea of agential activity makes no sense. However, a growing number of philosophers accept the idea, and I will soon give some reasons for thinking that they are right to do so. My primary topic is this: assuming that there is some kind of distinction between activity and passivity in action, how does that distinction arise? What makes some actions active, and others passive?

1.2. Activity, reflection, and deliberation

Philosophers who accept the notion of agential activity almost inevitably associate agential activity with reflective, self-conscious, or deliberative activity. As noted above, Velleman claims that deliberative actions are "paradigm cases of action" (2000, 124), and Wallace agrees, telling us that full-fledged agency is manifest in "choices" and "decisions" (2006, 149). Wallace continues,

Our motivations divide fundamentally into states of two different kinds. There are, first, motivations with respect to which we are basically passive, such as conscious desires, inclinations, yearnings, and various long-term dispositions.... Second, there are motivations that are not merely given, but that directly express our activity as agents, such as choices, decisions, and intentions to act. (Wallace 2006, 149)

Here, Wallace claims that phenomena such as choices and decisions are direct expressions of the agent's activity.

Many philosophers endorse similar ideas. Christine Korsgaard suggests that an agent is active when her action is produced by an episode of reflective, self-conscious deliberation. According to Korsgaard, "when you deliberate, it is as if there were something over and above all of your desires,

Claire's is not. Another relevant distinction is that Abe's act is not intentional, whereas Bob's and Claire's acts are intentional.

something which is *you*, and which *chooses* which desire to act on" (1996, 100). In other words, when I deliberate, I reflect on my potential actions, and choose to perform one of them. I experience my choice as expressive of me, attributable to me, in a way that my desires need not be. Richard Foley writes that "it seems clear that actions which are preceded by deliberation are paradigms of action" (1977, 59). Finally, to cite just one more example, Gary Watson claims that "practical decision, 'deciding to', is an active phenomenon" (2004, 140).

Here we have a number of philosophers who are engaged in quite different projects, and who operate with quite different sets of assumptions, yet who agree that actions produced by deliberation and reflective choice are paradigmatically active.⁷ In the following sections, I will argue that this agreement is founded on a mistake.

2. TEASING APART SEVERAL DISTINCT CLAIMS ABOUT REFLECTIVE AGENCY

2.1. The components of deliberative agency

To see why philosophers tend to assume that deliberative agency is paradigmatically active, let's begin by analyzing the notion of deliberation. The philosophers quoted above believe that there are two ways in which action can be brought about. Some actions are directly caused by the agent's motives; in these cases, episodes of choice are either absent or causally inert.⁸ Other actions are caused by the agent's choice.⁹ This view is explicit in Wallace, who writes that the "will," or the agent's capacity for choice, is "a capacity for active self-determination" (2006, 149).

⁷ To be clear, not all of these philosophers think that deliberation is a *necessary* condition for agential activity. Velleman, for example, argues that agential activity can be manifest in unreflective, skillful actions, such as playing a piano or riding a bicycle. However, all of these philosophers maintain that deliberative action is a paradigmatic case of agential activity, for two reasons. First, in ordinary conditions, deliberative action is *sufficient* for agential activity; if the agent deliberates and acts in accordance with this deliberation, she will manifest agential activity. Second, these philosophers claim that it is by analyzing the structure of deliberative agency that we will understand other types of agential activity (cf. Velleman 2004, 281–2).

⁸ I will use the term "motive" in a very encompassing way, to refer to any folk-psychological state that can cause action. Examples include desires, affects, moods, and emotions.

⁹ That is, the episode of choice is *one* of the causes of the action. It need not be the only cause.

In short, these philosophers assume that in normal circumstances, by choosing to A, an agent makes it the case that she will A. Of course, this isn't always true. Sometimes, my choices are overpowered by strong motives. I may choose to have only one drink at the party, and find myself unable to refuse a second or a third. I may decide that I need to spend three hours grading student papers, and find myself distracted by another task within the first hour. So, strong motives sometimes seem to trump choice. Additionally, the world can thwart my attempt to carry out my choice: I may choose to eat dinner at a favorite restaurant, yet arrive to find the restaurant closed. I may choose to drive to my office, yet find that my car won't start. In each case, the world foils my chosen course of action. However, these cases are atypical: in each case, we need to provide an explanation of *why* the choice didn't issue in the chosen action. Either the motives are unusually strong and overpowering, or the world is unexpectedly recalcitrant. Thus, these cases support the claim that in the usual case, we are capable of acting as we choose to act. If I choose to eat a grapefruit for breakfast, I will probably end up eating it; if not, we need some explanation of why my choice did not eventuate in action. Thus, the following claim seems true:

(Choice) Typically, if I am faced with two actions that it is possible for me to perform, A-ing and B-ing, and I choose to A, then I will A.^{10,11}

How does the agent's choice relate to her motives? In other words, if choice typically determines action, how do motives relate to choice?

In the tradition, we can find two views: either motives determine what we choose, or they do not. According to the first view, while a mental episode of choosing or deliberating may occur in me, the course of these events, and their outcome, are determined by the motives that I have.¹² The

¹⁰ Obviously, complications would arise if we attempted to spell out what is meant by "actions that it is possible for me to perform." As the focus of this chapter is elsewhere, I will not delve into these issues. All that matters, for our purposes, is that there are cases of the following sort: it is possible for me to have either grapefruit or cereal; I choose to have grapefruit; this choice eventuates in my having grapefruit.

¹¹ For philosophical arguments and empirical evidence in support of the claim that conscious decisions typically cause actions, see Mele (2009, chapter 7). Mele writes, "there is powerful evidence for the truth of the following thesis: the fact that an agent consciously decided to A or had a conscious intention to A sometimes has a place in a causal explanation of the corresponding overt intentional action" (2009, 144).

¹² A version of this view is defended by Hobbes, who writes, "When in the mind of man appetites and aversions, hopes and fears, concerning one and the same thing, arise alternately; and diverse good and evil consequences of the doing or omitting the thing propounded come successively into our thoughts; so that sometimes we have an appetite to it, sometimes an aversion from it; sometimes hope to be able to do it, sometimes despair, or fear to attempt it; the whole sum of desires, aversions, hopes and fears,

philosophers who associate agential activity with deliberative action reject this picture, instead adopting the second view. We can put the second view as follows:

(Non-determination) Choice is not determined by motives; an agent's motives could be the same, and yet she could choose differently.

For example, my decision whether to have grapefruit or cereal for breakfast is not determined by my desires; on two consecutive mornings I might have exactly the same desires, and yet choose grapefruit one day and cereal the next.

Wallace and Velleman are explicit about their commitment to Non-determination,¹³ but we can see that anyone who wants to maintain a distinction between actions caused by choice and actions caused solely by motives will be committed to it as well. After all, if choice were simply determined by motives, then there wouldn't really be two types of actions. Chosen actions would *also* be determined in a passive fashion, issuing merely from the agent's motives.

What reasons might we have for accepting Non-determination? A version of this claim was proposed by Leibniz, and defended by Locke and Kant. Leibniz famously claimed that motives "incline without necessitating."¹⁴ On this view, motives do not directly cause action. Rather, while motives incline or dispose us to perform certain actions, we are free to *choose* whether to act as the motives incline us to act.

continued till the thing be either done, or thought impossible, is that we call deliberation" (*Leviathan* VI). He continues, "In deliberation, the last appetite, or aversion, immediately adhering to the action, or to the omission thereof, is that we call the will..." (*Leviathan* VI). For a more sophisticated version of this position, see Schopenhauer, *Prize Essay on Freedom of the Will*, pp. 36–7.

¹³ Wallace writes "volitional motivations [e.g., choices] are *independent* of our given desires... What an agent chooses or intends to do is not a function of the given desires to which the agent is subject at the time" (2006, 150). Velleman's view is more complex. He argues that when an agent is faced with the choice of whether to A or B, the agent can sometimes make it the case that he will A by forming a self-fulfilling belief that he will A. The self-fulfilling belief need not be determined by the motives that the agent is reflecting upon, so a version of Non-determination is true. However, the self-fulfilling belief will be determined by a background motive which, Velleman argues, is omnipresent—the motive of self-understanding. This background motive determines which particular self-fulfilling belief the agent formulates. So, while the motives reflected upon do not determine choice, a background motive does determine choice. Thus, Velleman's view is a qualified version of Non-determination. For this view, see "Epistemic Freedom" in Velleman (2000).

¹⁴ See Leibniz's Fifth Letter to Samuel Clark, collected in Leibniz and Clark (2000, 37).

Locke and Kant offer a sophisticated defense of this position. They claim that when a self-conscious agent reflects on potential actions, she is committed to viewing her deliberation as capable of *suspending the effects of her motivational states*. Locke writes that the mind has "a power to *suspend* the execution and satisfaction of any of its desires." The mind can "consider the objects of [these desires]; examine them on all sides and weigh them with others. In this lies the liberty that man has" (Locke 1975, 263).

Kant endorses a similar model of deliberation, writing that human choice "can indeed be *affected* but not *determined* by impulses... *Freedom of choice* is this independence from being determined by sensible impulses" (*Metaphysics of Morals* 6: 213–14). Christine Korsgaard describes the Kantian model of deliberation as follows:

Our capacity to turn our attention to our own mental activities is also a capacity to distance ourselves from them, to call them into question... I desire and I find myself with a powerful impulse to act. But I back up and bring that impulse into view and then I have a certain distance. Now the impulse doesn't dominate me and now I have a problem. Shall I act? Is this desire really a *reason* to act? (Korsgaard 1996, 93)¹⁵

Locke and Kant are making a claim about motivation: by self-consciously reflecting on a motive, we can distance ourselves from this motive, thereby making the motive cease to "dominate" us, or "suspending" the motive. Put differently, reflecting on a motivational state suspends the efficacy of the motive. So, in *unreflective* action, we may simply be caused to act by our strongest motive. But, in *reflective* action, our motives operate as mere inclinations, which are incapable of causing us to act.

Thus, the primary point that Kant and Locke are making can be put this way:

(Inclination) In deliberative agency, motives incline without necessitating. The agent's motives could be the same, and yet she could choose differently.¹⁶

Further, Kant and Locke support Inclination with a claim about deliberation: they claim that when an agent reflects on her motives, she suspends the influence of these motives, and is then free to assess the rationality of acting as the motives suggest. We can put this claim as follows:

¹⁵ Compare Tamar Schapiro, who writes, "According to the Kantian picture of action, our inclinations need not determine what we do. They influence us, but we have the capacity to decide, freely and rationally, whether or not to act on them" (2009, 229).

¹⁶ Inclination differs from Non-determination in only one respect: Inclination includes a claim about how motives operate.

(Suspension) When an agent reflects on her motives for A-ing, she suspends the influence of these motives upon her assessment of whether there is reason to A.¹⁷

Although the details will have to be worked out, I think we can already see why Kant and Locke wish to distinguish reflective, deliberative actions from unreflective, non-deliberative actions.¹⁸ If the claims about Suspension and Inclination were true, then the reflective agent would have the capacity to rise above her motives, to cancel their effects and choose in complete independence of them. There would, then, be good reason for distinguishing unreflective actions, which are produced by the agent's motives, from reflective actions, which are produced in independence of the agent's motives.¹⁹

Versions of this Kantian/Lockean model of choice have become predominant in the literature on action, and indeed seem to be implicitly or explicitly present in each of the accounts reviewed in Section 1. In the following sections, I will ask whether the Kantian/Lockean model is tenable. I will argue that while Inclination is true, Suspension is false. In addition, I will argue that Suspension is required to support the idea that

¹⁷ We could state Suspension in a more modest form, by adding the word "typically" to the beginning of the definition. I will return to this point in Section 3.1.

¹⁸ There is a potential objection to my claim that Kant endorses Suspension. Kant often remarks that agents can be mistaken about how they are being motivated. For example, Kant claims that a reflective agent can believe that he is acting on duty, but actually be acting upon self-interest. This might seem to indicate that Kant actually rejects Suspension: for, in this case, the agent reflects on his motives, takes himself to be acting on one of them, but is actually acting on another. Reflection therefore seems not to suspend the effects of the latter motive (self-interest). (Thanks to Chris Korsgaard for raising this objection.) However, I think these kinds of example are compatible with Kant's acceptance of Suspension. I intend Suspension to be read as follows: when an agent reflects on a set of motives, she suspends the effects of *those motives*. Kant does not think that reflection suspends the effects of *all* of the agent's motives—in particular, reflection does not suspend the effects of the motives that the agent is not reflecting upon. This gives us a way of interpreting the passages on self-interest: Kant is claiming that background motives, upon which the agent is not reflecting, can influence the agent, and thereby make it the case that the agent is actuated by self-interested motives even when he takes himself to be acting on duty. (An example: I can reflect on the fact that helping Bob is the right thing to do, and neglect to attend to the fact that I have a strong desire to ingratiate myself with him; this latter motive, unnoticed, can influence my action. I take myself to be acting on the motive of duty, whereas I am actually influenced by a self-interested motive.) If this is the correct interpretation, then it is compatible with my attribution of the Suspension claim to Kant.

¹⁹ This point is often emphasized in discussions of free will. For example, Robert Kane writes, "Free will . . . is the power of agents to be the ultimate creators or originators and sustainers of their own ends or purposes . . . when we trace the causal or explanatory chains of action back to their sources in the purposes of free agents, these causal chains must come to an end or terminate in the willings (choices, decisions, or efforts) of the agents, which cause or bring about their purposes" (Kane 1996, 4).

deliberative actions are paradigmatically active. Thus, the Kantian/Lockean model of agential activity will have to be rejected.

2.2. Is Inclination true?

Let's begin with Inclination. Should we accept the Kantian/Lockean claim that we can choose in independence of our motives? If we reflect on the experience of deliberation, I think it will provide some support for the Kantian/Lockean analysis. Take my choice of what to have for breakfast. Here, I reflect on and assess my desires for cereal and grapefruit. Kant and Locke are claiming that when I reflect on these desires, I will have an experience of distance from them: I will experience myself as free to act on either of the desires. Richard Holton offers the following characterization of choice, which makes the point nicely:

Our experience tells us that choice is not determined by our beliefs and desires, or by any other psychological states—intentions, emotions, etc.—to which we have access. These could be the same, and yet we could choose differently. (Holton 2006, 15)

As a characterization of the phenomenology of choice, this point seems accurate. Though there are cases in which desire seems to compel choice, such as addictions, pathologies, and cases of extreme emotion, these cases are atypical. Typically, when we decide what to have for breakfast, or how to spend our evening, or what career to pursue, it feels as if we are free to choose between competing options.

Of course, the phenomenology could be misleading.²⁰ However, Kant points out that there is a powerful reason for taking the phenomenology at face value: if there is to be any such thing as deliberation at all, then agents must be capable of choosing in a way that is not determined by their motives. Kant argues that I not only experience my choices as undetermined by my motives; in addition, there is a sense in which I am *committed* to viewing my choices as undetermined. This is one way of putting Kant's point about the necessity of assuming freedom from the practical

²⁰ Perhaps the most obvious worry about the phenomenology is that it seems to view deliberation as an uncaused cause. But don't our choices have to be determined by *something*? And isn't the idea that motives incline without necessitating a denial of this? No—as a number of philosophers have noted, the claim that our choices are not determined by our *motives* does not imply that our choices are not determined by *anything* (cf. Holton 2006). One could deny that our motives determine choices, while still maintaining that other factors determine our choices. So one can divorce Inclination from general concerns about causal determination. One worry about taking the phenomenology at face value therefore disappears.

standpoint: deliberating about whether to A commits one to viewing oneself as having a say in whether one will A. If one's action were simply determined independently of deliberation—for example by one's motives—then one would not have such a say. So deliberation presupposes its own efficacy.

To see why the reflective agent is committed to viewing his reflection as undetermined by his motives, consider the difference between deliberating about what to do and thinking about what someone else should do. Suppose I consider whether my neighbor should act on his desire for another serving of ice cream. In engaging in this kind of reflection, I do not take my thoughts to have any bearing on what the neighbor will actually do. But this is quite different from the ordinary case of first-person deliberation about action. When I deliberate about whether I should act on a desire for ice cream, I take my answer to settle the question of what I *will* do. I do not regard my reflection as a passive, theoretical inquiry that has no bearing on its subject matter; I do not regard myself as making predictions while waiting to see how things turn out. Rather, if I decide that I should act on some desire, I also take this to settle the question of whether I will.²¹

That is why I am committed to viewing my deliberative reflection as suspending the workings of desire: if I did not view my deliberation as settling the question of what I will do, then my deliberation would be exactly analogous to the way in which I might think about what my neighbor should do. In other words, it would not be *deliberation* at all.

This Kantian argument establishes that we must conceive of our own deliberation as proceeding in a way that is not determined by the motives upon which we are deliberating. Moreover, we must conceive of this deliberation as eventuating in action. Of course, there could still be a mismatch between the way that we are compelled to conceive of deliberation, and what actually happens when we deliberate. But Kant's argument establishes a very powerful point: if there is to be any such thing as genuine deliberation—if anything in the world answers to our concept—then deliberation must proceed in a way that is not determined by motives. Inclination must be true.²²

²¹ Of course, any number of things can prevent my judgment that I should A from leading to my A-ing: weakness of will, unforeseen circumstances, accidents, failures. So the claim is not that when I judge that I should A, I then A. Rather, the claim is that when I judge that I should A, I take this judgment to settle the question of whether I will A.

²² Notice that I am *not* claiming that in order for agents to engage in *what they take to be* deliberation, Inclination must be true. As I note above, there could be a mismatch between the way in which we conceive of deliberation and what actually happens when we deliberate. My claim is simply that in order for deliberation *to be what we take it to be*, Inclination must be true. Thanks to David Enoch for pressing me on this point.

So we have two reasons for accepting Inclination. First, it is supported by the phenomenology. Second, the Kantian argument shows that it is an ineradicable feature of our conception of deliberation. Not surprisingly, these arguments fall short of a deductive proof that Inclination is true. They do, however, give us good reason for assuming that this claim is true. I think the burden of proof is therefore very much on those who wish to deny the claim. Thus, for the remainder of this essay, I will assume that Inclination is true.

3. SUSPENSION AND THE NATURE OF MOTIVATION

Having granted Inclination, let's now turn to Suspension. I will argue that Suspension is false. The problem with the claim that deliberation suspends one's motives is that it is ambiguous; it fails to specify what is meant by "suspending" one's motives. This ambiguity arises from a failure to distinguish two models of the way in which motives can operate: a simplistic model and a sophisticated model. According to the simplistic model, motives operate as mere pushes and pulls. Below, I will argue that if this were the *only* way in which motives operated, then Suspension would be true. However, I argue that motives also operate in a more sophisticated manner: motives influence reflection itself. I explicate and defend this model of motivation, showing that it is well supported both by philosophical considerations and by results from empirical psychology. Moreover, I show that the sophisticated model of motivation entails that Suspension is false.

3.1. Two ways of being moved

With this in mind, let's examine two ways in which attempts at deliberative suspension can fail. There appears to be a familiar type of failure: an agent's attempt at suspension can be *overpowered*. Consider Harry Frankfurt's description of such a case:

Sometimes people are unsuccessful even in strenuously conscientious efforts to avoid being moved into action by desires that they would prefer to be motivationally ineffective. For instance, someone may act out of jealousy, or out of a desire for revenge, although he disapproves of these motives and would strongly prefer that he not be driven by them. Unhappily, as it turns out, their force is too great for him to withstand; and *in the end he submits to it*. Despite his resistance, the unwelcome desire is effective in moving him to act. (Frankfurt 2004, 19; emphasis added)

Frankfurt describes a self-conscious, reflecting agent who is confronted with an attitude of which he disapproves. The agent struggles to avoid being moved by this attitude, but in the end, he "submits" to it; he is, perhaps, carried along by its force.

What kind of process is Frankfurt envisioning, when he imagines the agent confronting and submitting to his motive? The picture seems to be as follows:

1. A motive tempts the agent to pursue some end.
2. The agent reflects on this motive and his potential actions, considering whether he wants to act on the motive.
3. The agent decides not to act on the motive.
4. The agent struggles to avoid being caused to act by the motive.
5. The agent's resistance is overcome, and he is caused to act by the motive.

There are cases that seem to fit this description. Addictions provide good examples: an alcoholic struggles to resist his desire for another drink, but finds himself picking up the bottle once more. Perhaps cases of this sort occur even in more mundane, everyday situations: a decadent dessert lies before me on the table, and I crave it. I disapprove of this craving, decide not to eat the dessert, and struggle to resist the temptation. I try to put the thought of dessert out of my mind. But the temptation is too much for me: the thought of eating keeps recurring, my eyes are continuously drawn to the dessert. Eventually, I give in and pick up the spoon.

If this sort of case is a real possibility, then Suspension, in its unqualified form, would be false: there would be at least some cases in which deliberation failed to suspend the influence of motives. However, proponents of Suspension could respond in two ways. First, they could admit that this kind of case is possible, but point out that it is atypical, occurring only when the motive in question is unusually strong. Accordingly, they could qualify Suspension by claiming that it is true only in cases in which the motive is not unusually strong. Second, they could argue that Frankfurt simply mischaracterizes the case. In the jealousy and dessert cases above, I submit to the desire not in the sense that it overpowers me, but in the sense that I *decide to give in to it*, as I might give in to a pesky child, when resistance no longer seems worth the effort. In short, I am not *overpowered* by the desire; rather, I at first resist, but eventually *consent* to be moved by it.²³ If this is right, then Frankfurt's examples are compatible with Suspension.

²³ There are cases in which the person literally has no control over the behaviors that his body performs. For example, in cases of anarchic hand syndrome, the person's hand engages in goal-directed behaviors over which the agent has absolutely no control. But

Examples of the above sort do not impugn Suspension, then. However, I am going to argue that there is a different way in which an agent can be moved, unwillingly, by an attitude. I will show that this second form of motivation is far more problematic for proponents of Suspension.

Let's return to the example of jealousy. Frankfurt considers the case in which an agent clear-sightedly reflects on his jealousy and tries to resist its grip. But surely one of the most distinctive features of our attitudes is that *they influence reflection itself*. Part of what it is to be in the grip of jealousy is to see reasons for jealousy everywhere: in the fact that Melissa got home a few minutes later than usual; in the fact that she got off the phone rather quickly last night; in the fact that she is a bit quiet tonight. Accordingly, jealousy and other attitudes can move an agent not simply by overpowering his capacity to resist their pull, but by influencing his judgment. A jealous agent's attention will be drawn to certain features of his environment that another agent would scarcely notice. A jealous agent's trains of thought will return to details that another agent might regard as inconsequential. A jealous agent's deliberative process itself can be influenced by these attitudes; they can incline him to draw conclusions that are not supported by the evidence, to give excessive weight to certain features, and so on. All of this may occur without the jealous agent's recognizing that it is occurring.

Precisely because attitudes influence reflective thought, agents often fail to grasp the ways in which attitudes are motivating them. An agent who succumbs to jealousy is not always an agent who consents to be moved by his jealousy; indeed, an agent who succumbs to jealousy need not even *recognize*, much less consent to, a fully formed attitude of jealousy. More often, the jealous agent will struggle to resist the jealousy, but succumb to it in subtler ways. The attitude influences the agent's reflective thought itself: the agent experiences herself as having a reflective distance from the attitude, as scrutinizing the attitude and asking herself whether there is a reason to act on it; but, all the while, the attitude influences the agent's reflective thought in ways that she does not grasp.

This type of influence is easiest to detect when we look back at an action retrospectively. A person can be dissatisfied with his past actions not because he submitted to or was overcome by a recalcitrant attitude, but because his attitude blinded him, leading him to have a restricted, impoverished, or distorted conception of the options that were open to him. Looking back on my jealous spat with Melissa, the problem was not that I yielded to jealousy: the problem was that, in the grip of jealousy,

I see no reason to think that the case of desire is in any way analogous to these pathological cases.

I made so much of so little. The problem was that I saw my rage as *warranted* by the fact that Melissa got home a few minutes late. I now see that the rage was entirely unwarranted, that I was driven to rage in a way that I did not comprehend. In this way, an agent can act reflectively, yet still be moved by attitudes that operate in the background.

So we have distinguished two ways in which Suspension could fail:

(Motives as *forces*) Motives, if they operate as brute forces, could cause us to act in a way that we choose not to act.

(Motives as *influences*) Motives could influence our reflection in such a way that even when we attempt to achieve a distance from them, we fail to do so.

In cases of influence, motives operate *through* reflection, rather than independently of it. In a moment, I will show how these claims bear on the idea of deliberative suspension. First, though, let's explicate the idea that motives might operate as influences.

3.2. Exploring the way in which motives can operate as influences

There are at least three ways in which motives can act as influences upon reflection itself:

- (i) Motives can affect perceptual saliences.
- (ii) Motives can affect the way in which we conceive of our circumstances, our reasons for action, and the potential actions that lie open to us.
- (iii) Motives can affect the course of deliberation itself.

Influence of type (i) is perhaps the most obvious: motives affect perceptual saliences. When I'm hungry, my attention tends to be drawn to food. When I'm angry, my attention tends to be drawn to features of my environment that might justify or perpetuate the anger. The effects will probably be proportional to the strength of the motive: mild hunger won't generate as much effect on perceptual saliences as extreme hunger. I take it that this is uncontroversial and familiar.²⁴

²⁴ Indeed, the connection between motivation and perceptual salience is tight enough that some philosophers have attempted to analyze desire in terms of salience. For example, Mark Schroeder argues that "for *X* to have a desire whose object is *P* is for *X* to be in a psychological state grounding the following disposition: when for some action *a* and proposition *r* believed by *X*, given *X*'s beliefs *r* obviously helps to explain why *X*'s doing *a* promotes *p*, *X* finds *r* salient, and this tends to prompt *X* to do *a*, and *X*'s attention is directed toward considerations like *r*" (Schroeder 2007, 156–7). While I think this analysis is promising, my argument requires only the weaker claim that

Type (ii) influence occurs when motives affect the way in which we conceive of our circumstances, our reasons, and our potential actions. This requires more explanation. First, we need to distinguish the agent's circumstances from the way in which the agent *describes* or *conceives* of his circumstances. Anyone who reflects on a case will be entertaining some description of the facts of the case; this description will be partial, couched in these terms when it might have been couched in those terms. Moreover, this description will influence, constrain, and sometimes even determine the way that the person goes on.

This shows up even in relatively simple cases: until you stop seeing the geometrical problem in *this* way, you won't be able to solve it. It is also obvious in the case of discussion: often, changing someone's mind is achieved by getting him to use different descriptions of the same facts ("Don't think of it as stealing, think of it as taking a little something from an incredibly rich, greedy corporation"; "Don't think of it as making a joke, think of it as hurting Tom's feelings").

Stuart Hampshire makes this point quite clearly:

It is misleading to speak of 'the facts of the situation' in such a way as to suggest that there must be a closed set of propositions which, once established, precisely determine the situation. The situations in which we must act or abstain from acting, are 'open' in the sense that they cannot be uniquely described and finally circumscribed. Situations do not present themselves with their labels attached to them... (Hampshire 1949, 476)

In a footnote to the above passage, he adds:

The word 'fact', here as always, is treacherous, involving the old confusion between the actual situation and the description of it; the situation is given, but not 'the facts of the situation'; to state the facts is to analyze and interpret the situation. And just this is the characteristic difficulty of actual practical decisions, which disappears in the text-book cases, where the 'relevant facts' are pre-selected. (Hampshire 1949, 476)

Descriptions of situations, just in virtue of the fact that they characterize the situation in a determinate way, involve simplification and incompleteness. Certain details are emphasized at the expense of others. In this sense, descriptions are inescapably partial.²⁵

motives *generate* saliences; I am not committed to any particular analysis of what motives *are*.

²⁵ Frederic Schick also remarks on this phenomenon. Schick argues that an adequate psychological account of action will make room not only for beliefs and desires, but also for the way in which the agent "understands," "perceives," or "frames" her situation. See Schick (2003, 61 *et passim*).

Psychologists have studied this phenomenon for a number of years under the heading of "framing." Researchers have demonstrated that the way in which a situation is described has a profound impact on the decisions that agents make. For illustrative purposes, I will mention the classic example of this phenomenon: Kahneman and Tversky's "dangerous disease" case (Kahneman and Tversky 1981).

Case 1: There has been an outbreak of a dangerous disease. Doctors can adopt treatment Program A or treatment Program B, but not both. If Program A is adopted, 200 people will be saved. If Program B is adopted, there is a 1/3 probability that 600 people will be saved, and a 2/3 probability that no one will be saved.

Case 2: There has been an outbreak of a dangerous disease. Doctors can either adopt treatment Program A or treatment Program B, but not both. If Program A is adopted, 400 people will die. If Program B is adopted, there is a 1/3 probability that nobody will die, and a 2/3 probability that 600 people will die.

There is no factual difference between these two cases. Rather, they simply state the facts differently: the first case focuses on the number of people who will be saved, whereas the second one focuses on the number of people who will die. However, researchers obtained strikingly different results when asking respondents which scenario they would prefer. When presented with Case 1, 72% preferred Program A and 28% preferred Program B (sample size was 152). When presented with Case 2, 22% preferred Program A and 78% preferred Program B (sample size was 155).

The researchers demonstrated that this sort of bias is ubiquitous: the way in which we "frame" or conceive of our situation has a dramatic effect on our deliberation and choice (for further details, see Kahneman and Tversky 1981). With these points in mind, we can now explicate type (ii) influence. Our motives can manifest themselves by influencing the particular descriptions or conceptions that we employ, for the way in which we frame or conceive of situations depends, in part, on the motives that we have.

The most obvious way in which motives influence conceptualizations or descriptions is by affecting perceptual saliences: the way in which I describe a situation will be strongly influenced by facts about which features of my environment are salient to me.²⁶ Thus, a recent survey of the psychological work on the relationship between emotion and cognition concludes that

emotion may influence cognition [by] modulating which information in the environment reaches awareness. At any given time, we are bombarded with sensory input. Only a portion of this input is available for cognitive awareness . . . A number

²⁶ Thus, type (ii) influence occurs, in part, through type (i) influence.

of psychological studies have confirmed that attention and awareness may be influenced by the emotional content of a stimulus. (PHELPS 2005, 70)

For example, when an agent is hungry, she is more likely to be aware of information that relates to the presence of food in her environment, and as a result she is more likely to conceptualize her situation in terms relating to food. Or, to return to the earlier example, if an agent is jealous, he is more likely to conceptualize his environment in terms of jealousy.

However, there is evidence that motives have an even more pervasive effect on the way in which situations are conceptualized. I will cite just two examples. First, it has been demonstrated that the perceived distance between an agent and a goal increases as the effort associated with walking to the goal is increased. For example, the perceived distance increases when the agent is wearing a heavy backpack, or when the agent is tired (Witt et al. 2004). Second, studies have shown that participants' characterizations of other individuals are influenced by their moods. For example, subjects who watch a sad movie and are then asked to characterize a given person as likeable or unlikeable tend to characterize him as unlikeable, whereas subjects who watch a happy movie tend to characterize the same individual as likeable (Forgas and Vargas 2004, 357). As these examples demonstrate, motives exert a significant influence on the ways in which situations are conceptualized.²⁷

In sum, type (ii) influence occurs when motives affect the starting points of deliberation, by influencing our conceptualizations or descriptions of the case. This brings us to type (iii) influence: motives can also affect the *process* of deliberation. In other words, motives can affect the movement from a description of the situation to a conclusion about what to do.

There is a variety of work on individuals' desires to reach specific conclusions about themselves or others; much of this, in philosophical writing, is placed under the heading of self-deception. For example, I don't want to recognize that my lover is cheating on me, and I selectively process information to that end: I discount evidence of cheating, and heavily weight evidence of fidelity. Or, I want to view myself in a positive light, and end up downplaying or ignoring certain aspects of myself.

However, this phenomenon is hardly restricted to self-deception. There is a wealth of psychological research on the way in which motives influence reasoning and thinking. It is well known that motives or affects influence

²⁷ For a nice overview, see Forgas and Vargas (2004), which provides a host of evidence that "mood [or, more generally, affect] can have both informational and processing effects on cognition. Informational effects occur because mood influences the content of cognition (*what* people think). Processing effects occur because mood influences the process of cognition (*how* people think)" (Forgas and Vargas 2004, 351).

judgments by selectively influencing attention, memory, and association (Bower 1981; Clark and Waddell 1983; Isen 1984, 1987). In general, this happens in two ways. First, as we saw above, motives will "affect reasoning by affecting which information will be considered in the reasoning process" (Kunda 1990, 486). Second, motives can influence the reasoning process itself. For example, motivation impacts "evidence evaluation": people evaluate information that supports or contradicts positive self-evaluations in an interesting way, giving more credence to information that supports a positive self-conception, and less credence to information that threatens it. Motivation also impacts "information search": we are motivated toward "decreased processing and quick acceptance of favorable evidence, and increased processing and hesitant acceptance of unfavorable evidence" (Molden and Higgins 2005, 299). There is also a nice body of evidence showing that motives influence "memory search"—that is, motives influence which memories come to mind. For example, Sanitioso et al. (1990) induced one group of participants to view introversion as a desirable character trait, and another group to view extroversion as a desirable character trait. When asked to recall memories, the participants in the first group were significantly more likely to recall memories of their own introversion.

In sum, as the psychologists Molden and Higgins put it in a recent survey of the literature on motivated thinking, motivation works by "directing people's cognitive processes (e.g., their recall, information search, or attributions) in ways that help to ensure that they reach their desired conclusions" (Molden and Higgins 2005, 297). Moreover, several studies have shown that increased deliberation and reflection actually *increase* the degrees of these effects (Feshbach and Singer 1957; Forgas 1994).

Much more could be said on the particular ways in which each of these types of influence occurs. However, this brief overview will suffice for our purposes. The important point is that a wealth of psychological research supports a point that should be clear merely from a careful and judicious analysis of the phenomenology of motivation: motives do not manifest themselves simply as pushes and pulls, whose urgings we can choose to resist. The effects of motives are far subtler: our motives manifest themselves as colorings of thought and deliberation. They influence both the starting points and the process of deliberation itself, and thereby influence the outcome of deliberation (i.e. the choice). Put simply, motives operate *through*, rather than *independently of*, our rational capacities.

3.3. The ambiguity of Suspension

The prior two sections argued that motives operate *through* reflection and deliberation: part of what it is to have a motive is for one's reflection to be altered in a certain way. Let's ask how these claims bear on Suspension:

(Suspension) When an agent reflects on her motives for A-ing, she suspends the influence of these motives upon her assessment of whether there is reason to A.

Kant and Locke make two points here. First, they claim that reflection suspends the effects of the agent's motives, as we've seen above. Second, they claim that once the agent has begun reflecting on a motive, she can assess the rationality of acting on it without influence by the motive.

Let's start with the first point. I think we can agree that once I have begun reflecting on a motive, the motive very rarely acts as a brute force compelling me to act. In other words, cases of motives operating as overpowering forces—cases of the sort that Frankfurt describes—are at best rare. However, we have seen that motives can operate as influences upon reflection itself. Thus, the motive can continue to operate on me even as I am examining it. The jealous agent, reflecting on her jealousy, will see the jealousy as warranted, and she will see it as warranted precisely because she has *failed* to suspend its influence.

This brings us to the second point. We can also agree that the reflective agent typically looks for a reason to act on the motive. However, given that motives influence reflective thought, a problem arises: the reasons that the agent finds will be products of his motives. For example, when in the grip of jealousy, reflective assessment of one's jealous motives will typically vindicate these motives, precisely because the jealousy will manifest itself by inclining the agent to see jealous responses as warranted by the situation at hand. In other words, motives will affect the agent's perception of reasons.

Thus, it is a mistake to think that *either* the agent is pushed to action by some overpowering motive, *or* the agent acts reflectively in a way that is not determined by his motives. There is a third option: the agent can scrutinize his motives, decide that there is a reason to act on them, and yet, all the while, be in thrall to some motive. The effects of the motive needn't be construed as pushes and pulls that force an agent to act; rather, the motive moves the agent by influencing the agent's perception of reasons, inclining the agent to see action that fulfills the motive as rationally warranted.

Once we recognize that there are two ways of being moved (via force or via influence), I think we can see that Suspension is false. Grant that our motives rarely act as brute forces compelling us to act. If that were the *only* way in which motives could move us, then (some version of) Suspension would be true, for deliberation often does suspend this kind of influence.

However, there is another way in which motives can move us: they can pervasively influence the course of reflection itself, in particular by inclining us to see acting upon the motives as warranted. Reflection does not suspend this type of influence. On the contrary, reflection is often the *vehicle* for this kind of motivation, perpetuating its effects.

Put simply: Suspension will seem true only if we operate with an impoverished conception of motivation. Once we recognize the different ways in which motivation can operate, we see that Suspension is false.²⁸

4. RETHINKING THE CONNECTION BETWEEN AGENTIAL ACTIVITY AND DELIBERATION

4.1. There is no essential connection between deliberation and agential activity

I began by noting that philosophers have been increasingly drawn to the idea that deliberative agency is paradigmatically active. There are two ideas at work here. First, these philosophers believe that deliberation is typically sufficient for agential activity; excepting a few odd cases, if an agent deliberates, chooses to perform a certain action, and performs that action, then she manifests agential activity. Second, these philosophers believe that by investigating the structure of deliberative action, we will be able to understand other types of agential activity.²⁹ In the prior sections, I have asked why this should be so. I distinguished two claims:

(Inclination) In deliberative action, motives incline without necessitating. The agent's motives could be the same, and yet she could choose differently.
(Suspension) When an agent reflects on her motives for A-ing, she suspends the influence of these motives upon her assessment of whether there is reason to A.

²⁸ Note that I am not claiming that reflection *never* suspends the effects of motives. There may be some cases in which reflection does suspend the effects of motives. For example, when an agent reflects on weak and evenly balanced motives, such as the desire to have cereal for breakfast and the desire to have grapefruit for breakfast, it is entirely possible that reflection actually does suspend these motives. My argument is simply that Suspension—the claim that reflection either *always* or *typically* suspends the effects of motives—is false. (For the qualified version of Suspension, see Section 3.1.)

²⁹ As David Velleman puts it, the primary topic in action theory is “how to characterize the stereotypes or paradigms approximation to which determines the extension of the concept ‘action’”. The second is to characterize the dimensions along which instances can depart from the paradigm, and the contextual variables that determine how much of a departure is too much” (2004, 282).

If both of these claims were true, then there would be good reason to associate deliberative action with agential activity. Non-deliberative, unreflective actions would merely be caused by the agent's motives, whereas deliberative, reflective actions would be produced by the agent independently of her motives.

However, I have argued that Suspension is false. Let's now examine the implications of this conclusion. Given that Suspension is false, is there any reason for maintaining that deliberative actions are paradigmatically active?

The easiest way to answer that question is to consider cases in which the agent's choice determines what she does, and Inclination is true, but Suspension is false. Consider the following case, which I borrow from Nomy Arpaly:

I see a piece of cake in the fridge and feel a desire to eat it. But I back up and bring that impulse into view and now I have a certain distance... Is this desire really a reason to act? I consider the action on its merits and decide that eating the cake is not worth the fat and the calories. I walk away... This could be the inner monologue of a rational, autonomous being, and this is how it is usually presented. It could equally well be the inner monologue of an individual with severe anorexia nervosa, weighing eighty-five pounds... Imagine such a case: a woman who appears to herself to be in control of her desires, deciding between them on their merits, but who appears to her friends (or even to her future self, after having recovered from her anorexia or her irrational dieting) to be a person who is in fact at the mercy of her desires... The anorectic... experiences her psyche in terms of self-control, as if there were something that was *her*, choosing between her desires on the basis of their merits, giving her control over herself, while we have good reasons to believe that unconscious desire or emotion moves her in a manner not characteristic of well-exercised practical reason. (Arpaly 2004, 17–18)

Presumably, the anorectic agent does have the power to determine her action by choice: it is her decision not to eat the cake that leads her not to eat it. Moreover, we can assume that Inclination is true of her: her motives do not necessitate her choice, for she could have chosen, instead, to eat the cake. However, Suspension is not true of her. She steps back from her desire to be thin, reflects upon it, and attempts to assess it. She examines reasons for and against acting on it. She comes to a reasoned conclusion: she should not eat the cake, for she doesn't need the extra fat and calories. Yet this is manifestly not the conclusion that she would have drawn, had she actually managed to suspend the influence of her desire. Her conclusion is decisively influenced by the very motive that she takes herself to be suspending and rationally assessing: the desire to be thin. Her reflective scrutiny of her motive doesn't enable her to suspend this motive; on the contrary, it perpetuates the effects of this motive.

Precisely for this reason, it would be perverse to present the anorectic agent as an example of agential activity. On the contrary, she seems a paradigm of passivity: she takes herself to be determining her action via choice and rational reflection upon her motives, but her pathological desire is carrying her reflective thought in its wake.³⁰

Thus, the anorectic is a good example of how Suspension can fail. Moreover, notice that if the anorectic agent—who fulfills Inclination, but not Suspension—is not active, then Inclination alone is not a sufficient condition for agential activity. In other words, she illustrates that given the falsity of Suspension, reflective acts needn't be any more active than unreflective ones.

One might worry, however, that the anorectic is an aberrant case, an individual whose condition is a function of pathology. I think this would be a mistake; the anorectic is simply a more vivid example of the way in which motives operate in everyday, mundane cases. As the prior sections have argued, while our motives do not *necessitate* choice, they do *pervasively influence* choice, in ways that the agent often fails to grasp. To illustrate this point, consider a much more familiar kind of example:

Tom is a manager who must promote one of his employees in order to fill a recent vacancy. The leading candidates for the job are Candace and Dorothy. As it happens, Tom considers Candace extremely attractive, whereas he has no such feelings for Dorothy. However, Tom is a judicious and thoughtful manager, and he knows that perceived attractiveness should have no role in decisions about promotion. Indeed, he explicitly tells himself that he will not let Candace's attractiveness play any part in his decision. Instead, he tells himself, he will make the decision solely on the basis of Candace and Dorothy's performances in their current jobs. Carefully weighing the evidence, Tom ultimately concludes that while both Candace and Dorothy are admirable employees, Candace seems somewhat more capable, somewhat more attentive, and somewhat more driven than Dorothy. Consequently, he offers Candace the promotion.

So much for Tom's thoughts. Let's suppose, plausibly enough, that Tom's judgment that Candace is more capable, more attentive, and more driven is subtly influenced by his attraction to Candace. Impartial judges, who are not attracted to Candace, would view Dorothy as a marginally superior candidate.

³⁰ Here I am assuming that the anorectic does not admit to herself that she is acting on her anorexia; rather, she attributes her action simply to a desire to be thin. I am also assuming that she disapproves of her anorexia, so that, if she recognized its manifestations, she would disavow them. Of course, not all anorectics are like this: we could imagine an anorectic who openly admits to herself that she is avoiding the cake because she is anorectic, and who approves of this action. Arguably, such an anorectic would be active rather than passive. Thanks to Lisa Damm and Jennifer Hawkins for pressing me on these points.

There is no reason to doubt Tom's sincerity, in the above example. We can stipulate that he really does try, he really does *struggle*, to eliminate any traces of his motives on his judgment. We can further stipulate that he honestly thinks he succeeds; he thinks he has suspended the effects of the motive, and has made his decision impartially. Nonetheless, he fails to do so. The motive (attraction) exerts its influence on Tom's reflective thought, even as he tries to distance himself from it.³¹ Consequently, he seems to be a passive conduit for forces within.

This kind of case is ubiquitous; a host of psychological studies have demonstrated that it occurs all the time. Employees who are perceived as more attractive are more likely to be promoted, to receive better evaluations, to garner higher pay, and so forth.³² In many such cases, where the agent takes herself to be achieving distance from a motive that she disavows, but nevertheless is driven by the motive, the agent will strike us as failing to manifest agential activity.

4.2. A new account of agential activity

If Kant and Locke had been right about Suspension—if reflective deliberation actually did enable us to suspend our motives, and choose in complete independence of them—then there would be good reason for singling out deliberative actions as paradigmatically active. However, given that deliberation typically perpetuates the effects of our motives, there is no obvious reason for associating deliberative action with agential activity.

Does this show that the distinction between activity and passivity should be abandoned? If being moved by reflective thought isn't something *other than* being moved by motives, should we conclude that there's no principled reason for attempting to distinguish between activity and passivity in the production of action?

I think those conclusions would be premature, for they would clash with some of our most basic intuitions about agency. Consider the following case:

³¹ That is, Tom asks himself whether he should act on his attraction to Candace, decides that he should not, attempts to suspend the motive, and nonetheless is under the motive's influence as he reflects. Put simply, the motive tells him "promote Candace!" and this is exactly what he does.

³² Marlowe et al. (1996) and a number of other studies reveal that there is a demonstrable bias in promotion and hiring decisions toward candidates who are judged to be more attractive. Hammeresh and Parker (2005) discuss a study showing that teachers who are judged more attractive receive higher ratings from their students on course evaluations.

Amy is walking down the street on her way to a meeting. She's very hungry. However, she is in a rush. She decides that she should hurry on her way, and tries to put the thoughts of food out of her mind. As she walks along, she doesn't explicitly consider her hunger; her mind is occupied with thoughts of what she will say in the meeting. However, as she walks, her attention is drawn to the presence of food; she eyes the restaurants that she strolls by, she finds herself interrupted by thoughts of food, and so on. As she is about to pass a coffee shop, she suddenly stops and decides to go in for a snack.

Amy's hunger manifests itself by drawing her attention to food, by inclining her to think about food, and so on. Hunger influences reflection by making certain features of the situation salient, by inclining the agent to give weight to considerations relating to the acquisition of food, and so on. But in the case of ordinary, mild hunger these effects are typically unobjectionable, for several reasons. First, and most importantly, hunger typically attunes one to food only when one actually requires food. Thus, an agent who is being motivated by her hunger will typically approve of the way in which she is being motivated. Second, the effects are obvious; it is hardly a surprise to learn that one's hunger is leading one to focus on food. Third, the effects are benign; though the effects of hunger may be distracting, they typically do not interfere with one's capacity to reason and to pursue other goals. (Of course, things will be very different in cases of severe hunger.)

Contrast this with Arpaly's anorectic. The anorectic case has the same features as our case of hunger. In the anorectic case, a desire to appear thin exerts itself by affecting perceptual saliences, judgments, and reflective thoughts; in the hunger case, a desire to eat exerts itself by affecting perceptual saliences, judgments, and reflective thoughts. There is no *fundamental* difference between the ways in which motives operate in the anorectic case and the hunger case; they operate in exactly analogous ways, albeit with different intensities.

In both cases, then, the agent's motives channel and guide her reflective thought. Nonetheless, I think most of us will agree that there is a distinction, here: the case of hunger seems far less problematic than the anorectic case.

Thus, rather than concluding that there is no distinction between the active and the passive, we might conclude that there is a distinction, but that the distinction is not to be marked in terms of whether one's action was caused in a reflective or a non-reflective manner. Arpaly's anorectic is highly reflective and deliberate, but seems a paradigm of passivity. The hungry agent, by contrast, is minimally reflective and deliberate, yet seems entirely active.

I think these ideas can be developed into a new theory of agential activity. They seem to indicate that if there is to be a distinction between activity and passivity in the production of action, we will need something

that is more nuanced, something that comes in degrees, and that takes account of the complex interactions between reflective thought and affect. The relevant distinction will not be whether the agent suspends the effects of her motives, for that might never happen; rather, the relevant distinction will have something to do with *how* the agent's motives affect the agent's reflective thoughts.

I suggest that the following concept suffices to characterize the relevant difference between the anorectic and the hungry agent:

(Equilibrium) The agent A's, and approves of her A-ing. Further knowledge of the motives that figure in A's etiology would not undermine her approval of A-ing.

(Disequilibrium) The agent A's, and currently approves of her A-ing. However, if she knew more about the motives that figure in A's etiology, she would no longer approve of her A-ing.³³

The anorectic conceives of her action as resulting from an entirely reasonable desire to avoid excessive calories, but the more apt description of her action would attribute it to pathological desires to avoid food, inaccurate self-conceptions, disgust with her own body, and so on. If the anorectic were to become aware of this, it would likely change her attitude toward her action. Thus, she is in disequilibrium.³⁴ The hungry agent, on the other hand, would retain the same attitude toward her action were she to know more about the motives prompting it. Thus, she is in equilibrium.

A few clarifications are in order. First, the account should be understood as applied to agents, holding all else constant except giving the agent further information about the motives figuring in the etiology of the action under consideration. In particular, we do not want to consider cases in which the agent changes his values. For example, I would now disapprove of many of the actions that I performed, approvingly, as a child; but this does not show

³³ I borrow the term "disequilibrium" from Nozick, who uses it to define a related condition. According to Nozick, an action is in disequilibrium for a person if "(a) he does (or wants to do) it, yet (b) if he knew the causes of his doing or wanting to do it, this knowledge would lead him not to do it, or not to want to (or to want not to want to do it, or at least to a lessening of his want to do it . . .)" (Nozick 1981, 349). Otherwise, the act is in equilibrium. Nozick puts the notion of equilibrium to different work, suggesting that we might be able to explain *what it is to be good* in terms of *being in equilibrium* and meeting certain additional conditions (Nozick 1981, 350ff.). My task, here, is less ambitious: I argue that the version of equilibrium defined above offers a characterization of agential activity.

³⁴ Note that I am not claiming that *all* anorectics are in disequilibrium, just that the anorectic described in Arpaly's example is in disequilibrium. It seems to me that some anorectics are in equilibrium—a quick Internet search will reveal a profusion of websites in which well-informed anorectics, who seem to be fully cognizant of the etiology of their anorexia, nevertheless valorize their anorexia and encourage others to do the same. On my view, these anorectics are active. See also note 30, above.

that the actions were in disequilibrium, for at the time of action I may have wholeheartedly approved of them.³⁵

Second, notice that the only factor that we are changing here is the amount of information that the agent has about the *etiology* of the action. Agents sometimes disapprove of a past action not because they learn more about the act's etiology, but because they learn more about its *consequences*. I did not realize that my innocent, offhand remark would hurt Sarah's feelings. Now, seeing her upset, I regret the remark, and wish that I had not made it. But this does not show that my action was in disequilibrium.

Notice, third, that an agent whose actions are in disequilibrium would not necessarily want to *act* differently. She might be dissatisfied with her actions not because she disapproves of what she has done, but because she disapproves of her *motives* for doing what she has done. For example, suppose that Sally volunteers in a soup kitchen. She believes she is volunteering out of a desire to aid the impoverished beneficiaries. Yet a psychologically adept observer, well acquainted with Sally's character, would describe things differently: Sally takes satisfaction in feeling superior to the impoverished recipients, and her volunteering is in part motivated by this desire. Suppose Sally comes to realize that one of the desires motivating her action is the desire to feel superior. She finds this desire reprehensible, and she is no longer able to view her action of volunteering with approval. Thus, her action is in disequilibrium. However, it would be inaccurate to say that she wants not to volunteer. Rather, she still wants to volunteer, but she wants to volunteer out of beneficent motives, rather than self-serving ones. So, she is in disequilibrium not because she wants to act differently, but because she wants to act out of different motives.^{36,37}

³⁵ Equilibrium bears a resemblance to Harry Frankfurt's notion of wholeheartedness. Roughly, Frankfurtian wholeheartedness obtains when the agent bears a higher-order attitude of acceptance or approval toward his lower-order desires (cf. Frankfurt 2004). Equilibrium is more demanding: the approval in question must be stable in the face of further information about the action's etiology.

³⁶ Analogous points apply to more complex cases. Consider the following example, a version of which was suggested by Julia Driver: Bill screams at his wife in anger, and takes this action to be motivated by jealousy. In cool moments, he thinks this jealousy is unwarranted. He disapproves of screaming out of unwarranted jealousy, so his action is in disequilibrium. However, unbeknownst to Bill, there are good reasons for his jealousy: though Bill doesn't admit it to himself, there are compelling indications that his wife has been unfaithful. Suppose that Bill would approve of yelling at his wife out of *warranted* jealousy. What does my model imply about a complex case like this? I would say that this case is structurally analogous to the above case: the agent is in disequilibrium, because he wants to act out of different motives (in this case, he wants to act out of warranted jealousy rather than unwarranted jealousy).

³⁷ Notice that my account implies that an agent could be passive for long stretches of his life. Consider the following case. Bill meets and becomes strongly attracted to Sarah.

A final clarification: this account of equilibrium does not associate agential activity with *morally upstanding* or *praiseworthy* activity. An agent could be in equilibrium while performing an action that is manifestly immoral. For example, a sociopath might inflict pain on another agent because it pleases him to do so. It is entirely possible that even if the sociopath knew all of the motives that figure in the etiology of this action, he would continue to approve of it. Thus, he would be active in performing the action.

With these points in mind, we can see that disequilibrium constitutes a form of psychic conflict. An agent acts, and approves of his action. However, this approval is contingent upon his ignorance of the motives that are actually leading him to act. So there is a conflict between the agent's attitude toward the action as he takes it to be, and the agent's attitude toward the action as it is. Moreover, disequilibrium implies that one has motives that are influencing one in ways that one would disavow. Thus, there is an interesting form of conflict between the agent's reflective and unreflective aspects at the time of action.

I submit that equilibrium is a necessary condition for agential activity. Equilibrium seems to offer a characterization of the conditions under which an agent can be said to be in control of her action. The agent acts, approves of the act, and further knowledge of the action would not undermine this approval. To speak metaphorically, the agent's whole being is behind the action.³⁸

Sarah, as it happens, is an ardent and vociferous vegan, whereas Bill has long regarded veganism as unnecessary and perhaps even vaguely objectionable. However, Bill now worries that his meat-eating ways will ruin his chances with Sarah. Bill regards himself as a principled and steadfast person, who will not change his habits merely for the sake of a potential romance. Thus, he explicitly decides that his desire for a relationship with Sarah is not a good reason for becoming a vegan. However, reflecting on veganism's merits, Bill decides on other grounds to become a vegan. (Perhaps he tells himself that the suffering inflicted on factory-farmed animals is sufficient reason not to eat them.) Now, assume that, unbeknownst to him, Bill's romantic desire did influence this decision to such an extent that, were Bill aware of these facts, he would disapprove of his decision. Bill's vegan actions are then in disequilibrium. However, these vegan actions might proceed for many years, structuring large portions of Bill's life. Yet, on my account, they will be passive. Is this a problem? Is it odd to think that an agent could be passive for long stretches of his life? In response, two points are worth emphasizing. First, on my account it certainly is possible for an agent to be passive in *every* action that he performs. However—and this is the second point—such a case would be extremely unusual. In the vegan case, Bill is passive in his larger action (or project) of being a vegan, but he may be entirely active in the smaller actions that partially constitute this larger action. For example, when Bill goes to the grocery store to buy vegan food, when he orders tofu at the restaurant, and so forth, these events, considered as isolated actions, may be in equilibrium, and hence active. (Thanks to John Brunero for raising this question and suggesting a similar example.)

³⁸ Another image may be helpful. Intuitively, there is a distinction between a motive *influencing* the agent and a motive *changing* the agent: when considering a new motive, we sometimes want to say that the motive merely influences the agent, and at other

4.3. A final clarification regarding activity and passivity in reflective thought itself

There is a complication to address. My proposed account of agential activity asks us to consider whether the agent's attitude toward the action would be stable if the agent were aware of additional information about the action's etiology. Of course, *reflecting on an action's etiology* is itself an action, and is thus susceptible to the forms of influence by motives that I have discussed above. Consequently, when an agent asks herself whether her action is in equilibrium, it is possible for this act of equilibrium-assessment to itself be in *disequilibrium*.

An example will be helpful. Consider the famous quotation from Augustine, in which Augustine describes his efforts to extirpate his attraction to pride:

Even when I reproach myself for it, the love of praise tempts me. There is temptation in the very process of self-reproach, for often, by priding himself on his contempt for vainglory, a man is guilty of even emptier pride. (*Confessions*, Book X, §38)

Augustine is attempting to avoid the performance of any actions that are motivated by pride. Thus, suppose that Augustine is engaged in casual conversation with an acquaintance, and without prompting Augustine begins to speak of his own accomplishments. Call this action A. At the time of action, Augustine views himself as motivated simply by a desire to make conversation. Reflecting on the action at a later time, though, Augustine recognizes that his A-ing was in fact motivated by pride. Recognition of this motive for A-ing leads Augustine to disapprove of his A-ing. Thus, he judges A-ing to be in disequilibrium. Now, let B-ing be Augustine's higher-order action of reflecting on and assessing A-ing. Upon further reflection, Augustine notices that his B-ing is itself motivated by a particular form of pride: the pride that he takes in chiding himself for falling short of his ideals. Thus, B-ing—that is, Augustine's act of assessing

times—when the motive is particularly influential or deep-seated—that it changes the agent's character. My account recognizes this distinction. Some agents are simply influenced by their motives in such a way that, were they aware of this influence, they would reject it. An example is Arpaly's anorectic. These agents are passive. Other agents are so deeply changed by their motives that, were they aware of this influence, they would accept the influence of those motives. An example is the self-satisfied anorectic mentioned in note 34. These agents are active. (Thanks to David Enoch for asking me to address the distinction between change and influence.)

his own action and judging it to be in disequilibrium—is *itself* in disequilibrium.³⁹

As Augustine illustrates, the very action of *trying to determine whether one's own actions are in equilibrium* can itself be in disequilibrium. This might seem to engender a problem for my view. Do the equilibrium-assessments themselves have to be in equilibrium in order for the assessments to be accurate? And if so, won't this run the risk of launching us into an infinite regress?⁴⁰

In fact, this is not a problem for my account of activity as equilibrium. It *would* be a problem if whether a given (first-order) action, A-ing, were in equilibrium depended upon whether the (second-order) act of reflecting upon and assessing A-ing were itself in equilibrium. For, on that view, we would then have to ask whether the (third-order) act of reflecting upon and assessing the reflecting upon and assessing of A-ing were in equilibrium, and this would launch us into an infinite regress.

However, my account does not depend upon this kind of assessment. It does not matter whether, like Augustine, the agent actually engages in an episode of reflection and comes to some assessment of his past action. What matters is simply the status of a certain counterfactual claim: whether, if the agent were aware of further information about his motives, he would continue to approve of his action. We are not examining the agent's *actual* thoughts, we are examining a counterfactual—in which the agent has more information about the etiology of his action—and asking whether the agent's approval of the action then dissipates. In assessing this counterfactual, we take the closest possible world: we change *only* the amount of information that the agent has about the etiology of his action, and we see whether the agent's approval dissipates or remains stable.

So, to use the Augustine example: Augustine's original action of speaking of his own accomplishments is in disequilibrium, because he would (and in fact does) disapprove of it given further information about its etiology. His action of reflecting on and assessing this past action is also in disequilibrium: given further information about its etiology, Augustine's approval of this action would (and in fact does) dissipate. But the question of whether this higher-order action is in disequilibrium has no bearing whatsoever on the question of whether the original action is in disequilibrium.

³⁹ Like the soup-kitchen volunteer mentioned above, Augustine does not want to perform a different action; rather, he wants to perform the same action out of different motives. That is, he still wants to disapprove of his past prideful action, but he wants this act of disapproval not to be motivated by pride.

⁴⁰ Thanks to several participants in the Wisconsin Metaethics Workshop for pressing me on this point.

With this in mind, we can draw attention to a related point. Notice that whether the agent's approval of his action dissipates or remains stable depends on what other motives the agent has. Put differently, we are *not* attempting to determine how the agent would react to his own action in the (impossible) case of not being affected in any way by his own motives. Take an example. An angry agent may see that his current action of screaming at Tom is motivated by jealousy. In a cool moment, he would disapprove of screaming out of jealousy. However, in the midst of his anger, fully cognizant of why he is screaming, he approves of the action. On my view, this action is in equilibrium: in the closest world, where all that we have changed is that the agent has more information about the action's etiology, he will retain his approval, precisely because he will still be angry. This seems to be the right result: an agent who clear-sightedly acts out of jealousy, in full cognizance of what he is doing and why, seems quite different than an agent whose jealousy surreptitiously guides him in ways that he would disavow were he cognizant of them.

In sum: higher-order actions, such as the action of assessing one's actions for equilibrium, can themselves be in either equilibrium or disequilibrium. In other words, higher-order actions can be either active or passive. This is no surprise: given the above account of motivation, we should expect that reflective thought is not necessarily active. Like overt physical actions, our reflective thoughts are subject to the vicissitudes of activity and passivity. Yet the question of whether a particular act—be it higher- or lower-order—is itself in equilibrium is not determined by the agent's actual reflective thoughts concerning the act. Consequently, the problem of regress does not arise.

4.4. The advantages of this account of agential activity

The proposed account of agential activity as equilibrium has several advantages over the traditional account. First, it does not commit us to the (false) claim that deliberation is capable of suspending the effects of our motives. Second, it does entail that reflection—or, more precisely, reflective attitudes toward one's own actions—is partly determinative of whether one is agentially active. Third, the account of agential activity as equilibrium offers a correct characterization of certain paradigm cases: it judges the anorectic agent to be passive, and the hungry agent to be active.

The core idea of this account is that passivity does not involve being moved unreflectively or independently of reflection. After all, we are moved unreflectively all the time: relatively few of our actions are preceded or accompanied by anything like a bout of explicit reflection, and even when

we do reflect, this often has the effect of perpetuating, rather than suspending, our motives. Rather, passivity involves being moved in a way that *conflicts with* the reflective attitudes that one would have, if one recognized how one was being moved.⁴¹

5. CONCLUDING REMARKS

In this chapter, I have examined a pair of ideas that has grown increasingly influential in recent work on agency. The first idea is that agents play more and less active roles in the production of their actions. The second is that deliberative agency is paradigmatically active.

In an attempt to analyze these ideas, I have teased apart three distinct claims about deliberative agency: that choice determines action, that motives do not determine choice, and that agents are capable of deliberately suspending the effects of their motives. I have argued that while the first two claims are probably true, the third is false. The claim that we can suspend the effects of our motives will seem plausible only if we work with a deficient conception of motivation. However, I argued in favor of a more adequate conception of motivation, which reveals that motives operate as influences upon reflection itself. With this theory of motivation at hand, I argued that deliberating agents typically fail to suspend the effects of their motives. While it is true that the agent who deliberates is not simply compelled to act by his motivational states, his motives do continue to operate, in a subterranean fashion, even as the agent reflects on them. In some cases, the motives appear to decisively guide the agent's choice, in ways that he does not recognize.

As a result, I have argued there is no good reason for taking deliberative agency as paradigmatically active. The assumption that reflective, deliberative actions are paradigms of agential activity relies on an impoverished account of the way in which motives operate. Once we have a better account of motivation, we can see that reflective actions needn't be more active than unreflective actions. I closed by articulating and defending a new model of agential activity, which avoids these problems. According to

⁴¹ A central feature of this account is the idea that an individual whose action is in equilibrium cannot act akratically. If an agent performs an action of which she disapproves, she is not in equilibrium. For this reason, it would not be inaccurate to characterize my account as maintaining that agential activity is the opposite of akrasia. Notice that if we do characterize the account in this way, we must classify certain cases of surreptitious influence by motives—such as Arpaly's anorectic case and my hiring case, above—as cases of akrasia. Thanks to Chris Korsgaard for suggesting that my view could be characterized in this way.

this new account, an agent manifests agential activity when she approves of her action, and this approval is stable in the face of further information about the action's etiology. In short, agential activity concerns the relationship between the agent's reflective thought and the agent's unreflective motivation. I have argued that this new model of agential activity accurately captures our intuitions about certain paradigm cases, and is fully compatible with the more sophisticated account of motivation.

References

- Arpaly, Nomy, *Unprincipled Virtue* (Oxford: Oxford University Press, 2004).
- Augustine, *Confessions*, R. S. Pine-Coffin (trans.) (New York: Penguin Classics, 1961).
- Bower, G. H., "Mood and memory," *American Psychologist* 36 (1981), 129–48.
- Bratman, Michael, *Faces of Intention* (New York: Cambridge University Press, 1999).
- *Structures of Agency* (Oxford: Oxford University Press, 2007).
- Clark, M. S., and B.A. Waddell, "Effects of moods on thoughts about helping, attraction and information acquisition," *Social Psychology Quarterly* 46 (1983), 31–5.
- Feshbach, S., and R. D. Singer, "The effects of fear arousal and suppression of fear upon social perception," *Journal of Abnormal and Social Psychology* 55 (1957), 283–8.
- Foley, Richard, "Deliberate Action," *Philosophical Review* 86 (1977), 58–69.
- Forgas, J. P., "The role of emotion in social judgments: An introductory review and an affect infusion model (AIM)," *European Journal of Social Psychology* 24 (1994), 1–24.
- and P. Vargas, "The Effects of Mood on Social Judgment and Reasoning," in Michael Lewis and Jeannette M. Haviland-Jones (eds.), *The Handbook of Emotions* (New York: Guilford Press, 2004).
- Frankfurt, Harry, *The Reasons of Love* (Princeton, NJ: Princeton University Press, 2004).
- Hammerash, Daniel, and Amy Parker, "Beauty in the Classroom: Instructors' Pulchritude and Putative Pedagogical Productivity," *Economics of Education Review* 24 (2005), 369–76.
- Hampshire, Stuart, "Fallacies in Moral Philosophy," *Mind* 58 (1949), 466–82.
- Hobbes, Thomas, *Leviathan*, Edwin Curley (ed.) (Indianapolis: Hackett Publishing, 1994).
- Holton, Richard, "The Act of Choice," *Philosopher's Imprint* 6.3 (2006), 1–15.
- Ilsen, Alice, "Towards understanding the role of affect in cognition," in R. S. Wyer and T. K. Srull (eds.), *Handbook of Social Cognition*, Vol. 3 (Hillsdale, NJ: Erlbaum, 1984), 179–236.

- "Positive affect, cognitive processes and social behaviour," in L. Berkowitz (ed.), *Advances in Experimental Social Psychology*, Vol. 20 (San Diego, CA: Academic Press, 1987), 203–53.
- Kahneman, Daniel, and Amos Tversky, "The Framing of Decisions and the Psychology of Choice," *Science* 211 (1981), 453–8.
- Kane, Robert, *The Significance of Free Will* (Oxford: Oxford University Press, 1996).
- Kant, Immanuel, *The Metaphysics of Morals*, Mary Gregor (ed.) (New York: Cambridge University Press, 1996).
- Korsgaard, Christine, *The Sources of Normativity* (Cambridge: Cambridge University Press, 1996).
- "Self-Constitution in the Ethics of Plato and Kant," *Journal of Ethics* 3 (1999), 1–29.
- *Self-Constitution* (Oxford: Oxford University Press, 2009).
- Kunda, Ziva, "The Case for Motivated Reasoning," *Psychological Bulletin* 108 (1990), 480–98.
- Leibniz, G. W., and S. Clark., *G.W. Leibniz and Samuel Clark: Correspondence*, Roger Ariew (ed.) (Indianapolis: Hackett Publishing, 2000).
- Locke, John, *An Essay Concerning Human Understanding* (Oxford: Oxford University Press, 1975).
- Marlowe, C., S. Schneider and C. Nelson, "Gender and Attractiveness Biases in Hiring Decisions," *Journal of Applied Psychology* 81 (1996), 11–21.
- Mele, A. R. (2003), *Motivation and Agency* (Oxford: Oxford University Press, 2003).
- *Effective Intentions: The Power of Conscious Will* (Oxford: Oxford University Press, 2009).
- Molden, D. C., and E. T. Higgins, "Motivated Thinking," in K. Holyoak and R. Morrison (eds.), *The Cambridge Handbook of Thinking and Reasoning* (New York: Cambridge University Press, 2005).
- Nozick, Robert, *Philosophical Explanations*. (Cambridge, MA: Harvard University Press, 1981).
- Phelps, E. A., "The Interaction of Emotion and Cognition," in R. Hassin, J. Uleman, and J. Bargh (eds.), *The New Unconscious* (New York: Oxford University Press, 2005).
- Raz, Joseph, *Engaging Reason* (Oxford: Oxford University Press, 2002).
- Sanitosa, R., Z. Kunda, and G. T. Fong, "Motivated Recruitment of Autobiographical Memory," *Journal of Personality and Social Psychology* 59 (1990), 229–41.
- Schapiro, Tamar, "The Nature of Inclination," *Ethics* 119 (2009), 229–56.
- Schick, Frederic, *Ambiguity and Logic* (New York: Cambridge University Press, 2003).
- Schopenhauer, Arthur, *Prize Essay on the Freedom of the Will*, E. F. J. Payne (trans.) (New York: Cambridge University Press, 1999).
- Schroeder, Mark, *Slaves of the Passions* (Oxford: Oxford University Press, 2007).
- Velleman, J. David, *The Possibility of Practical Reason* (Oxford: Oxford University Press, 2000).

- Velleman, J. David, "Replies to Discussion on *The Possibility of Practical Reason*," *Philosophical Studies* 121 (2004), 277–98.
- Wallace, Jay, *Normativity and the Will* (Oxford: Oxford University Press, 2006).
- Watson, Gary, *Agency and Answerability* (Oxford: Oxford University Press, 2004).
- Witt, J. K., D. R. Proffitt, and W. Epstein, "Perceiving distance: A role of effort and intent," *Perception* 33 (2004), 577–90.