

ROBERT J. HOWELL



CONSCIOUSNESS AND THE LIMITS OF OBJECTIVITY

The Case for Subjective Physicalism

OXFORD

Consciousness and the Limits of Objectivity

This page intentionally left blank

Consciousness and the Limits of Objectivity

The Case for Subjective Physicalism

Robert J. Howell

OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,
United Kingdom

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. Oxford is a registered trade mark of
Oxford University Press in the UK and in certain other countries

© Robert J. Howell 2013

The moral rights of the author have been asserted

First Edition published in 2013

Impression: 1

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, by licence, 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 work in any other form
and you must impose this same condition on any acquirer

British Library Cataloguing in Publication Data

Data available

ISBN 978-0-19-965466-6

Printed by the MPG Printgroup, UK

To Lanie

This page intentionally left blank

Contents

<i>Acknowledgments</i>	viii
Introduction	1
Part I. Defining Physicalism	
1. The Base Problem	13
2. Supervenience and the Relation Problem	41
Part II. The Threat of the Subjective	
3. Phenomenal Knowledge and Acquaintance	57
4. Acquaintance and Objectivity	90
Part III. Saving Physicalism	
5. The Ontology of Subjective Physicalism	101
6. Deduction, Necessitation, and Acquaintance	123
7. Assessing Subjective Physicalism	152
<i>Bibliography</i>	174
<i>Name Index</i>	187
<i>Subject Index</i>	189

Acknowledgments

Several portions of this book have been adapted, often with significant changes, from previous publications. Parts of Chapter 1 were published in 2011 as “Physicalism, Old School,” in *Consciousness and the Mind Body Problem: A Reader*. Eds. Torin Alter and Robert J. Howell, New York: Oxford University Press. Most of Chapter 2 was published in 2009 as “Emergentism and Supervenience Physicalism,” *Australasian Journal of Philosophy*, 87:1. Portions of Chapters 3 and 4 involve arguments from “The Knowledge Argument and Objectivity,” *Philosophical Studies*, September 2007. Chapter 5 is a revision of “The Ontology of Subjective Physicalism” *Nous*, 43:2, 2009.

Research on this book was aided in part by research leave granted by SMU in the Fall of 2009 during which I was a visiting scholar at the New York University department of philosophy as part of the NYU Consciousness Institute. My sincerest thanks to both SMU and NYU for their support. While in New York City, I received valuable support from the Ferris/Kennedy home for wayward philosophers, and I cannot thank Josh, Elizabeth, and Coop enough for their hospitality and friendship.

I have benefitted enormously from discussions with many individuals over the years. Though this book did not grow out of dissertation work, my advisors Jaegwon Kim and Ernest Sosa have continued to be very nurturing, and discussions with Jaegwon in particular have (quite obviously) continued to have a strong impact on my thinking about the topics in this book. I also benefitted greatly from an NEH Seminar in Mind and Metaphysics led by John Heil. John and the participants in that seminar were crucial in helping me get some of the key ideas into shape and I thank them all. Thanks to audiences at the University of Alabama, Brown University, and Washington University in St Louis, as well as audiences at the Tucson Consciousness Conference and several APA sessions. I am very grateful for, and have learned much from, discussions with Roberta Ballerin, Kati Balog, Jacob Berger, Ned Block, Richard Brown, Max Deutsch, Josh Dever, Carrie Figdor, Brie Gertler, Chris Hill, Cory Juhl, Amy Kind, Uriah Kriegel, William Lycan, Pete Mandik, Andrew Melnyk, Kevin Morris, Brendan Murday, Alyssa Ney, Brendan O’Sullivan, David

Papineau, Derk Pereboom, David Pitt, Jim Pryor, Ted Sider, Jon Simon, David Sosa, Jonathan Sutton, and Jessica Wilson. Special thanks go to Peter Momtchiloff for his continued editorial guidance.

I have the great fortune to be in an extremely collegial department filled with very strong philosophers. Everyone in the SMU department has helped me with this book in one way or another. I owe Doug Ehring a great debt for helping me become a better writer and philosopher during the past ten years, and the leadership and insights of Eric Barnes and Steve Sverdlik have been very important to me as well. I am particularly lucky to have as colleagues three excellent philosophers of mind who keep me honest and make me better: Philippe Chuard, Justin Fisher, and Brad Thompson. Philippe and Brad both read complete drafts of the book, prompting crucial revisions. Problems with the book remain—I know, because Philippe told me—but it would be a much weaker book indeed were it not for the BrainPains crew.

Ken Daley and Clayton Littlejohn deserve special mention not only for their insights in many a philosophical discussion, but for their friendship and support. I owe Torin Alter more than I can say, not only for being an inspiring co-author and for reading draft after draft of my own work, but for being one of my dearest friends. I doubt this book would exist without him.

Although it deserves a larger forum, I can't help but take space here to thank one of the angels of our discipline. David Chalmers is of course a philosopher of the highest order, but if he had never put pen to paper he would remain one of philosophy's MVPs. His willingness to help others, especially younger philosophers who have yet to achieve prominence, raises the bar for all of us. Though he didn't know me from Adam when I sent him my first paper on consciousness almost a decade ago, he read it and got back to me with comments in less than 12 hours . . . from Canberra. Again and again he has read my work, usually quickly finding the flaws that sent me back to the drawing board. Throughout, though, he has always been encouraging and has become a true mentor and friend.

Finally, I offer my greatest thanks to my wife Lanie, to whom this book is dedicated. Her support has been crucial to all of my work, and her love and friendship continues to make me a better person in all ways. As I finish this book, Lanie is just days away from giving birth to our first child. That's not only given me joy, but it's helped keep things in perspective. The book is done, but I eagerly await our next chapter.

This page intentionally left blank

Introduction

Writing a book about the philosophical problem of consciousness is a bit like opening a restaurant in New York City. Surely there are only so many ways to prepare a fish! When it comes to consciousness, positions have been staked out, mapped out, and the logical space seems to be getting pretty tight. The problem is by now familiar: conscious states, such as the state of seeing red or feeling pain, have natures that seem to be beyond the reach of physical explanation. While physics and the physical sciences can explain all the pushings and pullings of the world, and those pushings and pullings can explain a surprising amount, it is difficult to imagine how they might fully explain what it's like to taste chocolate. These features of experiences can only be known "from the inside" by actually having the experiences. Faced with this difficulty, most philosophers have divided into three factions:

The hardliners: The hardliners deny the initial intuition, or at least its value in thinking about consciousness. Consciousness can be fully explained and understood by the physical sciences, and actually experiencing those states will add nothing to that understanding.

The epistemicists: The epistemicists maintain that though there is something one can learn by actually having an experience, this does not imply that physicalism is false. There is an epistemic gap between our understanding of physical states and our understanding of conscious phenomenal states, but there is no corresponding metaphysical gap.

The non-physicalists: The non-physicalists—who include dualists, panpsychists, proto-panpsychists, and the like—claim that this gap between our understanding of the physical world and our grasp of phenomenal consciousness does warrant a metaphysical conclusion. Physicalism must be false.

While there are ways of carving things up into thinner slices, this threefold division captures the major positions. One group denies the datum (or at least its importance as datum), another group accepts the datum and denies the inference from it, and the last group accepts both. Unless everyone has the basic dialectic incorrect, these positions seem to exhaust logical space. Each space is well populated by individuals with subtly different views, and excepting a few conversions, once philosophers choose a place to camp they rarely feel the need to move.

The majority of philosophers, it seems, are satisfied with the epistemicist position in one form or another, pleased to be able to have their cake and eat it too. While I will ultimately embrace a version of epistemicism, it is one of the main themes of this book that current epistemicists have not yet earned their peace of mind.¹ When subjected to scrutiny their view has a distinctly non-physicalist fragrance. In particular, I will argue that maintaining their view requires them to believe in a special epistemic relation, the relation of acquaintance, which is not easily squared with a purely physicalistic outlook. What's more, this commitment in turn forces them to deny that any objective depiction of the world can be complete. That is, physics must in some sense prove incomplete.

I suspect that if convinced of my arguments, many current epistemicists will find themselves sliding to a hardline stance. Denying "objectivism" will seem far too close to denying physicalism for most physicalists. I think this worry is real, and it should be recognized, but that there is ultimately a way to save the epistemicist project from complete collapse. It is to embrace a "subjective physicalism." Doing so, however, will require a bit more metaphysical spadework than most epistemicists have done. Even making sense of subjective physicalism takes a little doing, much less making it palatable.

This book attempts to show a way forward for the dedicated physicalist who is gripped by the problem of consciousness and wants to give the basic data their full due. It does so in part by developing a truly metaphysical definition of physicalism, untainted by epistemic elements, and it argues that only by questioning the non-physicalist's metaphysical arguments "upstream" of his position that the anti-physicalist arguments can be

¹ It is instructive, I think, that one of the earliest epistemicists—Terry Horgan (1984)—doesn't seem to have this peace of mind. The worries expressed in Horgan and Tienson (2001), for example, are of the sort that must still be dealt with.

resisted. It also attempts to explain the role of acquaintance in our apprehension of phenomenal properties, and to explain how the introduction of this relation solves problems without introducing worse ones. In my view, subjective physicalism is not so much an alternative to epistemicism as it is a more honest and more complete form of it. It is more honest, I claim, because it acknowledges its dependence on acquaintance and the loss of objectivism. It is more complete because it recognizes that simply “going epistemic,” even to the point of embracing acquaintance, doesn’t get one out of the anti-physicalist arguments unless some further ontological work is done.

Subjective physicalism occupies an admittedly uncomfortable space between traditional epistemicism and non-physicalism. It can seem so uncomfortable—a fact I discuss in the final chapter—that one wonders why one shouldn’t just give up the game and become a dualist or some other flavor of non-physicalist! The reason is simple: these views, far from bringing phenomenal experience to the fore, unintentionally relegate them to irrelevance.

Although the problem of mental causation is well rehearsed, it is worth reviewing. This problem provides a significant part of the motivation for a physicalistic account of the puzzle of conscious experience, and it will occupy the background of much of the discussion to come. The problem of mental causation is a result of the apparent inconsistency of the following independently plausible theses:

1. *Mental distinctness*: The mental is not identical with the physical.
2. *Physical adequacy*: Physical events have sufficient physical causes if they are caused at all.²
3. *Mental causation*: Some physical events are caused by mental events.
4. *Non-overdetermination*: Not every case of mental causation is a case of overdetermination.

If these theses are in fact inconsistent, any three of them constitutes a valid argument against the fourth. It is not unusual for physicalists to argue

² I prefer speaking of physical adequacy instead of “the causal closure of the physical” since the latter, but not the former, seems to imply that physical events only cause physical events. This problem arises, however, from the plausible view that a physical account can be provided for why physical events happen.

against mental distinctness by appealing to theses 2 through 4. Papineau puts the case succinctly:

Many effects that we attribute to conscious causes have full physical causes. But it would be absurd to suppose that these effects are caused twice over. So, the conscious causes must be identical to some part of those physical causes.³

It is somewhat ironic that the physicalist is in a better position than the dualist to acknowledge the centrality of conscious thought in our lives. In fact, the dualist faces a painful choice between isolation and infection: either conscious states exist in causal isolation without contributing anything new to the physical world (which is a consequence of rejecting either mental causation or non-overdetermination) or there are physical events that cannot be explained by the physical sciences. In the latter case consciousness is not simply a benign addition to the physical system but infects causal claims within the domain of physics.

It is important not to underestimate the significance of either isolation or infection. It is tempting to think that the epiphenomenality of consciousness might not be such a disaster. It might even be thought that contemporary science is showing us that consciousness does not play the causal role we think it does, and so the epiphenomenality of consciousness actually receives empirical support.⁴ This cannot be the case. Even if it turns out that conscious experiences do not play the roles in behavior that we think they do, the idea that they play no role at all is absurd. At the very least conscious states had better play a causal role in our knowledge of them and in our talk about them. The empirical tests, in fact, assume as much since they rely upon physical signs (such as verbal reports) to tell the experimenter when the conscious experience is present. Surely if we have to give up on the idea that our reports and descriptions of our conscious experiences are not in fact caused by them we should second guess our resistance to eliminativism. After all, it would be hard to explain how we are in epistemic contact with them at all!

The infection that would come with interactionist dualism is probably to be preferred to epiphenomenalism, but the consequences are pretty undesirable here as well. Consider the prospects for neuroscience.

³ Papineau (2002, p.17).

⁴ I have in mind experiments by Libet et al. (1983), Milner and Goodale (1995), and claims made in Ramachandran and Blakeslee (1998) and Wegner (2002). See Blackmore (2004) for an accessible presentation of some of this research.

Presumably many of the psycho-physical causal interactions would be between conscious events and neural events. Neuroscience only talks about the latter. Take a particular neural event N1 that has a non-physical cause C1. Since C1 falls outside of the domain of neuroscience and cannot be detected by its method, either the neuroscientist will think that N1 is uncaused, or she will claim that N1 is directly caused by some previous physical (probably neural) event N*. Either way, the neuroscientist is wrong. The greater the number of psycho-physical interactions, the worse off neuroscience is. The best we could hope for would be a model that allows us to make predictions about neural events, but we would have to give up on any claim that neuroscience was giving us the correct picture of things. The infection does not stop at neurons, however. Neurons are made of physical particles which are themselves made of physical particles. Neurons fire because of charge inequalities which are the result of the movement of ions across the cell membrane. This should be completely determined by facts such as the presence of positively charged sodium and potassium ions, the gating of ion channels, etc. All of these things should be governed by electro-chemical and physical laws, but if N1 really fires because of a non-physical conscious event, and there is no overdetermination, they simply can't be. Somewhere along the line there will have to be an event that either appears uncaused or appears to be caused by something that does not in fact cause it. This means that there is unlawful behavior somewhere, either in the activity of the potassium ions, or in the ion gates, etc. In any case, it will turn out that the electro-chemical laws are falsified by some chemical reactions in the brain. And of course the infection goes deeper, since chemical reactions are governed by physical laws.

The resulting infection is not only troubling because it goes all the way down to the basic sciences. It is most troubling, perhaps, because the infection would not be local. The potassium in our neurons is just like potassium elsewhere, and the atoms that make up that potassium are just like atoms elsewhere. It is not as if the chemical and physical laws hold for all atoms except for the ones in brains. If they do not hold of the atoms in brains, they do not hold of atoms at all. If interactionist dualism is true, our best picture of the world is fundamentally flawed.

The fact that non-physicalist views give us isolation of infection means that we should avoid them if we can. Subjective physicalism offers a way to do this, even as it gives up on a few other physicalist goals such as the

completeness of objective theories such as physics. Objective theories are only incomplete in a limited, epistemic sense, however. Put simply, subjective physicalism is the view that ontologically the world is entirely physical. Everything that there is supervenes upon the basic physical particles and properties. Nevertheless, not every feature of the world can be completely grasped by objective theorizing. This is not because there are ghostly features of the world at the bottom level, as panpsychism might have it, or that there are emergent phenomena at the more sophisticated level of minds. Both of these would be ontological conclusions. The claim is rather that at certain levels of complexity, at the level of knowers, there are ways of grasping the world that are different than those given by objective theorizing and that cannot be attained by such theorizing.

This book divides into three parts. The first, “Defining Physicalism,” sets the stage for the debate. This is important to do from the outset, since many definitions confuse issues of epistemology and the successes of science with the ontological doctrine of physicalism. By doing so, they can make logical space look smaller than it is. The problem of defining the physical has two parts. One is to define what it is for basic properties to be physical, and the other is to establish the relation everything else must bear to those properties if the doctrine of physicalism is true. I address these issues in Chapters 1 and 2. The second part of the book, *The Threat of the Subjective*, argues that all of the ways to resist the knowledge argument against physicalism without becoming a hardliner require denying objectivism. The discussion of the knowledge argument and our surprisingly limited devices in responding to it occurs in Chapter 3, and Chapter 5 teases out the implications for and the proper understanding of “objectivism.” It would seem that embracing objectivism leads rather quickly to the denial of physicalism. The third part of the book argues that this is not in fact the case by discussing several arguments which push in that direction. In Chapter 5, I discuss the conceivability argument and Max Black’s “presentation argument,” exposing along the way the optional metaphysical commitments that make those arguments plausible. In Chapter 6, I return to the knowledge argument—highlighting its common premise with the conceivability argument—and describe how acquaintance can be used to defuse those arguments without that relation itself falling prey to the anti-physicalist arguments. In Chapter 7, the final chapter, I take a look back, considering the costs and benefits of subjective physicalism, and comparing it to other views on the table.

Although I do tend to view everything here as a package, from the definition of physicalism to the role of acquaintance, it should be said that many of the different positions and arguments of the book can stand alone. Those philosophers who are either unbothered by the problem of defining physicalism, or are perfectly satisfied deferring to the science of physics, might have little interest in Part I of the book, for example. It is, I hope, a virtue of my defense of physicalism that it does not depend on anything particularly idiosyncratic about my definitions of that notion. Nothing, so far as I can tell, is hidden in these definitions that allows physicalism too easy a defense. Similarly, though my defense of the acquaintance strategy is couched in a two-dimensional semantics, philosophers averse to such views can probably *mutatis mutandis* preserve the central idea of the acquaintance theory in the semantics of their choosing. The account in this book tends to stay close to the definitions and semantics invoked by the anti-physicalists, not only because I find myself in sympathy with them, but because I am interested in a defense of physicalism that cannot be accused of changing the topic or of ignoring the dualist's insights. Perhaps some will find this too concessive, but if physicalism can be defended while still in a concessive mood, all the better.

This page intentionally left blank

PART I

Defining Physicalism

Why worry about defining “physicalism?” Many philosophers discussing the problem of consciousness don’t bother with the issue, but being clear about what this problem really is requires that we do. The problem of consciousness, in philosophy anyway, is raised by a group of arguments claiming that the existence of consciousness is not compatible with physicalism. If we are interested in whether or not these arguments succeed, indeed if we are interested in understanding what the arguments claim, we should become as clear as possible about what they mean by “consciousness” as well as what is involved in physicalism.

What should we look for in such a definition? If it is to serve our purposes—in evaluating the anti-physicalist arguments—it must at the very least divide the terrain in the way the arguments seem to presuppose. It would be silly to equate “the physical” with “the real,” for example, since the conclusion of these arguments would be self-contradictory. So, we will let our definition be guided by the arguments. Furthermore, if possible our definition should make it as easy as possible to evaluate the premises of the arguments. If the arguments claim that we can imagine a physical duplicate of our world which lacks consciousness, we should know what we are being asked to imagine, and ideally we should have some sense of how to imagine it in the detail required by the arguments. Vague definitions, therefore, which define the physical in terms of the conclusion of a future physics, are of no help. Our definitions should also help to contextualize the concerns of the anti-physicalists in a greater tradition of arguments. They are asking of consciousness, for example, what some philosophers and scientists once asked about life. So though we

should let our definition be guided by the arguments, it should not ignore the fact that the questions they are raising are parts of a larger question.

For these reasons and others, we cannot simply rest with the view that physicalism maintains that everything is physical. Such a flat-footed definition has two shortcomings. One failing is that it is not particularly informative, since it employs one of the notions most in need of clarification. The problem with understanding physicalism isn't really the "-ism," but the "physical." What does it mean for a property or thing to be physical? Many theorists think there is a further problem with this simple definition. Even if we could pin down what it meant to be physical, do we really want it to claim that everything—as in every property or object—is physical? Need it claim that numbers are physical objects, or that the property of being jovial is a physical property? Perhaps, but perhaps not. This will partly depend on how one defines physical, and just how strict one's notion of physicalism is. At the very least, this concern should lead us to be a bit more cautious in advancing our definition.

The problem of defining "physicalism" divides into three parts. First, the definition should specify the domain of physicalism. Should it cover numbers and other abstracta? Should it cover supernatural divine entities or simply concrete things of the sort we discover in scientific investigations? Call this the *domain problem*. Second, it should provide a precise characterization of what it takes for something to be physical in a narrow sense. Call this the *base problem*. Third, it should indicate a relation things must hold to the objects or properties described by the base notion if they are to be called physical in a more general sense. Call this the *relation problem*. So, for example, one could define physicalism as the view that every contingent property is identical with a property posited by current physics. The domain of this thesis is "contingent properties," the base notion is "properties posited by current physics" and the relation is "identity."¹

To see how each of these three parts can raise issues, we can look at some problems with this straw-man definition. First, suppose one were engaged in a debate about the physicality of numbers. This definition would essentially be useless, since it doesn't concern things like numbers. So, the domain one is interested in studying might well affect what is

¹ See Stoljar (2010).