

More information about this series at <http://www.springer.com/series/6540>

Carlos Muñoz-Suárez • Felipe De Brigard
Editors

Content and Consciousness Revisited

With Replies by Daniel Dennett

 Springer

Editors

Carlos Muñoz-Suárez
Departament de Lògica, Història
i Filosofia
Universitat de Barcelona
Barcelona, Spain

Felipe De Brigard
Center for Cognitive Neuroscience
Department of Philosophy
Duke University
Durham, North Carolina, USA

The image included in the Chapter 1, according to the U.S. Copyright Policies, is in the Public Domain due to copyright expiration, because its first publication occurred prior to January 1, 1923. In this case, in 1892.

Studies in Brain and Mind
ISBN 978-3-319-17373-3 ISBN 978-3-319-17374-0 (eBook)
DOI 10.1007/978-3-319-17374-0

Library of Congress Control Number: 2015940543

Springer Cham Heidelberg New York Dordrecht London
© Springer International Publishing Switzerland 2015

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made.

Printed on acid-free paper

Springer International Publishing AG Switzerland is part of Springer Science+Business Media (www.springer.com)

Foreword: Writing *Content and Consciousness*

Oxford in the mid-1960s dominated Anglophone philosophy as never before (and never since), and there were dozens of Americans, Canadians, Australasians, and South Africans (whites, of course, back then) eager to become certified practitioners of the then fashionable ordinary language philosophy. I was as enthusiastic as any, with Ryle's *Concept of Mind* and Austin and Wittgenstein as my beacons, but when I talked with my fellow graduate students, I discovered a disturbing complacency and lack of intellectual curiosity infecting their approaches. I remember in particular a meeting in my first term of the Ockham Society, a graduate discussion group. In the midst of a discussion of Anscombe's *Intention*, as I recall, the issue came up of what to say about one's attempts to raise one's arm when it had gone "asleep" from lying on it. At the time I knew nothing about the nervous system, but it seemed obvious to me that something must be going on in one's brain that somehow amounted to trying to raise one's arm, and it might be illuminating to learn what science knew about this. My suggestion was met with incredulous stares. What on earth did science have to teach philosophy? This was a philosophical puzzle about "what we would say," not a scientific puzzle about nerves and the like. This, it seemed to me, was as weirdly narrow an approach as setting out to learn all about horses by seeing what everyday folk had to say whenever they used the word "horse." It might help, mightn't it, to examine a few horses? My fellow philosophers of mind in Oxford were untroubled by their ignorance of brains and psychology, and I began to define my project as figuring out as a philosopher how brains could be, or support, or explain, or cause... minds.

I asked a friend studying medicine at Oxford what brains were made of and vividly remember him drawing simplified diagrams of neurons, dendrites, and axons—all new terms to me. It immediately occurred to me that a neuron, with multiple inputs and a modifiable branching output, would be just the thing to compose into networks that could learn by a sort of evolutionary process. Many others have had the same idea, of course, before and since. Once you get your head around it, you see that this really is the way—probably, in the end, the only way—to eliminate the middleman, the all-too-knowing librarian or clerk or homunculus who manipulates

the ideas or mental representations, sorting them by content. With this insight driving me, I could begin to see how to concoct something of a “centralist” theory of intentionality. (This largely unexamined alternative was suggested by Charles Taylor in his pioneering book, *The Explanation of Behaviour*.) The result would be what would later be called a functionalist, and then teleofunctionalist, theory of content in which Brentano and Husserl and Quine could all be put together, but at the subpersonal level. (The personal/subpersonal distinction was my own innovation, driven by my attempts to figure out what on earth Ryle was doing and how he could get away with it.) In order to do this right, I needed to learn about the brain, so I spent probably five times as much energy educating myself in Oxford’s Radcliffe Science Library as I did reading philosophy articles and books.

I went to Ryle, my supervisor, with my project, and to my delight and surprise he recommended that I drop the arduous B.Phil. program with its brutal examinations and switch to the D.Phil., a thesis-only degree. I was off and running, but the days of inspiration were balanced by weeks and months of confusion, desperation, uncertainty. A tantalizing source of alternating inspiration and frustration was Hilary Putnam, whose *Minds and Machines* (1960) I had found positively earthshaking. I set to work feverishly to build on it in my own work, only to receive, from my mole back at Harvard, an advance copy of Putnam’s second paper on the topic, “Robots: Machines or Artificially Created Life?” (not published until 1964), which scooped my own efforts and then some. No sooner had I recovered and started building my own edifice on Putnam paper number two than I was spirited a copy of Putnam paper number three, “The Mental Life of Some Machines” (eventually published in 1967), and found myself left behind yet again. So it went. I think I understood Putnam’s papers almost as well as he did, which was not quite well enough to see farther than he could see what step to take next. Besides, I was trying to put a rather different slant on the whole topic, and it was not at all clear to me that, or how, I could make it work. Whenever I got totally stumped, I would go for a long, depressed walk in the glorious Parks along the river Cherwell. Marvelous to say, after a few hours of tramping back and forth with my umbrella muttering to myself and wondering if I should go back to sculpture (my alternative career path), a breakthrough would strike me, and I’d dash happily back to our flat and my trusty Olivetti for another whack at it. This was such a reliable source of breakthroughs that it became a dangerous crutch; when the going got tough, I’d just pick up my umbrella and head out to the Parks, counting on salvation before supertime.

Ryle himself was the other pillar of support that I needed. In many regards, he ruled Oxford philosophy at the time, as editor of *Mind* and informal clearinghouse for jobs throughout the Anglophone world, but at the same time he stood somewhat outside the cliques and coteries, the hotbeds of philosophical fashion. He disliked and disapproved of the reigning Oxford fashion of clever, supercilious philosophical one-upmanship and disrupted it when he could. He never “fought back.” I tried to provoke him, in fact, with some elaborately prepared and heavily armed criticisms of his own ideas, but he would genially agree with all my good points as if I were talking about somebody else and get us thinking of what repairs and improvements we could make together of what remained. It was disorienting, and my opin-

ion of him then—often expressed, I am sad to say, to my fellow graduate students—was that while he was wonderful at cheering me up and encouraging me to stay the course, I hadn’t learned any philosophy from him.

I finished a presentable draft of my dissertation in the minimum time (six terms or 2 years) and submitted it, with scant expectation that it would be accepted on first go. On the eve of submitting it, I came across an early draft of it and compared the final product with its ancestor. To my astonishment, I could see Ryle’s influence on every page. How had he done it? Osmosis? Hypnotism? This gave me an early appreciation of the power of indirect methods in philosophy. You seldom talk anybody out of a position by arguing directly with their premises and inferences. Sometimes it is more effective to nudge them sideways with images, examples, and helpful formulations that stick to their habits of thought.

My examiners were A. J. Ayer and—an unprecedented alien presence at a philosophy “viva” occasioned by my insistence on packing my thesis with speculations on brain science—the great neuroanatomist, J. Z. Young from London. He, too, had been struck by the idea of learning as evolution in the brain and was writing a book on it, so we were kindred spirits on that topic, if not on the philosophy, which he found intriguing but impenetrable. Ayer was reserved. I feared he had not read much of the thesis, but I later found out he was simply made uncomfortable by his friend Young’s too-enthusiastic forays into philosophy, and he found silence more useful than intervention. I waited in agony for more than a week before I learned, via a cheery postcard from Ryle, that the examiners had voted me the degree.

I returned to the United States, to UC Irvine, my first teaching job, age 23. Now it was time to turn my dissertation into articles and a book. I revised the first chapter and sent it out as a journal article. There were a dozen submissions and a dozen rejections, with many revisions in between. Then Wilfrid Sellars, editor of *Philosophical Topics*, wrote me a nice letter saying that he was intrigued by the draft I had sent; once I clarified a few foggy points, he thought it would be fine. I sent him a clarified version within the week, and he wrote back to say that now that it was clear what I was doing, he thought it was not a publishable paper! A few more rejections and I gave up on that chapter and started several other projects, with no greater success. Perhaps I wasn’t going to make it as a philosopher after all.

One day Julian Feldman, an artificial intelligence researcher at UCI, came storming into my office with a copy of Hubert Dreyfus’s notorious RAND memo, “Alchemy and Artificial Intelligence.” What did I make of it? I read it and said I disagreed quite fundamentally with it. “Write up your rebuttal, please, and get it published!” Why not? I wrote “Machine Traces and Protocol Statements” and promptly published it in *Behavioral Science* (1968, my first publication), and my career as philosopher-laureate of AI had begun. I’d already been attracted to the field by Alan Ross Anderson’s pioneering anthology, *Minds and Machines*, and found at Irvine a small group of AI researchers who invited me to join them. Allen Newell came through town to give some talks and struck up a lively conversation with me, and I was hooked. Other colleagues at Irvine, in particular the psychobiologist James McGaugh, struck by my knowledge of, and interest in, theories of learning in neural systems, also took a vigorous interest in further educating me and getting me thinking about their work and its problems.

I decided I had to concentrate on turning the dissertation into a book, and I think these nonphilosophers contributed the most to the improvements, clarifications, and enlargements that distinguish *Content and Consciousness* from the naive stumblings in my D.Phil. dissertation.

In the summer of 1967, I sent the new manuscript to the famous Routledge & Kegan Paul series, the International Library of Philosophy and Scientific Method. This series of books, with their red covers and yellow dust jackets, included most of my favorite books in philosophy: Wittgenstein's *Tractatus*, Smart's *Philosophy and Scientific Realism*, and Sellars' *Science, Perception and Reality*, for instance. A year passed without a word from the new editor, Ted Honderich, who had taken over from A. J. Ayer. I didn't dare upset the applecart by complaining about the unresponsiveness. Finally, when I knew I was off to Oxford for a quarter sabbatical in the fall of 1968, I wrote a timid inquiry to Honderich, who discovered that the manuscript had been mislaid by the referee to whom he had sent it. He retrieved it, read it himself, and forthwith accepted it, pending revisions which I hastened to complete that autumn in Oxford. I was in heaven. But still, I couldn't talk about it to other philosophers. My problem was that my way of approaching the then standard issues in the philosophy of mind was too eccentric, too novel, to afford easy entry into a discussion. When somebody asks you what you're working on, you usually can't back them into a corner and harangue them for a couple of hours about your project, and I could imagine no more modest framing job that might bring interlocutors to where I was. After all, my attempts to publish the first chapter showed that the first ideas I needed to get across were bound to be misunderstood and had already been misunderstood in half a dozen versions by some of the best philosophers of mind in the field. So I was a very lonely and uncertain philosopher those first few years at Irvine, spending more happy hours talking AI or psychobiology than philosophy. In spite of my presumably sterling pedigree as a student of Quine and Ryle, I felt like an outsider, a dark horse candidate that one should probably not bet on. The acceptance of the manuscript by Honderich, and his further invitation to write an essay on free will ("Mechanism and Responsibility," in which I introduced the terminology of the intentional stance and intentional systems), gave me new confidence, however.¹

When *Content and Consciousness* was published, in 1969, J. Z. Young sent me a nice note telling me to ignore the review in the *Times Literary Supplement*, which I hadn't seen until he drew my attention to it. This was my first review, and it was a stinker. Reviews in the *Times Literary Supplement* those days were all anonymous, but years later I learned that it had been written by D. W. Hamlyn, and to my dismay one of his chief criticisms was about the style, which I had thought to be refreshingly unlike other philosophy books of the day. Young's note did cheer me up, however, and soon the book got two wonderful reviews: J. J. C. Smart did a long "Critical Notice" in *Mind*, and R. L. Franklin wrote an even more positive long

¹The preceding paragraphs are drawn, with minor revisions, from my essay "Autobiography," published in *Philosophy Now* (London, July 2008).

review in *Australasian Journal of Philosophy*. These two reviews put my book in the limelight, and soon I began to field inquiries and invitations from all over the Anglophone philosophical world. Gilbert Harman at Princeton was one of the book's first supporters, as I learned when a former student of mine who had gone on to Princeton wrote me a note telling me that my book was being featured in his course. Richard Rorty was another enthusiastic reader. Princeton, in fact, was the epicenter of interest, and I was invited to give a talk there in December of 1970, the first professional talk of my career. (I presented "Intentional Systems" to an audience that included, in addition to Harman and Rorty, Alonzo Church, Donald Davidson, David Lewis, Thomas Nagel, Max Black, and quite a few other luminaries. I was terribly nervous, but the reception was cordial and constructive. Rorty had a reception for me afterwards at his house, beginning a lifelong friendship.)

The book was chosen for an "author meets critics" session at the annual meeting of the Eastern Division of the American Philosophical Association in December of 1972 more than 3 years after it was published—the world moved more slowly then. Michael Arbib, one of the first computational neuroscientists, and Keith Gunderson, philosopher and poet, were the critics. (I think Arbib's and Gunderson's talks were never published, but my response is included as a chapter in *Brainstorms*.) Since my book was thus featured in a symposium, I expected the American publisher, Humanities Press, to have it prominently displayed at their table in the book exhibit room, but to my dismay they didn't have a single copy to show or sell. When I confronted the proprietor with this anomaly, his response aggravated my bad mood: "A symposium? So that's why people have been coming around all day asking for it!" He had no copies because it wasn't a new book. I later learned, moreover, that Routledge & Kegan Paul, which had an arrangement with Humanities Press to print their copies with a different front page bound in, had decided some months before the symposium that my book wasn't going anywhere and had remaindered the rest of their stock to Humanities Press for a dollar a copy. So when the book did take off and have a good sale in the United States, I got the handsome royalty of ten cents a copy. But that didn't matter to me; the book was being read and discussed in courses and seminars.

One of my favorite responses to the book came from Arthur Danto, whom I had not yet met. He sent me a nice note about how much he had enjoyed the book, and learned from it, but then he went on to draw my attention to one of the embarrassing errors in it. I had misexplained Quine's famous example "Giorgione was so-called because of his size." I had supposed that Giorgione meant Little George, not Big George. Danto enclosed a copy of the letter he had just sent to Quine, informing him that in any case Quine was wrong, too! According to Vasari, Danto noted, Giorgione was so-called dalle fattezze dalla persona e dalla grandezza del animo—because of the features of his face and the greatness of his soul. "However," Danto went on, graciously, "it is not my intention to wander either into questions of physiognomy or grammar, but to report a factual error which would be minor in the case of someone who did not bear the awful responsibility of stocking philosophers with what meager facts they may claim." I immediately wrote to Quine, apologizing for butch-

ering his example, but taking some pleasure in Danto's discovery that Quine himself had perpetrated a falsehood. Quine immediately wrote back a friendly letter, enclosing a copy of his reply to Danto:

Dear Danto,

I much appreciated your generous and amusing letter of August 17. I find a perceptible gain in taking on a new fact and getting rid of a dud, whatever the chagrin over being caught out. But the present case leaves me in doubt. My dictionary gives "fattezze" as "bodily proportions," among other things... Vasari softens the blow by ringing the animo in too, but I would set that down to the animo of Vasari.

Over the years Quine delighted in finding embarrassing factual errors in my books, and it became a running joke. I never caught him out in a factual error, but when I sent him a list of the factual errors in *Darwin's Dangerous Idea* (the book dedicated to him) that he had missed, we had a good laugh over it.

It seems to me that for all its flaws, *Content and Consciousness* had enough things right to make it an excellent platform on which to build further philosophical work. Or better, it has been, for me, a sort of philosophical kitchen, stocked with almost all the utensils and containers, all the ingredients and methods, from which I have concocted the rest of my work. And over the years I have enjoyed watching other philosophers gravitating inexorably towards versions of the views I first spelled out there. What seemed outrageous and even incomprehensible to many of my colleagues 40 years ago makes much more sense today.

Medford, MA, USA

Daniel Dennett

Acknowledgments

The Editors: We thank Gualtiero Piccinini, the editor of this excellent book series, for his constant patience and invaluable help. We want to thank also the authors of the articles included in this volume, as well as those who helped review their manuscripts.

Carlos Muñoz-Suárez: I'm grateful to Dan Dennett for his support and willingness to discuss and update his ideas. It has been a pleasure to work with Felipe De Brigard. I thank him for his remarkable patience and outstanding collaboration. I also thank Pete Mandik and Don Ross for their support from the earliest stage of this project. I'm specially grateful to Aura Marina Suárez for her unconditional love and support. To my family and friends, thanks for sailing with me across the seas of love and knowledge. Finally, I thank the PERSP Project (Consolider-Ingenio project CSD2009-00056) for funding my life in Barcelona.

Felipe De Brigard: First, I would like to thank Dan Dennett for being a source of philosophical inspiration and for his constant support throughout the years. Thanks also to Carlos Muñoz-Suárez, my coeditor, for getting this project going. Finally, I would like to thank my wife, Anne, my son, David, the one the way, and our dog, Jeffers, for their unending love. You guys are awesome.