University of Pennsylvania
ScholarlyCommons

Publicly Accessible Penn Dissertations

2017

The Ethics Of Meaning

Robert B. Willison
University of Pennsylvania, willison@gmail.com

Follow this and additional works at: https://repository.upenn.edu/edissertations

Part of the Ethics and Political Philosophy Commons, and the Other Languages, Societies, and Cultures Commons

The Ethics Of Meaning

Abstract

This dissertation develops an ethics of meaning. In the first chapter, I offer an account of meaning that comprehends its many varieties—natural, cultural, linguistic, literary, and ethical meaning, for example—by appeal to the structural role meaning plays in the practice of interpretation. In Chapter 2, I develop a distinctive account of the concept of ethical meaning ("meaning" as it's used in the phrase "the meaning of life"). In Chapter 3, I develop a new account of irony on the basis of the comprehensive-interpretive account of meaning introduced in Chapter 1.

Degree Type
Dissertation

Degree Name
Doctor of Philosophy (PhD)

Graduate Group
Philosophy

First Advisor
Michael Weisberg

Keywords
Ethics, Irony, Meaning

Subject Categories
Ethics and Political Philosophy | Other Languages, Societies, and Cultures | Philosophy

THE ETHICS OF MEANING

Robert B. Willison

A DISSERTATION

in

Philosophy

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2017

Supervisor of Dissertation

_____

Michael Weisberg

Professor of Philosophy

Graduate Group Chairperson

_____

Samuel Freeman, Avalon Professor of the Humanities

Dissertation Committee

Elisabeth Camp, Associate Professor of Philosophy, Rutgers University
Karen Detlefsen, Professor of Philosophy and Education, University of Pennsylvania
Alex Guerrero, Associate Professor of Philosophy, Rutgers University

For my mother, who makes so much of the meaning she finds everywhere

ACKNOWLEDGMENTS

could I have hoped for a more wonderful committee. Alex Guerrero read and discussed my work with great care and terrific intelligence; I always left his office with new ideas and renewed excitement for philosophy. He was also unfailingly thoughtful and kind. Adrienne Martin has always inspired me with her courage, honesty, and toughness. Michael Weisberg seems to pass part of what's best about him to each of his students—Emily Parke's conscientious respect for science as it's actually practiced; Karen Kovaka's broad knowledge and social conscience; Carlos Santana's fearlessness: I'm glad to join their ranks. Karen Detlefsen is a deeply moral, insightful, resilient philosopher, and an imposing role model. She works so hard to do what's right, and succeeds.

In addition to the members of my committee, Nat Adams, John Arthur Baldwin, Devin Curry, Sam Graham-Felsen, David Hills, Abe Kunin, Ben Mathis-Lilley, Sucosh Norton, Shane Slattery-Quintanilla, and Eric Todrys read and commented on portions of this dissertation. Thanks to them all.

Kierkegaard said that being human doesn't come easily, but it's even harder when you share your life (and, especially, your living space!) with me. I'm surpassingly grateful, therefore, for the profound tolerance, uncompromising loyalty, and soulful companionship of Louise Daoust. Louise's philosophical influence is everywhere in this dissertation, but especially at its foundation, in Chapter 1. I owe much of my interest in and understanding of Gibson and Dretske to Louise. More importantly, though, there's no more philosophically important or fulfilling activity than sharing time and conversation with a genuinely original, miraculously good human

being. I couldn't have asked for a better teammate. She's like Kevin Garnett and Isaiah Thomas put together.

Somewhere in their acknowledgements, philosophers usually stipulate that the people thanked shouldn't be held responsible for any mistakes the author has made. I want explicitly to except my mother from any such disclaimer: her influence on me is far too thorough for her to escape responsibility for my failings. If, in the pages that follow, I sometimes see connections where there are none, become melodramatic in my pessimism, or insufficiently rigorous in my moments of optimism: blame her. To my eye, these appear as the occasionally necessary excesses of a great spirit rather than as vices, because they belong to her. But even if you don't share my bias, I hope that these imperfections, to the extent that they are present, are outweighed by the virtues my mother taught me: a deep charity toward the strangest of views, a refusal to dismiss the unpleasant, and faith in humanity.

If I can be understood, it's in a language of my mother's making. I am who I am because of her. Given all of that, it's of comparatively little consequence that I dedicate this dissertation to her, with my deepest love and gratitude. Nevertheless, I do.

ABSTRACT

THE ETHICS OF MEANING

Robert Willison

Michael Weisberg

This dissertation develops an ethics of meaning. In the first chapter, I offer an account of meaning that comprehends its many varieties—natural, cultural, linguistic, literary, and ethical meaning, for example—by appeal to the structural role meaning plays in the practice of interpretation. In Chapter 2, I develop a distinctive account of the concept of ethical meaning ("meaning" as it's used in the phrase "the meaning of life"). In Chapter 3, I develop a new account of irony on the basis of the comprehensive-interpretive account of meaning introduced in Chapter 1.

TABLE OF CONTENTS

PREFACE
The Ethics of Meaning

On September 10, 1989, two days before her son's eighth birthday, Ann Weiss drove to the local shopping center to order a cake for the party. She chose chocolate, her son's favorite, and had his name, "SCOTTY," inscribed under an image of a spaceship and launching pad, with a backdrop of frosted white stars.

On the morning of his birthday, walking with a friend to catch the school bus, Scotty was hit by a car. Though he managed to walk home on his own power and tell his mother what had happened, he soon fell into a deep, unresponsive sleep. Ann and her husband, Howard, took Scotty to the hospital, where, after conducting some tests, the doctor assured both parents that Scotty was simply suffering from a concussion and a shock, and would wake soon.

After twenty four hours, Scotty still wasn't awake. This was no cause for serious concern, Dr. Francis said: they'd run extensive tests, and, aside from a small hairline fracture in the skull (not itself life-threatening), nothing seemed to be wrong. Extended sleeps like these can be a symptom of shock. Both Dr. Francis and Howard encouraged Ann to take a break from her vigil at Scotty's bedside—to return home, have something to eat, take a shower. (Howard had done the same a few hours earlier.) Sensing that Howard wanted some time alone—a respite from worrying and talking together—Ann reluctantly agreed. Maybe if she wasn't watching Scotty every second, she thought, he'd finally wake up and be alright.

When Ann returned to the hospital a few hours later, just after sunrise, Scotty finally

opened his eyes. He glanced at his parents, who were speaking to him tenderly, but showed no

sign of recognition. Then, "his mouth opened, his eyes scrunched closed, and he howled until he

had no more air in his lungs. His face seemed to relax and soften then. His lips parted as his last

breath was puffed through his throat and exhaled gently through the clenched teeth." The doctors

later told Ann and Harry that Scotty's death had been caused by a hidden occlusion: a "one-in-a-

million circumstance."

> The doctor walked them to the hospital's front door. People were entering and leaving the
> hospital. It was eleven o'clock in the morning. Ann was aware of how slowly, almost
> reluctantly, she moved her feet. It seemed to her that Dr. Francis was making them leave
> when she felt they should stay, when it would be more the right thing to do to stay. She
> gazed out into the parking lot and then turned around and looked back at the front of the
> hospital. "No, no," she said. "I can't leave him here, no." She heard herself say that and
> thought how unfair it was that the only words that came out were the sort of words used
> on TV shows where people were stunned by violent or sudden deaths. She wanted her
> words to be her own.[1]

§1 Introduction

If you followed the endnote, you'll know that the story I just told is a fiction, adapted from a short

story by Raymond Carver. I began that way for two reasons. First, the scene evokes the problem

that motivates the three essays collected in this dissertation. As Stanley Cavell described it,

> The problem is no longer how to do what you want, but to know what would satisfy you.
> We could also say: Convention as a whole is now looked upon not as a firm inheritance
> from the past, but as a continuing improvisation in the face of problems we no longer
> understand. Nothing we now have to say, no personal utterance, has its meaning
> conveyed in the conventions and formulas we now share. In a time of slogans, sponsored
> messages, ideologies, psychological warfare, mass projects, where words have lost touch
> with their sources or objects, and in a phonographic culture where music is for dreaming,
> or for kissing, or for taking a shower, or for having your teeth drilled, our choices seem to

be those of silence, or nihilism (the denial of the value of shared meaning altogether), or statements so personal as to form the possibility of communication without the support of convention—perhaps to become the source of new convention. And then, of course, they are most likely to fail even to seem to communicate. Such, at any rate, are the choices which the modern works of art I know seem to me to have made.[2]

Second, I think that beginning as I did was an ethical risk. I'm not sure that what I did was right, or fair. For instance, I tried to disguise my source material: the precise date at the beginning of the story is my fabrication, not Carver's. I hoped it would lead readers to wonder whether I was relaying a news story. I didn't quote directly until the end of my adaptation, when I reached the point that mattered for my own rhetorical purposes. I drew on someone else's imaginative brilliance to begin a piece that should be my own. Finally—and, I think, most troublingly—the subject is so difficult. When I read it, in Carver's words, I feel terrified. Sometimes I cry. I empathize with Ann, and I know the desperate frustration she feels at the emptiness of her culture's expressive resources matters much less to her, standing in that hospital parking lot, than it does to Carver, or to me. She has deeper sorrows.

Writers want to draw their readers into the emotional universe that animates their efforts. Readers want to feel the person on the other side of the text, and to be acknowledged as partners in communication. But the line between craft and manipulation can be difficult to locate, and it's hard not to take more than you deserve. My deliberation about how to begin this dissertation was one of the countless exercises in the ethics of meaning that we owe to one another as members of a shared community of meaning. This dissertation is about those obligations, and the nature of the relationships that generate them. I don't think I've earned my opening. I hope in this case I was right to conclude, finally, that it would be alright to ask for your generosity.

§2 Methodology

In "Philosophy and Civilization," John Dewey argues that the work of philosophers is tied, inextricably, to the predicaments of their own cultural and historical moment. The movement of time, and the study of philosophy's history, he says, have revealed as illusion the conviction of past philosophers that their own work had finally transcended local bias, and captured eternal truths. "Philosophers are parts of history, caught in its movement; creators perhaps in some measure of its future, but also assuredly creatures of its past."[3] If this is true, Dewey thought, philosophers should embrace their obligation to help their societies negotiate "the enduring junctions of a stubborn past and an insistent future," recognizing that their "problems and subject matter grow out of the stresses and strains in the community life in which a given philosophy arises."[4] Bertrand Russell admired Dewey, but disdained what he thought was an abdication of philosophical ambition. Truly "scientific" philosophy, Russell argued,

> is something more arduous and more aloof, appealing to less mundane hopes, and requiring a severer discipline for its successful practice…Philosophy is general, and takes an impartial interest in all that exists. The changes suffered by minute portions of matter on the earth's surface are very important to us as active sentient beings; but to us as philosophers they have no greater interest than other changes in portions of matter elsewhere.[5]

Russell's position is out of style, but I think they both have a point. Still, it's easier (and less lonely) for me to imagine filling the role that Dewey prescribes. A Deweyan philosopher, then, should be alert to "the stresses and strains" in her community life—the challenges that new developments pose to our old strategies for understanding ourselves and cooperating with one another. A philosopher who shares this Deweyan conception, and who wants to do useful work, should

therefore begin by asking: what are today's new problems, which require of us adjustment,

clarification of meaning, new understanding?

§3 Can We Mean What We Say?

I read David Foster Wallace as having developed one particularly compelling answer to that

question. His novel, *Infinite Jest*, begins with a waking nightmare. Hal, one of the novel's two

main characters, is being interviewed for an athletic scholarship by a trio of university deans.

Narrating the scene for the reader, Hal is grotesquely articulate. The room's polished conference

table, he tells us, shines "with the spidered light of an Arizona noon," and bits of dust and

sportcoat lint, excited by the air conditioner vents, "dance jaggedly in the slanted plane of

windowlight, the air over the table like the sparkling space just above a fresh-poured seltzer." The

dean to Hal's left wears a smile with "the impermanent quality of something stamped into

uncooperative material"; the dean at center "seems to have more than the normal number of

eyebrows." But, within the scene, Hal suffers from a terrible condition. Whenever he attempts to

communicate, he appears to his audience to be suffering an awful psychotic episode. After

witnessing Hal attempt to speak, one dean likens him to "a writhing animal with a knife in its

eye"; another compares the experience to watching "a stick of butter being hit by a mallet."

"Try to listen," Hal pleads, as the panicked deans, trying to immobilize him, push his face

into the room's parquet floor. "I'm not a machine. I feel and believe… I am not what you see and

hear."[6]

The scene is a dramatic expression of a strain of meaning skepticism that runs

throughout David Foster Wallace's work: he worried that many of us were coming to share Hal's

condition. Midway through the 20th century, Stanley Cavell famously wondered if we must mean what we say. By the century's end, Wallace had replaced Cavell's "must" with his own "can".

Specifically, Wallace's skepticism is *comprehensive*, *historical*, and *tragic*. *Comprehensive* because the skepticism is of "meaning" across its many senses: not just linguistic meaning, but, for example, the kinds of meaning we often speak of lives and of artworks as having. *Historical* because the skeptic's claim isn't *eternal*: he doesn't claim that the target of his skepticism never was; he claims, instead, that it no longer is. Finally, *tragic*, because unlike a *benign* skepticism, in which the skeptic reassures us that we are just as well or better off once freed from the tyranny of his target, conceding to the tragic skeptic means acknowledging a tragic loss.

§4 Contents

The title of this dissertation is at least as aspirational as it is representative. Nonetheless, the three essays that follow form the beginning of a response to the predicament Wallace dramatized.

Chapter 1, **The Comprehensive-Interpretive Account of Meaning**, develops a theory of meaning, grounded in the practice of interpretation, that that integrates meaning's several senses (e.g. natural, linguistic, and ethical meaning), and provides the conceptual foundation for an ethics of meaning. A fully comprehensive account should show, for example, how the kind of meaning a sentence can have relates to the kind of meaning a life can have, and it should illuminate what the enterprise of historical natural science, concerned to uncover truths about long-past token events on the basis of the traces they have left, shares with the enterprise of cultural anthropology, whose hope is to make sense of "social expressions on their surface enigmatical."[7] (It should show, that is, the connection between natural and cultural meaning.)

The comprehensive-interpretive account of meaning I develop in Chapter 1 lays the foundation for both of the dissertation's subsequent essays, on the ways we use "meaning" in ethical discourse, and on the nature and ethics of irony.

Chapter 2, on **Ethical Meaning**, applies the comprehensive account of meaning presented in Chapter 1 to advance our understanding of the way we use meaning, in ethical discourse, to articulate a particular set of aspirations and anxieties—as when we hope to lead meaningful lives, or worry that our lives our meaningless. The account of meaning I advocate explains the *depth* and *elusiveness* commonly associated with ethical meaning, and complicates the common assumption that meaning in life is necessarily desirable or valuable. It clarifies the sense in which ethical meaning is "objective," and recasts the relationship between meaning and characteristic subjective states (like fulfillment in its presence, and alienation or anxiety in its absence) as symptomatic rather than constitutive. It also clarifies the relationship between meaning, moral duty, and value in general. The chapter concludes with a discussion of the processes by which we come to find and to lose meaning in our own lives.

Finally, the third chapter, **Irony as Dynamic Reversal**, develops a new account of irony. The dynamic reversal account of irony improves on the leading contemporary accounts by distinguishing irony's form from its mechanism, providing a basis for continuity between irony's kinds (e.g. verbal, dramatic, and situational irony), and offering a richer understanding of irony's uses. One important insight that arises from these clarifications is that ironists can be sincere in ways that traditional accounts, which identify irony with the insincere mechanisms often used in its production, obscure. This, in turn, suggests that irony has ethical potential that's

unappreciated by critics like Wallace, who famously considered it "an agent of great despair and stasis in U.S. culture."[8] As Jonathan Lear has recently argued, irony is also a fundamental mode of human experience.[9] The dynamic reversal account of irony shows how irony in language can be used sincerely to understand and share ironic experience. It connects the discourse about irony in language with discourses about irony in culture, politics, literature, and self-understanding, giving us a more complete picture of irony's nature and uses.

§5 Conclusion

Describing his musical ambitions in an interview, John Coltrane said:

> I want to be able to bring something to people that feels like happiness. I would love to discover a process such that if I wanted it to rain, it would start raining. If one of my friends were sick, I would play a certain tune and he would get better; if he were broke, I would play another tune and immediately he would receive all the money he needed. But what those pieces are, and what way do you have to go to arrive at knowing them, I don't know.[10]

Originally, I'd hoped this dissertation would give a definitive answer to Wallace's anxieties about meaning disintegration. Instead, it manages only to lay the groundwork for a response. In the meantime, there's some consolation to be found in Viktor Frankl's observation, that "insofar as the feeling of meaninglessness is concerned, we should not overlook and forget that, per se, it is not a matter of pathology; rather than being the sign and symptom of a neurosis, it is, I would say, the proof of one's humanness."[11]

Still, it can't hurt to hope for something more. My own hope is that the project I've started here will be one small contribution to a joint effort that renders our community of meaning less like the one Hal experienced on the office floor, being smothered by the panicked

dean with too many eyebrows, and more like the one Dewey described in one of my favorite

passages in philosophy:

> Not only is social life identical with communication, but all communication (and hence all genuine social life) is educative. To be a recipient of a communication is to have an enlarged and changed experience. One shares in what another has thought and felt and in so far, meagerly or amply, has his own attitude modified. Nor is the one who communicates left unaffected. Try the experiment of communicating, with fullness and accuracy, some experience to another, especially if it be somewhat complicated, and you will find your own attitude toward your experience changing; otherwise you resort to expletives and ejaculations. The experience has to be formulated in order to be communicated. To formulate requires getting outside of it, seeing it as another would see it, considering what points of contact it has with the life of another so that it may be got into such form that he can appreciate its meaning. Except in dealing with commonplaces and catch phrases one has to assimilate, imaginatively, something of another's experience in order to tell him intelligently of one's own experience. All communication is like art. It may fairly be said, therefore, that any social arrangement that remains vitally social, or vitally shared, is educative to those who participate in it.[12]

Or, as one Paolo Freire's students put it: "Before this, words meant nothing to me; now they speak

to me and I can make them speak."[13]

[1] Raymond Carver, "A Small, Good Thing," in *Short Cuts: Selected Stories*, 1st edition (New York: Vintage, 1993), 112–14.

[2] Stanley Cavell, "Music Discomposed," in *Must We Mean What We Say?: A Book of Essays* (Cambridge University Press, 2002), 201–2.

[3] John Dewey, *Philosophy and Civilization* (P. Smith, 1968), 3–4.

[4] John Dewey, *Experience and Nature* (Courier Dover Publications, 1958), vi.

[5] Bertrand Russell, *Our Knowledge of the External World* (New York: Routledge, 1993), 26.

[6] David Foster Wallace, *Infinite Jest: A Novel* (Little, Brown and Company, 1996), 3–15.

[7] Carol E. Cleland, "Prediction and Explanation in Historical Natural Science," *The British Journal for the Philosophy of Science* 62, no. 3 (September 1, 2011): 551–82; Clifford Geertz, "Thick Description: Toward an Interpretive Theory of Culture," in *The Interpretation of Cultures: Selected Essays* (Basic Books, 1973), 5.

[8] David Foster Wallace, *A Supposedly Fun Thing I'll Never Do Again: Essays and Arguments* (Little, Brown and Company, 1997), 49.

[9] Jonathan Lear, *A Case for Irony* (Harvard University Press, 2011).

[10] Jamie Howison, *God's Mind in That Music: Theological Explorations through the Music of John Coltrane* (Wipf and Stock Publishers, 2012), 197.

[11] Viktor E. Frankl, *Man's Search for Meaning* (Beacon Press, 2006), 141.

[12] John Dewey, *Democracy and Education* (Free Press, 1997), 5–6.

[13] Quoted in Richard Shaull, "Foreword," in *Pedagogy of the Oppressed: 30th Anniversary Edition* (Bloomsbury Publishing USA, 2014).

THE COMPREHENSIVE-INTERPRETIVE ACCOUNT OF MEANING

Understanding "Meaning" across Discursive Contexts

§1 Problem(s) of Meaning

"Meaning" enjoys a broad range of usage in everyday discourse—so broad that it can present

ambiguously even in a single context. Consider the following late-night electronic chat between

Alice and James, who are old friends. As usual, Alice is having problems with her on-again, off-

again boyfriend, Larry:

BAD BOYFRIEND

  Alice: arthur and i were reminiscing last night and i just felt like wow, i am old.
  James: yeah, you are old, there's no denying it.
  you're almost as old as I am
  Alice: i know.
  James: we both should already have kids and be preparing for death
  Alice: no shit.
  i need a baby daddy. :)
  James: luckily advanced technologies have facilitated long empty lives for us utterly
divorced from nature or meaning
  Alice: just find the right girl -- and throw yourself into it and love her with all your
might. thats what i'd advise you
James: sure I'd be happy to. "find a billion dollars on the street and then just SPEND THE
SHIT OUT OF IT"
  I mean, OK, will do.
  Alice: thats what kills me with larry.
  i found the billion dollars
James: yeah, that is truly tough
  Alice: and so did he - i question many things but i don't question how
much he leaves me. that part seems pretty true.
  and when we are together its like lighter than light
  James: pretty bad Freudian slip there Alice
  possibly the worst one I've ever seen, no offense

What should we say is the meaning of

(1) i question many things but i don't question how much he leaves me.

—Alice's Freudian slip? There's a perfectly respectable sense in which **(1)** means that Alice doesn't question how much Larry leaves her. But, in the context of the conversation, we, like James, recognize easily enough that's not what Alice really *meant* by **(1)**—she meant, instead, that she doesn't question the strength of Larry's *love* for her (the implication being that Larry's love is strong, whatever other problems they may have). But if we ask Alice's therapist, *she* might tell us that **(1)** really means that, despite Alice's protests, Alice isn't secure in Larry's love at all. A psycholinguist, meanwhile, might take **(1)** to mean that Freudian slips can occur in written as well as spoken discourse. Yet few if any of these various interpretations of **(1)**'s meaning are genuinely at odds: the psycholinguist, depending on her sensibilities, might be perfectly happy to assent to the therapist's interpretation of **(1)** without feeling obligated to recant her own; and the therapist *relies* on **(1)**'s literal meaning to discern its true psychological import.

Things don't end there. Just as a single object (like Alice's utterance) can mean many different things, in many different ways, so lots of different sorts of things (not just linguistic expressions) can mean—at least if everyday language is any guide. A deceased father's wristwatch means a lot to his daughter; the dark clouds in the distance mean to the hikers that it'll rain later that afternoon. Events can have meaning, too, in the way that Caesar's crossing of the Rubicon meant the end of the Roman Republic. And, as one of James's cynical remarks above suggests, we sometimes think of human lives as having a kind of meaning, as well.

Despite the usual pressures, in theoretical discourse, for terms to be used more narrowly, "meaning" has managed to replicate its apparent versatility in philosophical discourse. In his 1954 Presidential Address to the Western Division of the American Philosophical Association, Paul Henle declared that "the problem of meaning is the problem of this century," marshaling as evidence the centrality of the theory of meaning to both pragmatism and positivism, as well as the powerful way the concept of meaning was being used as "the key to the solution of other problems."[1] Henle's claim hasn't lost much plausibility since the century's turn: meaning still confronts philosophers of language and mind as a central problem (or, as we'll see later, as a set of related problems), and philosophers across the sub-disciplines continue to regard meaning as a powerful tool for addressing other problems. In ethics, for example, Susan Wolf finds in meaning the right notion to illustrate the poverty of prevalent models of moral psychology, while Kwame Anthony Appiah uses meaning to explain one important way in which social identities like race, gender, and nationality matter for our ethical projects.[2] Meaning has even enjoyed a theoretical renaissance beyond philosophy's institutional borders. As the cultural anthropologist Clifford Geertz observed in 1973 (with a sense of chagrin that belied his own enthusiasm), "Meaning, that elusive and ill-defined pseudoentity we were once more than content to leave philosophers and literary critics to fumble with, has now come back into the heart of our discipline. Even Marxists are quoting Cassirer; even positivists, Kenneth Burke."[3]

Confronted by a concept with such broad purchase and enduring significance, however, we might reasonably suspect that meaning has won and held its "problem of the century" status by adopting the same attitude of benign indifference to issues of integrity that politicians find

useful when they're running for office. Instead of a single, crucial "problem of meaning," then, the suggestion is that there's a gaggle of problems of meaning, loosely and unsystematically related—a different face for each constituency. Jerry Fodor, for instance, thought the problem of meaning was to answer the question "How can anything manage to be about anything?" and argued that this question applies, in the first place, to individual mental representations.[4] On this view, the issue of linguistic communication becomes a secondary, largely derivative affair. Donald Davidson, in contrast, thought the right way to construct a theory of meaning was to begin by asking for a theory of what a competent speaker of a particular natural language would need to know to successfully interpret any meaningful string of that language, taking special care to account for the fact that language users have finite mental resources, but natural languages allow for the production of an infinite set of meaningful strings.[5] Both versions of the problem sound daunting until we consider Max Weber's: to meet the disenchanted intellectual's "inner need…to endow his life with a pervasive meaning, and thus to find unity within himself, with his fellow men, and with the cosmos."[6] Good luck with that one.

The worry is analogous to one that Appiah recently expressed about "culture," a notion he thinks "has attained ubiquity: but at the cost of conceptual purchase. What shall it profit a word if it shall gain the whole world and lose its own soul?"[7] Is there, on top of the apparently many problems of meaning, a problem of "meaning"?

§1.1 Two Approaches: Restrictive and Ecumenical

Broadly, there are two ways to address the worry. One is to take a *restrictive* approach, limiting meaning's theoretical ambit to a single, paradigm domain. (Once we have an account that works there, we may or may not find it worthwhile to extend the notion, perhaps metaphorically, to less central domains.) In philosophy, restrictivism was already the dominant trend by 1954, when Henle noted that all of the approaches to meaning he intended to discuss "take language as the typical or even the only case of meaning to be considered."[8] As Quine put the point in his pithy survey of the philosophical origins of the problem of meaning in "Two Dogmas of Empiricism": "Only linguistic forms have meanings."[9]

Alternatively, we can take an *ecumenical* approach to meaning, taking seriously the breadth of "meaning"'s everyday use, and trying to develop a theory of meaning consonant with that versatility. [10] From this perspective, linguistic meaning is a special case of a broader category, and should be explained, at least in part, in terms of the features of that broader category. John Dewey's work offers a useful illustration of an ecumenical approach. In *Democracy and Education*, for example, Dewey writes as if all sorts of nonlinguistic entities have meanings, including experiences, actions, and objects. In order to properly communicate an experience we've had, he says, we need to "get outside of it, seeing it as another would see it, considering what points of contact it has with the life of another so that it may be got into such form that he can appreciate its meaning."[11] Later, he distinguishes mental acts from mere "adjustment[s] to a physical stimulus" according to the fact that the former "involve response to a thing in its *meaning*"—so, when someone "bump[s] into a stone," he "kicks[s] it to one side purely

physically"; but when he "put[s] it to one side for fear some one will stumble upon it," he "respond[s] to a meaning which the thing has."[12]

You might wonder, however, if we need to choose between these approaches at all. Isn't keeping in mind that both theorists and everyday speakers use "meaning" in different ways all that's necessary? When Fodor and Zenon Pylyshyn, for example, assert that "like the Loch Ness Monster, meaning is a myth," we shouldn't read them as casting aspersions on W. E. B. Du Bois's ambition, in *The Souls of Black Folk*, to "show the strange meaning of being black here in the dawning of the Twentieth Century."[13] (They only want to cast aspersions on philosophers who persist in believing that concepts have senses, or intensions.) Solving the problem of "meaning," then, might require nothing more than some good common sense and a willingness to restrain your inner middle-schooler, resisting the temptation to contrive substantive disputes from harmless semantic variance.

This is certainly a sensible prescription, but there are interesting dimensions of the problem of "meaning" that it doesn't address. We can grant that Du Bois would feel perfectly comfortable continuing to write about meaning, in his sense, even if Fodor and Pylyshyn's skepticism about meaning, in their sense, were vindicated. Philosophers' appetite for petty intellectual combat notwithstanding, we can even be reasonably confident that Fodor and Pylyshyn would be comfortable letting Du Bois go on as he did, as well. But Dewey has good reason to resist the sort of restrictivism we saw Quine express above. For Dewey, there's nothing incidental about the fact that language, objects, and experience (even in the rich, general sense in which, for example, some people have experienced being black in America) can all be described as

6

meaningful: in fact, Dewey's naturalistic account of linguistic meaning relies on that conviction. Conversely, from a restrictive perspective, the ecumenical ambition to provide a theory of meaning that accounts for anything approaching the full range of its everyday application seems likely to confuse more than it clarifies, creating an unhelpful incentive to glide thoughtlessly over important distinctions that everyday usage too rarely has a reason (or the resources) to notice. The problem of "meaning", then, is a problem in what Alexis Burgess and David Plunkett have called "conceptual ethics"—a question about which concepts we ought to use and (depending on your views about the relationship between words and concepts) which concepts we should use our words to express.[14]

§1.2 Agenda

My goal in this chapter is to vindicate the ecumenical approach to the problem of "meaning" by providing an account of "meaning" that comprehends its various discursive modes. I contend, with Dewey, that the apparently diverse usage of "meaning" in everyday discourse is grounded in a deeper conceptual integrity; and, further, that appreciating the character of that integrity helps us better understand the special sorts of meaning—of objects, actions, artworks, events, sentences, lives—that are so related. Nonetheless, I hope also to accommodate the genuine insight that, I presume, provides much of the motivation for restrictivist approaches: the recognition that particular varieties of meaning, like linguistic meaning, raise important philosophical questions that can't be answered simply by reflecting on the character of meaning in general. Instead,

answering these questions requires an appreciation of the special character of the variety of meaning in question.

The substance of my effort, below, is divided into two major parts. In the first, §2, I develop my comprehensive account of meaning, arguing that the basis of meaning's conceptual integrity across its diversity of usage is the structural role it plays in the practice of interpretation. (For this reason, I call it the "comprehensive-interpretive account of meaning"—the "CI account" for short.) Roughly: meanings are the outputs of acts of interpretation.[15] The second part, §3, aims to demonstrate that the CI account achieves the integration it advertises, focusing on its application to (and implications for) so-called "natural" meaning, and meaning in philosophy of language and mind. A subsequent chapter discusses "meaning" as it's used in ethical contexts.

§2 The Comprehensive-Interpretive Account of Meaning

Dewey's strategy for integrating the varieties of meaning is a familiar one, though he isn't its familiar source. "It is the characteristic use to which the thing is put," he says, "which supplies the meaning with which it is identified…meanings spring up because both persons are engaged as partners in an action where what each does depends upon and influences what the other does."[16] Here, Dewey anticipates the organizing principle of one of the most important movements in 20[th]-century philosophy of language: that we ought to think of meaning in terms of activity rather than ontology. According to this suggestion, understanding meaning requires, in the first place, characterizing a particular kind of action, rather than a peculiar sort of entity. Put grammatically, "meaning" as a verb is explanatorily prior to "meaning" as a substantive. [17]

The comprehensive account I propose below preserves the spirit of this suggestion, but treats *interpretation* rather than meaning itself as the fundamental practice in whose terms a theory of meaning is best articulated. As we've seen, even in its verb form, "mean" can take a host of (grammatical) subjects that aren't themselves agents (as when written words or tracks in the snow mean), and the way those subjects mean doesn't always derive straightforwardly from the way agents do or can mean *by* them. Many, in fact—like most instances of tracks in the snow—aren't even vehicles by which any agent means anything at all.[18] Instead, what unites meaning across its varieties is the distinctive role it plays in the practice of interpretation. Because this is the fundamental commitment of the CI account of meaning, the substance of the CI account depends on a detailed account of the nature of interpretation.

In short: *interpretation is a functionally structured, mentally mediated activity; undertaken by organisms in an environment for the purpose of guiding action, belief, or affection; subject to norms of validity and fruitfulness*. I'll unpack this definition in steps.

§2.1 Dial "M" for Meaning

This scenario is drawn, with slight modification, from the plot of *Dial 'M' for Murder*:

> **BROKEN LOCK:** Tony Wendice, a retired professional tennis player, wants to dispose of his wife, Margot. He knows Margot has been having an affair with Mark, an American mystery novelist, but it's only Margot's personal wealth that keeps Tony living in the style to which he's become accustomed, so divorce isn't a good option. He therefore arranges for Swann, an old college acquaintance with a checkered past, to murder Margot. On the fateful night, Tony invites Mark (who makes a poetic alibi) to a stag party, leaving Margot home alone. Tony's provided Swann with a key, so that Swann can enter the apartment quietly after Margot has retired for the night, hiding, strategically, behind the telephone. Then, at a pre-arranged time, Tony calls Margot at home, waking her and luring her to the phone, where Swann is well-placed to strangle her. Tony has already broken the lock

on a large window that leads out to the garden, so that the killing will appear to have taken place in the course of a random robbery.

Unfortunately for Tony, Margot fights Swann off, ultimately killing him with a pair of sewing scissors. Unpleasantly surprised when he hears Margot's panicked voice return on the other end of the line, Tony doesn't lose his head. He instructs Margot not to do anything until he arrives, races home, and retrieves his spare latchkey from Swann's pocket before calling the police. (He also formulates a new plan: to frame Margot for Swann's murder.) Chief Inspector Hubbard of Scotland Yard is tasked with unraveling the circumstances of Swann's death.

Interpreting the evidence at the scene of the crime is one of the most important strategies Hubbard can employ to accomplish his task. The broken window lock, for example, is a salient detail—one that can be interpreted in several ways. One natural conclusion is the one that Wendice initially hoped Hubbard would reach, that

> **(2)** the window was Swann's method of entry into the apartment.

But Hubbard dismisses **(2)** on the basis of other evidence. The night of the killing was a rainy one; had Swann trudged through the garden to enter through the window, he'd have left footprints in the wet ground outside, and mud stains on the floor inside—but neither are present. Furthermore, Hubbard matches fibers left on Swann's shoes with the doormat outside the flat's front door, where Swann wiped his feet before entering. Hubbard therefore concludes, instead, that the broken lock means

> **(3)** someone attempted to create the misleading appearance that Swann entered by way of the window.[19]

Working from this relatively simple scenario, we can unpack in phases the features of interpretation I listed at the outset of §2.

§2.2 Agency and Environment

Interpretation is *agential*—undertaken by organisms, situated in an environment, who possess some general set of goals in light of which at least some of their behavior (including their interpretive activity) can be understood.[20] We interpreters, in Dewey's words, are "beings who suffer and endeavor." "Since we are creatures with lives to live, and find ourselves within an uncertain environment, we are constructed to note and judge in terms of bearing upon weal and woe—upon value."[21] Acts of interpretation, then, take place in an *evaluative context*: a complex of the cognitive, affective, or practical goals that motivate or rationalize an act of interpretation, along with the constraints and potentialities presented by the interpreter's environment. In BROKEN LOCK, for example, some salient features of the evaluative context are (1) Hubbard's cognitive goal—to uncover more completely the circumstances of Swann's death; (2) the potential unreliability of the testimony of the situation's chief living participants, who, in many possible scenarios, would have good reason to lie; and (3) the presence of various pieces of physical evidence, like the broken lock, that might reveal to a sufficiently careful observer important facts about the circumstances of Swann's death.

§2.3 The Structure of Interpretation

Next, interpretation is *functionally structured*, in the bare sense that it can be usefully described as taking an input and issuing an output (or, in incidentally apt terminology, taking an argument and issuing a *value*). I'll refer to an interpretation's input/argument as the *object of interpretation*, and to its output/value (if the interpretation goes well) as a *meaning*.[22] (An act of interpretation attributes a meaning to an object.) In BROKEN LOCK, the object of interpretation is the broken

11

window lock, and Hubbard considers both **(2)** and **(3)** as potential meanings. In the case as I described it, his interpretation *does* go well, and he correctly recognizes that **(3)** is a meaning of the broken lock, whereas **(2)** isn't.

§2.4 Validity and Fruitfulness

Many processes are functionally structured in the capacious sense specified above, and we can appreciate the particular character of interpretation by comparing it to some of its neighbors. Consider, for example, inspiration, which can also (at least in many cases) be described as taking an argument and issuing a value.[23] Sherlock Holmes, for example, is often presented as being inspired by some incidental object or event—say, a naïve or bumbling comment of Watson's— with a sudden insight that solves the case. But the content of Holmes's insight, in such cases, is neither the meaning of Watson's utterance, nor what Watson meant to convey by his utterance. What distinguishes cases of inspiration from cases of interpretation?

One good answer is that interpretation is subject to a norm of *validity*, but inspiration isn't. An object of interpretation, *qua* object of interpretation, constrains the range of interpretive values (meanings) in a way that objects of inspiration (i.e., objects that inspire) don't constrain inspirational values: it makes sense to call some instances of interpretation "invalid," whereas calling an instance of inspiration "invalid" is a category mistake. Validity, in turn, is clearly determined by some sort of relationship that holds between an object of interpretation and the meaning attributed to it. Imagine, for example, that Hubbard interpreted the broken window lock to mean that

(4) Wendice was trying to frame his wife for Swann's murder.

(4) is both true and valuable to Hubbard's project of uncovering the details on Swann's death, but it's nonetheless a bad interpretation of the broken window lock, because Wendice destroyed the lock to deceive the police about Swann's method of entry (under the assumption that Swann would kill Margot successfully)—not as part of his plan, subsequently devised, to frame Margot for Swann's murder. We can call the relation that an object of a valid interpretation must bear to its meaning a *licensing relation*. For the moment, I won't specify exactly what sort of relation licenses a given interpretation (canons of validity might even vary by classes of interpreted objects). Noting that some licensing relation is necessary is sufficient for my current purposes.

In addition to being governed by a norm of validity, interpretation is governed by a norm of *fruitfulness*: its output ought to be relevant with respect to the evaluative context within which the interpretation is made.[24] This feature of interpretation is reflected in everyday discourse by the way we call potential objects of interpretation "meaningless" if we think they afford no insight with respect to the implied evaluative context. The broken window lock is a meaningful aspect of the scene of the crime in BROKEN LOCK; the way the chairs are arranged around the dining table, by contrast, isn't. It isn't that the chairs' arrangement carries no interpretively available information; it's that none of that information is helpful (or relevant) with respect to Hubbard's task of uncovering the circumstances of Swann's death.

Here, *inference* is an illustrative contrast case. Like interpretation, inference is subject to a norm of validity, but, unlike interpretation, inferences, *qua* inferences, are equally good regardless

of their fruitfulness. If (to return to the characters of BAD BOYFRIEND) Alice tells Larry to "take a hike," both

   (5) Alice speaks English.

and

   (6) Alice is ending their romantic relationship.

are equally good inferences, but (presuming something like the usual ensemble of interests in exchanges like these) **(6)** is a better interpretation of Alice's utterance than **(5)**.[25, 26] In everyday discourse we use a variety of terms for the quality that makes one interpretation more fruitful than another (relevance, significance, etc.), but I'll generally refer to this quality as *import*.

§2.5 Mental Mediation

Finally, interpretation is *mentally mediated*, rather than mentally immediate. When Hubbard first observes the window lock, he simply *sees* that the lock is broken in a way that he doesn't just *see* that someone has attempted to create the misleading appearance that the window was Swann's method of entry. Items are objects of interpretation rather than, say, objects of perception when they are, in Clifford Geertz's formulation, "on their surface enigmatical": their import is not immediately apparent.[27] NHL defenseman Mike Green's vivid discussion of the intimacy a hockey player needs to have with his stick puts the point into relief: "If I pause to interpret what I'm sensing when the puck is on my stick, that extra split second can be the difference between a shot and a goal, a win or a loss or getting my head taken off. So the stick has to feel like a piece of you."[28] As Green implies, stereotypical interpretation involves a kind of conscious effort that basic

sensory perception, for example, doesn't. Interpreters are aware of and attend to objects of interpretation *as media*, possessed of features independent of the meanings they disclose, in a way that perceivers (we now think) generally do not attend to the "sense data" posited by certain indirect theories of perception.[29]

This point, however, requires elaboration. In the next section of this chapter, I'll turn to the task of showing the way the general account of meaning I've developed here applies to meaning's more specialized senses, with a focus on meaning in the philosophy of language and mind. But the way I've described mental mediation above raises a worry. Subjectively, processing basic linguistic utterances, made in a language in which we're fluent, doesn't seem to require the sort of conscious cognitive effort Green worries will slow him down on the ice. This, I think, is what A.P. Martinich has in mind when, in his widely used philosophy of language reader, he registers some initial surprise that analytic work done on "linguistic understanding…has been presented under the rubric of theories of interpretation or translation"—three phenomena he thinks display significant "intuitive differences." Interpreting an utterance, he thinks, seems to *presuppose* "knowledge of what the speaker clearly means," instead involving a *use* of that knowledge to explain "those aspects of the speaker's meaning that are difficult or dubious" or "what place the speaker's utterances…have within a larger cultural context."[30] If the characterization of interpretation I've given above were to disqualify the process by which a hearer comes to understand a linguistic utterance from counting as an instance of interpretation, then an account of meaning couched in terms of interpretation would simply ignore large swaths of language comprehension.

An elaboration of mental mediation should assuage this worry. Interpretation embraces certain instances of automatic processing in something like the way intentional action embraces certain instances of reflex. Usually, reflexive behaviors are presented in contrast to intentional action, as in a passage from Dewey that I quoted above: "A noise may make me jump without my mind being implicated. When I hear a noise and run and get water and put out a blaze, I respond intelligently; the sound meant fire, and fire meant need of being extinguished…When things have a meaning for us, we *mean* (intend, propose) what we do: when they do not, we act blindly, unconsciously, unintelligently."[31] But many reflexive behaviors—"reflexive" in the sense of automated, "automated" in the sense that they don't involve conscious reflection on the "meanings" (in Dewey's sense) that guide them—are nonetheless intentional and intelligent: consider the way that Roger Federer responds to a 130 mile-per-hour serve. His actions are as automatic (or nearly as automatic) as a patellar reflex, but, unlike a patellar reflex, they're voluntary, sensitive to his conscious goals, and sensitive to background/contextual knowledge (like the particular strengths and weaknesses of his opponent). Information processing can be interpretive, even when it's relatively automated, to the extent that it shares those additional features of reflex actions. Consider again the contrast class of basic visual perception. Knowing observers of the Müller-Lyer illusion may *judge* that the two lines are equidistant while still experiencing the *appearance* that one is longer than the other. Interpretations—even ones made relatively automatically—are revisable in a way that brute appearances sometimes aren't.[32]

In this vein, a second point, mentioned above, bears repeating. Even in relatively automatic cases of interpretation, interpreters are *aware* of the objects of their interpretation *as*

media—objects with properties independent of (and sometimes even in some phenomenological disharmony with) those essential to their relevant meanings. A person who encounters a painting so life-like that she is deceived into believing that she sees, directly, the depicted scene hasn't thereby *interpreted* the painting. As Richard Wollheim has observed, this "twofold" quality of interpretation—the simultaneous awareness of features of medium and meaning—is essential to our ability to appreciate poetry, and the harmonies it effects between the sounds and meanings of words.[33] It also explains how even conventionalized, relatively automatically processed instances of figurative language like metaphor and irony can nonetheless still depend, in Elisabeth Camp's formulation, "on a felt gap between what the speaker says and what she means."[34]

§2.6 The CI Account Summarized

First, a clarification of the preliminary formulation of the CI account I offered in §1—that "meanings are the outputs of acts of interpretation." Though that slogan captures the central idea that meaning, across contexts, should be understood in terms of the role it plays in the practice of interpretation, it also suggests, misleadingly, that an object's meanings *come into being* only when it is actually so interpreted. That suggestion is misleading on two counts. First, as the norm of validity implies, interpreting an object as having a certain meaning doesn't cause it to have that meaning—otherwise, misinterpretation would be impossible. Second, an object has the meanings it does relative to potential rather than actual acts of interpretation. Just as a substance can be nutritious for an organism whether or not it's actually ingested and metabolized, an object can be

meaningful with respect to a potential interpreter in an evaluative context whether or not it is

actually interpreted.[35]

An object's meanings, then, are its distinctively interpretive *affordances*: the ways that it

can be profitably used to guide action, theory, and affection by organisms with the ability to

interpret it.[36] Interpretation, in turn, is a functionally structured, mentally mediated activity,

undertaken by an organism in an evaluative context, subject to norms of validity and fruitfulness.

The former of these norms is met when a posited meaning is related, in the right way, to the

object of interpretation; the latter is met to the extent that the posited meaning has import with

respect to the evaluative context in which the interpretation is made. "Meaning" is unified across

its broad range of senses according to its structural role in the practice of interpretation.

§3 Demonstrating Comprehensiveness

I hope that my efforts above have been sufficient to convince you that the comprehensive-

interpretive account of meaning deserves to be called "interpretive"; now I'd like to convince you

that it deserves to be called "comprehensive," as well. I'd like to show not only that CI account

embraces meaning's important varieties—natural, linguistic, literary, cultural, ethical—but also

that it provides a rubric that helps us think through some of the important philosophical

questions that these varieties raise. In other words, I want to vindicate a claim I made at the outset

(§1.2), that appreciating the character of meaning's integrity generates valuable insights into its

particular varieties.

I'm not going to distribute my attention evenly, though. Instead, I'm going to focus

primarily on the question of how the general account of meaning I've provided bears on the

discussion of meaning in the philosophy of language and mind. Initially I'd planned otherwise—

to show, in detail, the way the CI account applies to all the varieties of meaning I listed in the last

paragraph—but equal treatment would have required something more like a dissertation than a

dissertation chapter. Given the need to focus, I have two reasons for dedicating the bulk of my

attention to the issue of linguistic meaning. First, restrictive impulses seem strongest—mostly, in

fact, just taken for granted—in the case of linguistic meaning. Cultural and literary theorists, for

example, are more apt to suppose that their theoretical target is "meaning" in some more

expansive sense, while philosophers of language seem generally less inclined to suppose that

considerations about other varieties of meaning bear significantly on their efforts, or that their

conclusions about meaning in mind and language should be thought to apply to meaning in its

other varieties (if there even are other "varieties," speaking precisely). Nor do I think this

asymmetry is merely a disciplinary idiosyncrasy: there are substantive reasons that make language

a special case. Therefore, addressing "meaning" in the context of language and mind is the most

challenging test case for an account that aspires to be comprehensive. My second reason is

simpler: a subsequent chapter of this dissertation focuses on irony, and irony, like meaning, has

often been treated as a primarily linguistic phenomenon.

Still, my attention will be disproportionate, not exclusive: I'll have something to say about

other varieties of meaning. In particular, I'll begin with a moderately involved treatment of

natural meaning, in part because that notion plays an important role in delimiting the domain of

the kind of meaning that, according to one major tradition, should primarily concern philosophers of language. And, in the process of addressing some questions about the nature of linguistic meaning, I'll have occasion to note some interesting implications of the CI account for "meaning" as it's used in other contexts.

§3.1 Varieties of Meaning

To this point I've been relying mostly on informal, commonsense intuitions to support the idea that "meaning" has at least superficially disparate senses, but the conceptual apparatus associated with the CI account provides the resources to differentiate between meaning's varieties more systematically. In fact, the CI account provides several dimensions according to which meaning might be classified into kinds, but the most obvious of these is by object-type. The suggestion is that "natural" meanings arise when we interpret distinctively natural objects, "linguistic" meaning arises when we interpret distinctively linguistic objects, "literary" meaning involves the interpretation of works of literature as such, and so on.

This suggestion, however, raises questions. How should the relevant object types—natural, linguistic, literary, and so on—be individuated? And, when we sort meaning into varieties using this strategy, how well do the resulting categories actually correspond with intuitive or theoretically significant kinds of meaning? For example, consider the way that the suggested object-type classification scheme intersects with Grice's seminal classification of meaning, into so-called "natural" and "nonnatural" kinds. Many objects that aren't stereotypically "natural," like quarterly business reports, can have "natural" meaning, in Grice's sense.[37] Moreover, Grice

argued that *beginning* an investigation into "the kind of meaning language has" by looking for something *distinctively* linguistic had led to significant errors, like a too-close identification between the meaning of a linguistic utterance and its conventionally-coded content. Instead, Grice argued, linguistic utterances are just one of many ways to convey nonnatural meaning.[38]

My task in §3—to show that the CI account of meaning does in fact comprehend meaning's important varieties—would be challenging but straightforward if the important varieties were already well and uncontroversially defined, so that all that was necessary was to show how this or that particular definition of "meaning" falls under the more general account given in §2. The questions I just raised indicate, however, that the task is more involved. So does an observation I made in §1—that even within specialized discourses, like the study of language, "meaning" often functions ambiguously. For instance, in *Interpreting Figurative Meaning*, an important critical survey of the empirical study of figurative language, Raymond Gibbs and Herbert Colston claim that "understanding what any figurative utterance means is not simply a matter of getting to a particular figurative meaning, but understanding what a speaker pragmatically intends to achieve by use of that trope."[39] Readers might justifiably wonder what "a particular figurative meaning" could be, if it's not what a particular "figurative utterance means": clearly the two can't be the same if understanding the former isn't sufficient for understanding the latter.[40] A significant part of the task of demonstrating that the CI account is genuinely comprehensive will be to achieve a clearer understanding of the varieties it aspires to comprehend.

§3.2 Natural Meaning

In §1 we noticed that Dewey, our representative ecumenicalist, wrote without qualm about the "meanings" of nonsymbolic objects (like a loose rock lying on a walkway) and substances (like water). If any sense can be made of that way of talking, it likely involves an appeal to natural meaning—a notion that Grice influentially identified with the sense of "meaning" employed in sentences like

> **(7)** Those tracks in the snow mean that M3 [a wolverine living in Glacier National Park] recently passed this way.

and

> **(8)** Dark clouds mean rain.[41]

As we've already noticed, natural meaning (meaning$_N$) is one half of the distinction at the foundation of Grice's theory of meaning; nonnatural meaning (meaning$_{NN}$) is the other. As Grice draws the distinction, natural meaning is *factive*: if x means$_N$ that p, x entails p. (If it was M3's nemesis M6 who left the tracks in the fresh snow, then, no matter how misleading the appearance, they didn't mean$_N$ that M3 had been on the spot within the past six hours.) Nonnatural meaning, by contrast, is *nonfactive*: x may mean$_{NN}$ that p without entailing that p. (So, when uttered, the sentence "*American Beauty* is a wonderful movie" generally means$_{NN}$ that *American Beauty* is a wonderful movie, despite the fact that it also means$_N$ that the person who utters it has questionable taste.)[42]

Natural meaning is an important notion for both restrictive and ecumenical approaches to meaning. On the restrictive side, Grice's way of framing this distinction has provided one

extremely influential method for isolating the proper target of inquiry. The task for philosophers of language who hope to provide a "theory of meaning," according to this way of thinking, is to provide an analysis, or grounding account, of nonnatural meaning—the distinctive (and distinctively mysterious) way that speakers use language to mean. On the (relatively) ecumenical side, Grice's framework has seemed to provide a promising starting point for theorists interested in emphasizing the ways that the human ability to discern and communicate meaning (particularly in language) is *continuous* with (or, at least, emerged from) the way structures bear, and organisms recover, meaning in the rest of the natural world. Grice himself offered a rough outline of how he thought such a project might proceed in "Meaning Revisited," and philosophers like Dorit Bar-On[43] and Ruth Millikan[44] (among many others) have made more thorough efforts in that direction since.

An initial challenge that these more ecumenical projects face is that Grice, perhaps because his primary interest was in nonnatural meaning, never offered a full-fledged, substantive account of natural meaning. Instead, he contented himself with providing "recognition tests" for distinguishing cases of meaning$_N$ from meaning$_{NN}$ (like the factivity distinction above), along with occasional suggestive comments: meaning$_N$, he says, is "closely related to the idea of [being] a natural sign," and "means$_N$" is perhaps interchangeable with "has the consequence," or "evidences."[45] But, as their vagueness suggests ("natural sign" is left undefined; the difference between "has the consequence" and "evidences" can, in relevant contexts, be momentous), these proposals are closer to marking the starting point of an inquiry than they are to expressing its deliverances.

Nor is natural meaning's importance limited to its role in distinguishing (as a contrast class) and explaining (as an evolutionary development) nonnatural meaning. If, as Grice plausibly suggested, natural meaning is intimately associated with natural signs, then it is a central component of everyday activities like hunting and predicting the weather, as well as larger-scale theoretical enterprises like historical natural science, which Carol Cleland has described as the testing of "hypotheses concerning long-past, token events" according to "their capacities to explain puzzling associations among traces discovered through fieldwork."[46] For example, interpreting the theoretical import of the composition of the clay at the Cretaceous-Paleogene (K-Pg) boundary of the geological record—unusually rich with iridium, shocked minerals, and Ni-rich spinels—is the basis of our knowledge that a meteor strike, rather than disease or volcanism, was the primary cause of the extinction of the dinosaurs.[47]

The CI approach to meaning offers a valuable guide for developing a substantive, independently attractive account of natural meaning. To show this, I'll introduce two theoretical frameworks with some promise for assimilating Grice's notion of natural meaning—Fred Dretske's semantic theory of information, and James J. Gibson's theory of ecological affordances—and evaluate each strategy according to the implications of the CI account.

§3.2.1 Natural Meaning as Semantic Information

One popular philosophical exposition of the relation between trace and event that underwrites the sort of inference discussed above—one that is, in fact, frequently associated with Grice's notion of natural meaning—is Fred Dretske's semantic theory of information.[48] According to this strategy

of assimilation, the "natural meaning" of an object, event, or state of affairs is just the *information* it carries—where information, Dretske says, is "an objective commodity…whose generation, transmission, and reception do not require or in any way presuppose interpretive processes."[49] The chief innovation of Dretske's theory was to combine this robust objectivity (interpretive- and mind-independence) with certain semantic features: to extrapolate from Claude Shannon's mathematical theory of information, which describes how information can be quantified, a method for specifying information's quality—that is, its content.

The strategy begins by explaining how the quantity of information at a receiver (r) (a measure of the extent to which the actual condition of r has reduced the possible conditions of r, abbreviated "$I(r)$" in Dretske's notation) can be *about* a source (s). The quantity of information at r that is *about s* ("$I_S(r)$") is a measure of the extent to which the information generated at r depends on the information generated at s (the "degree of lawful (nomic) dependence between these two points"). Specifying the content quantified by $I_S(r)$ is then just a matter of offering a description, in propositional form, of the condition at s upon which $I_S(r)$ (lawfully) depends.[50] This content, under the assimilation we're currently entertaining, is r's natural meaning with respect to s: what r means$_N$ about s is just the content quantified by $I_S(r)$. Roughly (but not misleadingly): a tree's trunk having n rings means$_N$ that the tree has lived n years because former depends causally (by processes governed by natural law) on the latter; the abnormal abundance of iridium in the layer of sediment at the K-Pg boundary means$_N$ that a large meteor struck Earth at the time the sediment was deposited for the same reason.

§3.2.2 Natural Meaning as Ecological Affordance

An alternative strategy for understanding natural meaning is to integrate it into the idiom of

ecological psychology, a tradition pioneered by James J. Gibson. In *The Ecological Approach to*

Visual Perception, Gibson argues that physics (among other disciplines) is insufficient to the task

of explaining visual perception: its concepts work at the wrong "level of description" for the

undertaking.[51] Instead, he advocates an "ecological approach"—one constitutively attuned to the

behaviors of organisms in their environment. The essence of the distinction, as Gibson makes it,

is that "[t]he world of physical reality does not consist of meaningful things," whereas "[t]he

world of ecological reality does. If what we perceived were the entities of physics and

mathematics," he continues, "meanings would have to be imposed on them. But if what we

perceive are the entities of environmental science, their meanings can be *discovered*."[52]

The constituents of ecological reality, in other words, are distinguished by (organized

according to) their implications for the natural activities of the organisms in the environment

they compose: their "affordances," as Gibson calls them. This is just a longer restatement of the

claim that ecological reality consists of "meaningful things." For example, "Water," Gibson

explains,

> affords bathing and washing, to elephants as well as to humans. Streams of water can be
> dammed, by beavers as well as by children and hydraulic engineers. Ditches can be dug
> and aqueducts built. Pots can be made to contain water, and then it affords pouring and
> spilling. Water, in short, has many kinds of meaning.[53]

After providing a similar accounting of fire's affordances, Gibson notes that it, "like water, has

many kinds of meaning, many uses, many values."[54] The implication of assimilating an object's

natural meaning to its ecological affordances manifests in this interchangeable use of terms—

"meaning…uses…values": an object's affordances, and hence, its meanings "are what it *offers* the

animal, what it *provides* or *furnishes*, either for good or ill."[55]

§3.2.3 The CI Approach to Natural Meaning: Critical Evaluation I

Suppose that the CI account competently characterizes meaning in general, and that natural

meaning is rightly understood as a variety of that general phenomenon. What should we conclude

about the two assimilation strategies I've just described? For reasons I'll detail below, neither

proposal—that natural meaning be assimilated to information, in Dretske's sense, or to

affordances, in Gibson's—is fully consistent with the CI approach to meaning. Ultimately,

though, I think this is a virtue of the CI approach, rather than a defect. The revisions to each

proposal required to make natural meaning consistent with meaning in general yield a clarified

rather than disfigured understanding of natural meaning, and suggest more coherent ways of

integrating natural meaning into the Dretskean and Gibsonian paradigms.

The disharmonies between understanding natural meaning as a variety of meaning on the

CI model, on one hand, and as identical with Dretske's semantic information, on the other, were

probably readily apparent in my initial presentation of Dretske's theory. Information, Dretske

held, is "an objective commodity…whose generation, transmission, and reception do not require

or in any way presuppose interpretive processes."[56] Meaning, I've argued, should be understood

according to the structural role it plays in the practice of interpretation. This fundamental

difference manifests in at least two important discontinuities between Dretske-information and

CI-meaning. First, as we saw in §2.4, *meanings* have import with respect to a broader practical, theoretical, or affective context. In BROKEN WINDOW, not all of the objects in the crime scene were equally meaningful, and this discrepancy was based not merely in the *quantity*, but on the *quality* of the information they carried—its relevance to the larger task in whose context interpretation was being undertaken. Second, as I argued in §2.5, the use of "meaning" implies something about the way import comes to be (or could come to be) appreciated: by interpretive effort.[57] The question of what information a signal carries about its source is sensitive to neither of those considerations: "The receiver of the signal may be more *interested* in one piece of information than he is in any other, he may succeed in *extracting* one piece of information without another, but these differences are irrelevant to the information the signal contains."[58] So, the CI approach suggests we should resist identifying natural meaning with Dretske-information. Is the suggestion good?

It is. First, it strains intuition to suppose that a signal means *p* (in any sense) if *p* isn't even *in principle* accessible to a potential interpreter on the signal's basis. Yet Dretske underscores the point that not even in-principle interpretive accessibility is required for a signal to carry information about a source.[59] (The same point applies to our intuitive resistance to calling useless information meaningful.) More importantly, distinguishing natural meaning from information in the ways the CI approach prescribes better respects the integrity of Dretske's own theory. Dretske repeatedly emphasized the importance of "distinguish[ing] information sharply from the concept of *meaning*—at least the concept of meaning relevant to semantic studies of language and belief."[60] The discussion above shows not only that he was right, but that he needn't have

qualified. Information's status as an "objective commodity"—one that, perhaps, can provide at least a partial basis for a non-circular explanation of various mental and semantic phenomena—can be preserved *alongside* the traditional association between meaning (even "natural" meanings) and relevance.

Moreover, although the CI approach discourages *identifying* natural meaning with information (an identification Dretske didn't himself propose), it does suggest a way of integrating the two notions. We might, that is, regard Dretske's analysis of carrying information as an account of the *licensing relationship* required for valid interpretations of natural meaning—so that interpreting a signal/object x as meaning$_N$ that p is valid iff x carries the information that p. More fundamentally, distinguishing natural meaning from information makes it possible to ask whether "carries information," in Dretske's sense, *is* the proper account of the licensing relationship for natural interpretation. According to Dretske's theory, a signal carries the information that a source is in a certain state only if the probability of the source's being in that state, conditional on the signal, is 1. Though Dretske has good reasons for imposing this requirement on information transmission, it might be too stringent as a condition for the validity of natural interpretation. Suppose, for example, that paleontologists and geologists have rightly concluded, on the basis of their careful interpretation of the composition of the sediment at the K-Pg boundary, that a meteor strike killed the dinosaurs—but that the conditional probability of a meteor strike given the sediment's composition, even accounting for additional evidence possessed by the scientists, although very high, is less than 1.[61] We might want to regard the scientists' interpretation as valid nonetheless. Similar considerations have led theorists who try to

provide an account of the naturalistic emergence of intentional mental states in terms of "natural signs" to object to using Dretske's stringent definition of information as their basis. For example, Ruth Millikan asks "what use signals carrying bits of natural information [in Dretske's sense] could be to an organism," noting that "the mere fact that a signal carries certain natural information seems not to bear on whether a creature could learn anything from encountering that signal or know anything by virtue of harboring that signal in its brain."[62] Instead, she argues, a theory that "will help to explain how real animals manage to obtain *useful* information will need to introduce statistical considerations about the environment in some controlled way…Nearly all the kinds of information needed by us, and by all other organisms as well, for securing what we need in an inclement world, is information that cannot possibly be acquired without leaning on certain merely statistical frequencies."[63] Distinguishing the natural meaning "closely associated with the idea of being a natural sign" (as Grice characterized it) from information makes it possible to ask important questions about how these notions relate, rather than discarding either because it doesn't meet the requirements of the other.

§3.2.4 The CI Approach to Natural Meaning: Critical Evaluation II

The considerations introduced above make an assimilation of natural meaning to Gibsonian affordances look more promising. Import (value or relevance for an organism in an environment) and accessibility are, after all, central to the idea of affordances, which, in Gibson's words, "point both ways, to the environment and to the observer."[64] An object that carries information either wholly inaccessible to an organism, or wholly irrelevant to any of its interests, affords nothing (so

30

far as it's been described). But natural meaning (understood as consistent with the CI account) resists assimilation to ecological affordances for a different reason.

Put simply, affordances arise with respect to a broader range of activities than (CI) meanings do. For Gibson, the environment is "meaningful" with respect to the full range of interests and behaviors of the animals who inhabit it: air affords respiration for land animals, water affords locomotion for fish, cragged cliffs afford refuge for a mountain goat, but a frustrating obstacle for a predator, and so on. Furthermore, Gibson argues, most of these affordances can be perceived directly, as his Gestalt forebears suspected. The Gestalts rejected prevailing theories of perception according to which "*no* experiences were direct except sensations…Bare sensations had to be clothed with meaning," recognizing, instead, that "the meaning or the value of a thing seems to be perceived just as immediately as its color."[65] This, Gibson argued, is because phenomenal objects (or, as he might prefer, ecological objects) are individuated by their affordances—

> what we perceive when we look at objects are their affordances, not their qualities...Phenomenal objects are *not* built up from qualities; it is the other way around. The affordance of an object is what the infant begins by noticing. The meaning is observed before the substance and surface, the color and form, are seen as such.

Hence, "the theory of affordances is a radical departure from existing theories of value and meaning." [66]

The CI account advises us to depart less radically. Rather than following Gibson in identifying the meanings of objects with their implications for the interested behaviors of an animal in an environment *simpliciter*, the CI account restricts meaning to affordances *only* with

respect to the practice of interpretation. Interpretation, meanwhile, is always undertaken in some

broader practical, theoretical, or affective context: interpretation *guides* intelligent action, belief

formation, and affection. Instead of being *in* the composition and layout of surfaces, as directly

perceivable Gibsonian affordances are, *meanings* are about something *other* than the surface

qualities of the objects that bear them; instead of being directly enjoyed or suffered, meanings are

implicated in larger contexts of belief, action, or feeling. To the extent that an object is cognitively

identified with or individuated by some immediate import, the need for any appeal to its

"meaning" with respect to that import is obviated. (Judgments of an object's meaning are

revisable under critical reflection without ceasing to be judgments of (a) *that* object's (b)

meaning.) Water affords locomotion to a fish without recourse to any interpretive process, and

there's little to be gained by conflating meaning with this very general sense of the way features of

an environment can facilitate or hinder an organism's interested behaviors. "Affordance" already

captures that important idea well.

Parallel to our discussion of Dretske, then, a simple assimilation of affordances to natural

meanings ultimately does Gibson's own theory no favors. But, again, there's a natural way to

integrate an understanding of meaning consistent with the CI account into a broadly Gibsonian

idiom: by regarding meanings as distinctively *interpretive* affordances. The suggestion is

intermediate between Gibson's style of treating meaning, value, and affordance as roughly

interchangeable terms, and a finer distinction he makes *within* the category of affordances, when

he picks out "a very special class of artificial objects" called "human displays"—"surface[s] that

ha[ve] been shaped or processed so as to exhibit information for more than just a surface itself."[67]

Gibson wrote as if the dichotomy between import trafficked in intentional communication and import perceived directly was exhaustive:

> The concept of information with which we are most familiar is derived from our experiences of communicating with other people and being communicated with, not from our experience of perceiving the environment directly. We tend to think of information primarily as being sent and received, and we assume that some intermediate kind of transmission has to occur, a "medium" of communication or a "channel" along which the information is said to flow… like the spoken and written words of language, *man-made*. The ambient stimulus information available in the sea of energy around us is quite different. The information for perception is not transmitted, does not consist of signals, and does not entail a sender and a receiver. The environment does not communicate with the observers who inhabit it. Why should the world speak to us?[68]

But as Dretske hoped to explain, even objects that haven't been intentionally processed for the purpose (and are hence "natural" in the relevant sense) may "exhibit information for more than just [their] surface[s]." When this information is relevant and accessible by means of interpretation to an organism engaged in a practical, theoretical, or affective project—when a surface has, to modify Gibson's phrase, *interpretive* affordances—we have the basis for an understanding of natural meaning consistent with the CI account.

Ultimately, then, natural meaning, understood according to the precepts of the CI account, resists strict assimilation to Dretske's notion of information and to Gibson's notion of ecological affordances. Instead, the CI account recommends an intermediate way of understanding natural meaning that draws powerful insights from both traditions, without disfiguring either. Unlike information, natural meaning isn't insensitive to the environmental conditions, capacities, and interests of its interpreters; neither, however, is it implicated merely in the immediate doings, sufferings, or perceivings that environmental objects afford. Instead,

33

natural meaning is grounded in the informational relationships that natural objects, states, and events share with one another, whose basis Dretske strove to explain. Such relationships, however, precisely, they are understood, can generate *interpretive affordances*—meanings—without relying on the sort of intentional productions that give life to manufactured symbols, like strings of language. This is why, even if nature doesn't speak to us, we can sometimes listen in.

§3.3 "Linguistic" Meaning

In Ursula Le Guin's *Earthsea* stories, wizards work magic by their knowledge of the "Old Speech," also called the "True Speech"—the language of dragons. The words of the Old Speech give their speakers power because they're the words of original creation: they compose "the language Segoy spoke who made the islands."[69] Everything in Earthsea has one, and only one, true name in the Old Speech: the name by which it was made to be. The wizards of Earthsea spend their lives learning these names, and the methods by which they can be discovered—many for the power it gives them, some few simply to understand and to love the world in which they live. No wizard can use the Old Speech to lie. In Earthsea, like John's Gospel, in the beginning was the word.

If we lived in Earthsea, I wouldn't need to include a separate section on linguistic meaning—or, at least, the reason I'd have for doing so would be very different. The languages we know aren't like Le Guin's Language of the Making. Instead, we live in the world Galileo described, in which "names and attributes must be accommodated to the essence of things, and not the essence to the names, since things come first and names afterwards";[70] the world in which, as Juliet protested, "a rose by any other name would smell as sweet."[71] Like any other meaningful

object, linguistic utterances have interpretive affordances: under the right circumstances, they can tell us that M3 recently passed this way just as well as his fresh footprints in the snow, and they can inform us of a tree's age just as well as the rings in its trunk. But footprints and tree rings relate to their meanings in a way similar to the way that the objects of Earthsea relate to their names: tracks of any other shape would *not* be M3's. The relationship between units of language and the meanings they can be used to communicate seems more arbitrary, more contingent. Granted that the noises others make, and the shapes they inscribe, can have meaning, just as the CI account characterizes it; one difficulty is, *how*? As Paul Horwich puts the question, "How is it possible for…intrinsically inert ink-marks… to reach out into the world and latch on to a definite portion of reality?"[72]

Unfortunately, Horwich's question isn't the only one that philosophers of language and mind have supposed a "theory of meaning" (in these fields, usually meant restrictively) ought to answer. We've already seen that Davidson wanted to know how a system composed of finite elements, employed by agents of limited cognitive powers, could, in principle, be used to generate an infinite (or, at any rate, uncountably large) number of meaningful sentences. And Grice drew our attention to the peculiarly non-factive way that linguistic utterances can "mean": the sense in which the sentence

(9) "Cassio was muttering about Desdemona in his sleep"

uttered by Iago to Othello, means that Cassio was muttering about Desdemona in his sleep—even if, in fact, Cassio never said anything in his sleep. More basically, there's something extraordinary about the way that language apparently enables us to produce objects with relatively precise

interpretive affordances for one another at will. Nature may not speak to us, but how is it that we speak to one another?

This variety of questions (and the ones I mentioned above are just the beginning) helps explain my reason for using scare-quotes in this section's title ("'Linguistic' Meaning"). It's not that I find the notion suspicious; it's just that it's difficult to pin down. The object-type classificatory strategy I proposed in §3.1 makes the issue seem straightforward: "linguistic" meaning arises when we interpret distinctively linguistic objects. But as I noted then, that strategy is both provisional and incomplete. Provisional, because the strategy is only useful for current purposes if it tracks, at least roughly, the theoretically interesting kind in question; incomplete, because we still require a specification of what counts as a distinctively "linguistic" object.

Neither issue admits of a simple, uncontroversial resolution. It's difficult to delimit the domain of the "linguistic" without begging many of the philosophical questions that a theory of meaning ought to adjudicate; moreover, as Elisabeth Camp has observed, even philosophers whose trademark claims rely on the distinction between language and nonlanguage haven't always provided a sufficiently thorough accounting of the distinction's supposed grounds.[73] Further, having arrived at some specification of the distinctively linguistic, we may discover that the features of linguistic meaning centrally in need of philosophical explanation aren't uniquely possessed by linguistic objects after all, in which case "linguistic meaning" will have turned out to be an imperfect name for the philosophically relevant variety of meaning in question.

David Lewis offers a sensible suggestion about how to proceed: "In order to say what a meaning *is*, we may first ask what a meaning *does*, and then find something that does that."[74] The

problem is that meaning is supposed to do so much: to explain how language facilitates communication, to explain how inert marks or even brain states "reach out into the world and latch on to…reality," to underwrite notions of translation, synonymy, analytic truth—and so on. Obviously, the theory of meaning in philosophy of language and mind is implicated in so broad an array of issues that an attempt to treat them all in a single section of a single dissertation chapter would be quixotic if it weren't so certain to be bereft of romance and poetry.

Fortunately for us both, the purpose of this section isn't to resolve (or even to address) all of the philosophical mysteries that the kind(s) of meaning associated with language generate. Instead, I hope only to show, in broad strokes, the way a notion of linguistic meaning can be integrated into the CI account, and, second, that this integration provides a useful rubric for ordering our thinking about some of the important issues at hand.

§3.3.1 The Metaphysics of Meaning

One question that has been relatively central to the theory of meaning in the last century or so is whether it's appropriate to think of the meanings of linguistic objects as a kind of entity, and, if so, to say what kind of entity they are. Referential theories of meaning are among the simplest and most resilient sort of entity theory, with the basic, intuitive idea being that the meanings of words are the things in the world they refer to, or stand for.[75] The British Empiricists, by contrast, thought that the meanings of words were mental entities. As Hobbes put the point, "Words so connected as that they become signs of our thoughts, are called SPEECH, of which every part is a name";[76] as Locke had it, "The use, then, of words, is to be sensible marks of ideas; and the ideas

they stand for are their proper and immediate signification."[77] Frege advanced influential

arguments against the sufficiency of either position: meanings couldn't be ideas because ideas are

subjective, whereas meanings are publicly shared ("capable of being the common property of

many people"); meanings couldn't simply be the objects in the world words refer to because

(among other reasons) this trivializes identity claims (like "water is H_2O") that state empirical

discoveries rather than analytic truths. Instead, Frege taught philosophers to identify the

meanings of words with their "senses"—abstract but objective "modes of presentation."[78] Finally

there are the philosophers who find the effort to identify meanings with entities altogether

misguided. For Quine, Fregean senses were "creatures of darkness"—ontological extravagancies

without standing from a physicalist perspective, and therefore best exorcised.[79] Davidson objected

to positing entities in a theory of meaning not because such entities "are abstract or that their

identity conditions are obscure, but because they have no demonstrated use."[80] Does the CI

account shed light on these issues?

Not as directly as it might initially seem. In the introduction to §2, I associated the CI

account with the tradition that advocates understanding meaning in terms of "activity rather than

ontology"—an affiliation that seems to foreshadow an allegiance with the anti-entity position. I

went on to develop an account of meaning, grounded in the structural role it plays in the practice

of interpretation, according to which meanings are *interpretive affordances*—relational

properties—of objects (grammatically understood). Isn't that, then, what the CI account has to

say about the metaphysics of meaning? Whether or not these relational properties get counted as

"entities" of some kind will depend on the precise way that term ("entity") is construed in one's

metaphysics, along with the way that the favored metaphysical approach regards the various possible grammatical "objects" (events, experiences, states-of-affairs, relationships, artworks, etc.) of interpretation. But, regardless of the way those details are adjudicated, the basic contribution of the CI account seems clear: meanings are relational properties of objects, as detailed in §2.

The problem with this sort of *direct* application of the CI account to the issue in question is that it ignores the point I made in §3.3: that, at least in one of its important phases, the search for a theory of meaning in the philosophy of language *assumes* that linguistic utterances frequently have the sort of interpretive affordances the CI account associates with meaning in general, and asks, instead, *how* this can come to be, what contribution *language* makes to investing otherwise apparently arbitrary or inert sounds, gestures, or inscriptions with the rich interpretive affordances they apparently have. And it's with respect to *this* question that the positions I canvassed above ought, at least in large part, to be understood. Associating meaning-entities with words, for example, is usually part of an effort to explain how words can be combined to compose sentences that have the interpretive affordances that they do.

There are two things worth noticing about this observation. First, it highlights the danger of a thoughtlessly direct application of the CI account to every question associated with the theory of linguistic meaning (as well as the correlative point that restrictivism is overstated rather than unmotivated). Second, it's a starting point for showing the value of understanding these more particular questions *in light of* a comprehensive understanding of meaning in general. By way of illustrating the first of those two points, consider two strings of inscriptions William Lycan introduces at the beginning of his introductory text on philosophy of language:

(10)   w gfjsdhkj jiobfglglf ud

(11)   Good of off primly the a the the why.[81]

Lycan calls both **(10)** and **(11)** meaningless, and an application of the CI account seems to be able

to account for that judgment easily enough: it's hard to imagine a context in which encountering

an utterance of either would afford much guidance for affect, action, or theory.[82] But how would a

direct application of the CI account make sense of the intuitively clear distinction Lycan draws

between **(10)** and **(11)**—that the latter contains "individually meaningful words," whereas the

former doesn't? If an instance of **(11)** affords little or no import to an interpreter in a given

context, and is thereby rightly called "meaningless," on what basis do we rightly call the words it

contains "meaningful"? If the whole has no import, then how could its parts? The point isn't that

there's not a good (and relatively obvious) answer to this question; it's that a flat-footedly direct

application of the CI account to the question of linguistic meaning makes it seem as if there isn't.

Arriving at the obvious answer involves returning to the suggestion I made in §3.1, and

raised some difficulties with in §3.3: to regard linguistic meaning as arising from distinctively

linguistic objects. We can still make sense, in CI terms, of the kind of "meaning" the word-

forming collections of letters in **(11)** have but the gibberish collections in **(10)** don't: the

collections in **(11)**, regarded as types, have standard uses with respect to the general task of

producing and interpreting English sentences; English sentences, in turn, have affordances for the

task of interpreting utterances made by means of those sentences on particular occasions. If we

begin with the (relatively) uncontroversial idea that language is a tool or set of tools that allows

for the intentional production of objects with certain interpretive affordances (like one person

uttering to another "M3 recently passed this way" in a context in which that assertion would typically be made), the CI account identifies an object's distinctively *linguistic* meaning with the general role it plays in facilitating that process. In this sense, the mistake of treating linguistic meaning as CI meaning *simpliciter* is analogous to the mistake I attributed to Gibson, above, when his overly broad use of "meaning" conflated interpretive affordances with affordances *simpliciter*. Complementarily, we should recognize a potential inadequacy of the suggestion, as formulated, that linguistic meaning arises from (should be identified with) the interpretation of distinctively linguistic objects. If I correctly infer a speaker's nation of origin on the basis of the accent she has when she says

(12)   "M3 has just been by here"

it's implausible to deny that I have thereby interpreted a linguistic object: in fact, my interpretation seems overwhelmingly likely if not necessarily to have involved a recognition that (12) was a deployment of language. But "the speaker is from Bulgaria" is certainly no part of the *linguistic* meaning of (12) in the sense that interests philosophers of language—instead, what we want to isolate is the contribution language, understood as a generally available communicative tool, makes to (12)'s possessing the interpretive affordance that motivated the speaker to produce it.[83]

Subtracting the talk explicitly related to the way these issues might be integrated into the CI perspective, most or all of what I've said in the last few paragraphs will register for philosophers of language as, first, familiar, and, second, the entry-points (rather than resolutions) of longstanding controversies. This reflects the fact that my effort thus far has been to show that

41

these issues can be understood as falling within the scope of a general account of meaning along the lines of the one I offered in §2. I think this process, however, has already provided some evidence that strongly restrictive approaches to meaning are inadequate. Take two of the familiarities I canvassed above. The idea that the "meanings" of words (if there are such things) should be understood according to the way they contribute to the meanings of the sentences in which they appear traces back, at least, to Frege's belief in the primacy of the sentence (also sometimes called his "context principle"):

> That we can form no idea of its content is therefore no reason for denying all meaning to a word, or for excluding it from our vocabulary. We are indeed only imposed on by the opposite view because we will, when asking for the meaning of a word, consider it in isolation, which leads us to accept an idea as a meaning. Accordingly, any word for which we can find no corresponding mental picture appears to have no content. But we ought always to keep before our eyes a complete sentence. Only in a sentence have the words really a meaning…It is enough if the sentence taken as a whole has a sense; it is this that confers on its parts also their content.[84]

But why (or in what sense) do sentences deserve to be regarded as primary, if we (standardly) apprehend their meanings *on the basis* of our knowledge of the meanings of their parts? The answer is that only sentences, uttered in a context, have interpretive affordances in the broader sense discussed in §2: interpreters can use them to guide belief, action, or affect in an evaluative context, often in ways equivalent to the way they can use natural signs, like footprints in the snow. (In many evaluative contexts, the import of encountering footprints in the snow or an utterance of **(12)** will be roughly the same.) Words, by contrast, don't generally have interpretive affordances in concrete situations unless they form part of a larger sentence: the interpretive value

they provide is parasitic on the interpretive value of the sentences in which they appear—and this latter kind of interpretive value isn't possessed exclusively by linguistic items.

A second familiarity I introduced above, by way of an attempt to isolate a distinctively *linguistic* sense of "meaning", was related to Grice's warning against conflating the "total signification" of an utterance with its conventional meaning: roughly, the distinction between an utterance's full suite of interpretive affordances, and the subset of that full suite made directly available by an interpreter's pre-existing knowledge of the language in which the utterance is made.[85] So each of these important observations already relies for its justification (Frege) or articulation (Grice) on a concept of meaning not exclusively possessed by linguistic objects.

In §3.3, I claimed not only that linguistic meaning could be integrated into a more comprehensive approach, but also that taking a more comprehensive approach could enlighten us with respect to controversies about meaning that frequently arise *within* the philosophy of language and mind. I think the paragraph above does some work toward keeping that promise, by illustrating the fact that at least some of the central (if not wholly uncontroversial) insights of philosophy of language already involve a more ecumenical stance than statements like "only linguistic forms have meanings" allow. (And this observation, as I suggested in §3.2, is only amplified when we consider the explosion of recent work concerned with the evolution of language and the intricacies of animal cognition.) Nonetheless, to conclude, I'd like to return to aspects of the substantive questions I raised at the beginning of this section, considering not just how these issues can be understood in light of a more general approach to meaning, but also how

taking that more ecumenical perspective can generate worthwhile insight. The related questions

I'll address are: (1) are meanings "in the head"? and (2) is meaning inherently social?

§3.3.2 Are Meanings "in the Head"?

As we saw in §3.3.1, Hobbes had a simple, intuitive account of linguistic meaning, according to

which names are "signs of our conceptions…not signs of the things themselves." In fact, Hobbes

thought, communication is only a secondary function of names; they are first and primarily aids

for individual thought: "So that the nature of a name consists principally in this, that it is a mark

taken for memory's sake; but it serves also by accident to signify and make known to others what

we remember ourselves."[86] A significant portion of 20[th] century philosophy of language (extended

back, for convenience's sake, to include Frege's work in the late 19[th] century) was dedicated to

refuting these claims, and the widespread dissatisfaction with the first (that names are signs of

conceptions rather than things) is captured well by Hilary's Putnam's famous slogan, "meanings

just ain't in the *head*!"[87]

Putnam's central reasons for protesting as he did are reflected in the CI account of

meaning. He formulates his protest as a rejection of three claims, held jointly:

> 1. Every word he uses is associated in the mind of the speaker with a certain mental
> representation.
> 2. Two words are synonymous (have the same meaning) just in case they are associated
> with the *same* mental representation by the speakers who use those words.
> 3. The mental representation determines what the word refers to, if anything.[88]

Mental representations alone can't fix reference, Putnam points out, because both environmental

and social factors contribute to fixing a term's reference, as well. For example, Putnam purports to

be personally unable to distinguish between elm and beech trees, yet he can use the terms "elm" and "beech" in discourse to refer to the two different types, because he can "always rely on experts" to do the necessary distinguishing for him. (This case is supposed to demonstrate that reference is socially determined.)[89] More radically, Putnam claims that the environment itself contributes to the meaning of natural kind terms. He imagines a Twin Earth, identical in every respect with Earth, except that the stuff called "water" on Twin Earth is a fundamentally different chemical compound (though indistinguishable at the level of everyday phenomena) than the stuff called "water" on Earth. Even supposing that Earthlings and Twin Earthlings have not yet developed their chemistry sufficiently to recognize the difference between the two substances (and so associate the same mental representations with the term "water"), their respective "water" words, Putnam thinks, have different references. "Meaning is interactional," Putnam concludes. "The environment itself plays a role in determining what a speaker's words, or a community's words, refer to."[90]

The CI account lends some credence to these anti-mentalistic insights. Interpretive affordances aren't "in the head"—they're properties of objects (including words) as those objects relate to potential interpreters, situated in environments. Words, furthermore, have the interpretive affordances they do on the basis of the role they play in a system (language) that facilitates overt communication: the intentional production of objects that offer particular interpretive affordances to an audience. Given that the practical, theoretical, and affective projects that lead humans to engage in interpretation are often communal rather than individual

45

undertakings, it should be no surprise that we can often rely on one another's expertise when we use terms to refer, sentences to assert, and so on.

The role that Putnam assigns to the environment in determining a term's reference is a more complex issue. On the CI account, reference is, in the first place, something agents do. If objects can be said to "refer," it's because interpretive agents refer *by* them. To alter slightly another thought experiment of Putnam's: if seaweed pulled out by a tide leaves channels in the sand that coincidentally inscribe "CHURCHILL," those markings don't refer to Churchill; but if Churchill himself, in an act of vanity, traced the same shape, it would.[91] In the example about beeches and elms, speakers are able to use a communal tool to refer to objects they can't themselves fully distinguish because they can rely on the abilities of other members of their community when they undertake joint action. But in cases like Twin Earth, Putnam suggests that speakers can refer to kinds that *no one* in the language community, nor any grouping of its members in concert, can currently distinguish with complete precision. Is that really possible?

Ultimately, from the CI perspective, this is a question about what can be done with language: a question that requires careful reflection on language and its use, rather than one that can be answered merely by appeal to a general account of meaning. Granted, the CI account of meaning regards an interpreter's environment (including her social world) as generating the problems in which the need for interpretation arises, as well as providing her with meaningful objects and even the resources necessary for the intentional production of more meaningful objects. But it doesn't follow from these generalities about the way that our environment contributes to meaning that *speakers* can index their terms to natural essences that no one can

46

currently characterize (as Putnam portrays pre-chemistry Earthlings as able to do, with respect to water).

Nonetheless, I think there's reason to believe that speakers can do something *like* this, at least. An underappreciated fact about language (underappreciated in its theory, but not its practice) is the fact that we use language not merely to communicate to others what we already know about the world, but also to structure our (often collaborative) efforts to *learn* about the world. Scientific terms, for example, are often introduced into discourse long before anyone has a precise (or even wholly correct!) understanding of the full suite of their referents' properties: consider the history of terms like GENE and ELECTRON.[92] Nonetheless, Putnam is probably right to insist that the *reference* of these terms doesn't change with every new piece of knowledge gleaned, or correction made, in the process of inquiry: Bohr's theory of electrons in 1934 was a successor to his theory in 1900—both theories of *electrons*—rather than a new theory about something else. More importantly for my current purposes, whether or not processes like these ultimately vindicate Putnam-style externalism about reference-fixation, they proceed along the lines that the CI account should lead us to expect. Presenting her "cognitive-historical" approach to conceptual innovation in science, Nancy Nersessian notes that

> Novel concepts arise from attempts to solve specific problems, using the conceptual, analytical, and material resources provided by the cognitive-social-cultural context in which they are created…A naturalist recasting of the problem of conceptual change in science shifts the focus of the problem from the conceptual structures themselves to the nature of the practices employed by human agents in creating, communicating, and replacing scientific representations of a domain. That is, it shifts the focus from the products to the processes, from the structures to the practices… To be successful in building an account of conceptual innovation, thus, requires both a model of the scientist

qua human agent and knowledge of the nature of the practices actually used in creating and changing conceptual structures.[93]

Meanings aren't *just* in the head. As the CI account insists, meaning, across its varieties, is a property of objects that compose an environment, understood in relation to the interpretive agents who inhabit that environment. A danger of studying meaning in purely linguistic contexts (or in mental contexts cashed out in linguistic terms) is the temptation it offers to focus exclusively on the semantic mappings between ready-made, abstract symbolic structures and the reality these structures seem to "represent," as if such relations could ultimately float free of the natures of the practical agents who employ them, and the environmental and social conditions in which they arise and evolve. But even if we succeed in resisting the temptation to treat language in this incomplete way, we still risk misunderstanding *meaning* if we restrict our attention to the role it plays in language. So, at any rate, I'll argue in the next section.

§3.3.2 Is Meaning Necessarily Social?

I started the last section by claiming that a significant portion of 20[th] century philosophy of language can be understood as a repudiation of Hobbes's thoroughly mentalist, thoroughly individualist theory of meaning. According to Hobbes, not only are the meanings of the words we use ideas in our minds, rather than things in the world, but the role they play in overt communication is only "accidental"—they are, first and foremost, aids to individual cognition: "sensible moniments" for storing and ordering our otherwise inconstant, fading, and desultory thoughts.[94] In response to claims like these, philosophers have more recently insisted on the fundamentally *social* nature of meaning. Perhaps the most famous instance of this sort of

insistence is Wittgenstein's *Philosophical Investigations*: his recommendation in that work that we understand meaning in the context of the social interactions ("language games") in which it arises, and his argument against the very possibility of a private language.[95] We can find similar, anticipatory sentiments in Dewey's work. According to Dewey, meaning comes into being *with* communication: "When communication occurs, all natural events are subject to reconsideration and revision; they are re-adapted to meet the requirements of conversation, whether it be public discourse or that preliminary discourse termed thinking. Events turn into objects, things with a meaning." Furthermore, "When the introspectionist thinks he has withdrawn into a wholly private realm of events disparate in kind from other events, made out of mental stuff, he is only turning his attention to his own soliloquy. And soliloquy is the product and reflex of converse with others; social communication is not an effect of soliloquy."[96]

With respect to *language* as it has evolved and actually exists, a picture like Wittgenstein's and Dewey's is almost certainly closer to the mark than Hobbes's. Putnam's observation about the linguistic division of labor shows one way that the social functions of language give rise to possibilities of meaning that are inexplicable in terms of a purely individualistic theory, and most current theories of the evolution of language (and the cognitive capacities underpinning it) place primary emphasis on its embeddedness in and facilitation of various sorts of social interaction.[97] Nonetheless, appealing to a more comprehensive view of meaning shows why objecting to Hobbes by claiming that meaning is *inherently* or *necessarily* social overstates the case. Instead, *linguistic* meaning inherits its sociability from the sociability of *language*. But because meaning in general, understood fundamentally in terms of the activity of interpretation, is a broader

phenomenon than the sort of "nonnatural" meaning characteristic of intentional communication, it's a mistake to restrict all meaning-related phenomena to those who engage in social discourse. In fact, not even the intentional use of symbols (objects intentionally manufactured to possess interpretive affordances) needs to be social in the way that Dewey seemed to suggest.

The first point is simple: objects have interpretive affordances (meanings) with respect to a wide range of projects, and only some of these are essentially social undertakings. An organism might have many asocial interests, for example, in whose light an ability to predict changes in the weather by interpreting natural signs is useful. Even the sort of interpretation that might be involved in tracking prey on the basis of scent or footprints isn't *social*, in Dewey's sense—there is no "community of meaning," no established "cooperation in activity in which there are partners, and in which the activity of each is modified and regulated by partnership."[98] If we're right to conceive of language as a publicly available tool for intentional communication, then the evaluative contexts within which it is produced and interpreted will necessarily be social (parasitically when not directly), as Dewey and Wittgenstein argue. But this observation doesn't extend to meaning in general.

If that argument doesn't itself redeem Hobbes, who was, after all, making claims about language himself, it's also worth noticing that even intentionally produced symbols can be used in the asocial way he suggests, even if such "marks" aren't the primary use of linguistic items. There is no deep metaphysical reason to suppose (nor, as far as I know, any strong empirical evidence to believe) that a contrived system of "sensible moniments" wouldn't aid the memory in just the way Hobbes suggests. A Robinson Crusoe figure, for example, might easily find it valuable to devise a

system of markings for differentiating the caves on his island—one mark for a cave that provides a useful storage space or shelter, another to indicate caves that are the dens of hostile animals, etc. These markings possess interpretive affordances but they imply no sociability, barring a trivialization of "social" to include the diachronic relations of a single self.

Tempering the critique of Hobbes in this way has implications for the familiar strategy of dismissing epistemological worries about "private" Cartesian mental states by observing that mental states have semantic content (on this strategy, a kind of meaning) and meaning is necessarily social, and therefore necessarily public. Clifford Geertz employs a strategy like this in his defense of the scientific credentials of cultural anthropology, which he characterizes as "not an experimental science in search of law but an interpretive one in search of meaning."[99] Having provided a "semiotic" account of culture (that is, having defined culture in terms of meaning) Geertz justifies culture as a proper object of scientific study by arguing that culture is neither "inside the head" nor hopelessly subjective—"Culture is public because meaning is."[100] I suspect that Geertz has things the wrong way around: cultural meaning is public because *culture* is. The robust sociability that supports Geertz's Ryle-inspired argument for the publicity of culture is a feature of the sort of *object* involved in cultural interpretation—using Geertz's own term, "social expressions"—rather than a feature of meaning itself.

§3.3.4 Meaning and "Content"

My goal in this chapter was to show that the varieties of meaning manifest in such a broad range of our discourse can be coherently understood as species of a common genus, and, consequently,

that an *ecumenical* approach to meaning—one that recognizes rather than ignores or denies MEANING's conceptual integrity—ultimately clarifies rather than obscures many of the issues we use "meaning" to articulate. For a concluding illustration, I want to return to one of the provocative claims I began with—Jerry Fodor and Zenon Pylyshyn's pronouncement that "like the Loch Ness Monster, meaning is a myth."[101]

Fodor and Pylyshyn occasionally write as if they expect to be the subjects of the next Michael Lewis explainer-turned-Hollywood-blockbuster:

> As we understand the jargon of Wall Street, a "contrarian" is someone whose practice is to buy when everybody else sells and to sell when everybody else buys. In effect, contrarians believe that what more or less everybody else believes is false more or less all of the time. This book is an essay in contrarian semantics.[102]

So it's dispiriting to learn, twenty pages after the Loch Ness Monster comparison, that ("of course") Fodor and Pylyshyn don't intend to deny that communication, paraphrase, translation, or (as far as I can tell) any other everyday activity that involves meaning or its preservation "really happen." Instead, they just warn us not to expect a "first-class" scientific theory that will tell us, definitively, which of two translations of *Pride and Prejudice* is better, or whether two terms are "strictly speaking" synonymous.[103] When we discover that the sense in which Fodor and Pylyshyn deny that meaning exists leaves nearly all of the ways that "meaning" is deployed in everyday discourse unaltered, the comparison with the Loch Ness Monster seems a lot less swashbuckling. It might also lead us to question the utility of using the term "meaning" as they do—defined in terms of a relatively technical theory of "conceptual content," rather than with respect to the way the word is used in far more heavily trafficked (and important) discursive contexts. Things might

be clearer (and ultimately less dispiriting, even if momentarily less exciting) if they just said that they don't think that the content of mental representations is best understood in terms of intensions.

In fact, this is another area where the mistake of understanding meaning primarily according to its linguistic manifestations causes confusion. From the CI perspective, the issue that concerns Fodor and Pylyshyn isn't an issue about *meaning* at all. The kind of primitive, innate, purely individualistic "concepts" posited by the representational theory of mind Fodor and Pylyshyn advocate are not themselves objects of interpretation: we don't encounter superficially opaque "mentalese" symbols and then interpret their import. Setting aside the question of whether this sort of representationalist program is ultimately the right paradigm for the study of the mind, gliding thoughtlessly between the way that mentalese symbols have "content" and terms in natural language do is dangerous. If there are mentalese symbols, they are neither the objects nor the tools of public dispute (consider "freedom"), nor pieces of our cultural inheritance (like "gene"); we don't use them, together, in shared enterprises; we don't use them to insult one another, or to express our sense of irony; they aren't a medium in which we make art. As we've already noticed, the social functions of language give its objects special interpretive features, and the fact that we *do* interpret language raises distinctive questions. Except insofar as we identify thought with subvocalizations of natural language, then, the strategy of modeling the medium of thought on our medium of communication needs to be undertaken carefully—neither can be straightforwardly read off of the other. Importing the term "meaning" as a residue from debates

53

about linguistic content, only to adopt a nihilistic stance, is therefore likely to confuse much more than it clarifies.

Conclusion

Though the concept MEANING is versatile, a fundamental integrity underlies its deployment across discursive contexts. That integrity, I've argued, is best understood according to the structural role that meaning plays in the practice of interpretation. Meanings are the (licensed) outputs of potential acts of interpretation: an object's distinctively interpretive affordances. Interpretation is a functionally structured, mentally mediated activity; undertaken by organisms in an environment for the purpose of guiding action, belief, or affection; subject to norms of validity and fruitfulness.

Instead of obscuring important distinctions, taking an ecumenical approach to the theory of meaning—one that integrates, rather than segregates, meaning's varieties—helps us make the important distinctions clear. Varieties of meaning can be classified according to various features of the kind of interpretation in question—most naturally, according to the kind of object being interpreted. Beginning from a comprehensive perspective helped us recognize the shortcomings of assimilating natural meaning with either semantic information or ecological affordances *simpliciter*. It helped to highlight and characterize some of the ways that accounting for linguistic meaning generates distinctive challenges, and to clarify the ways that environmental and social factors can contribute to the way terms can be used to refer. It warns against conflating questions about meaning with questions about content at large.

My reasons for developing and presenting the CI account, however, weren't just to provide an occasion to show the way that old controversies can be recast (if not resolved) in new terms. Instead, I think that a comprehensive understanding of meaning is essential to a full appreciation of meaning's ethical dimensions—in both the way that meaning, as a practice, generates distinctive sorts of ethical challenges, and correlative ethical obligations; and also in the way that the concept MEANING is deployed in ethical discourse (like when we aspire to lead "meaningful" lives). In the next chapter, I offer an analysis of "meaning" in ethical contexts, showing that the CI account of meaning helps us to better understand worries meaning in and of our lives. In the final chapter, I turn my attention to irony, showing how the comprehensive-interpretive account of meaning provides the resources for developing a satisfying account of irony's nature, in whose terms we can better understand the aesthetic and ethical controversies irony generates.

---

[1] Paul Henle, "The Problem of Meaning:," in *American Philosophical Association Centennial Series* (Philosophy Documentation Center, 2013), 24–25, http://www.pdcnet.org/oom/service?url_ver=Z39.88-2004&rft_val_fmt=&rft.imuse_id=apapa_2013_0229_0241_441&svc_id=info:www.pdcnet.org/collection.

[2] Kwame Anthony Appiah, *The Ethics of Identity* (Princeton University Press, 2007); Susan Wolf et al., *Meaning in Life and Why It Matters* (Princeton, N.J.: Princeton University Press, 2012).

[3] Clifford Geertz, "Thick Description: Toward an Interpretive Theory of Culture," in *The Interpretation of Cultures: Selected Essays* (Basic Books, 1973), 29.

[4] Jerry A. Fodor, *Psychosemantics: The Problem of Meaning in the Philosophy of Mind* (Cambridge, Mass.: A Bradford Book, 1989), 99.

[5] Donald Davidson, "Theories of Meaning and Learnable Languages," in *Inquiries into Truth and Interpretation*, 2 edition (Oxford: Clarendon Press, 2001), 3, 8.

[6] Max Weber, *Economy and Society: 2 Volume Set* (University of California Press, 1978), 506. If you worry you might be among the afflicted, Weber elaborates: "The salvation sought by the intellectual is always based on inner need, and hence it is at once more remote from life, more theoretical and more systematic than salvation from external distress, the quest for which is characteristic of nonprivileged strata. The intellectual seeks in various ways, the casuistry of which extends into infinity, to endow his life with a pervasive meaning, and thus to find unity within himself, with his fellow men, and with the cosmos. It is the intellectual who conceives of the 'world' as a problem of meaning. As intellectualism suppresses belief in

magic, the world's processes become disenchanted, lose their magical significance, and henceforth simply 'are' and 'happen' but no longer signify anything. As a consequence, there is a growing demand that the world and the total pattern of life be subject to an order that is significant and meaningful."

[7] Appiah, *The Ethics of Identity*, 120.

[8] Henle, "The Problem of Meaning," 25.

[9] Willard Van Orman Quine, "Two Dogmas of Empiricism," in *From a Logical Point of View: Nine Logico-Philosophical Essays, Second Revised Edition* (Cambridge, Mass: Harvard University Press, 1980), 22.

[10] As I hope is apparent, by "everyday use" I mean roughly what many philosophers use the modifier "ordinary" to indicate. I prefer "everyday" to "ordinary" because I dislike the suggestion that using language to think and communicate philosophically (or scientifically, or mathematically, etc.) isn't an ordinary way of using language. I hope "everyday" captures better the distinction between "ordinary" (as traditionally used) and "philosophical" or "theoretical" use as I see it: not a distinction between contrast classes, but a distinction between a general class (everyday usage) and one of its subclasses (philosophical or theoretical usage). I think the idea that philosophical discourse is, in some important sense, extraordinary is bad in both directions: bad in the direction toward which Russell, in his more extreme moments, strayed, regarding philosophical usage as extraordinary and superior (because of its greater rigor and clarity); and bad in the direction toward which Wittgenstein, in his more extreme moments, strayed, regarding philosophical usage as extraordinary and inferior (because empty and alienated from the living, immediate, communal interests that make language meaningful).

[11] John Dewey, *Democracy and Education* (Free Press, 1997), 5–6.

[12] Ibid., 28–29. This broad conception of meaning plays a foundational role in Dewey's systematic philosophy, and is central to his influential conceptions of both democracy and education (the former of which has a prominent contemporary advocate in Elizabeth Anderson). "A democracy," Dewey contended, "is more than a form of government; it is primarily a mode of associated living, of conjoint communicated experience." Its mark is the existence *and mutual recognition* of "varied points of shared common interest," so that "each [individual] has to refer his own action to that of others, and to consider the action of others to give point and direction to his own." Where, under conditions of systematic inequality and segregation, such shared interests (and their recognition) decay, "the influences which educate some into masters, educate others into slaves. And the experience of each party loses in meaning, when the free interchange of varying modes of life-experience is arrested. A separation into a privileged and a subject-class prevents social endosmosis. The evils affecting the superior class are less material and less perceptible, but equally real. Their culture tends to be sterile, to be turned back to feed on itself; their art becomes a showy display and artificial; their wealth luxurious; their knowledge overspecialized; their manners fastidious rather than humane." (84-87) Unsurprisingly, given the tight link Dewey made between democracy and education, meaning is similarly essential to the way he defined the latter: "We thus reach a technical definition of education: It is that reconstruction or reorganization of experience which adds to the meaning of experience, and which increases ability to direct the course of subsequent experience." (77)

[13] Jerry A. Fodor and Zenon W. Pylyshyn, *Minds without Meanings: An Essay on the Content of Concepts* (Cambridge, Massachusetts: The MIT Press, 2014), 58; W. E. B. Du Bois, "The Souls of Black Folk," in *Du Bois: Writings* (Library of America, 1996), 358.

[14] Alexis Burgess and David Plunkett, "Conceptual Ethics I," *Philosophy Compass* 8, no. 12 (December 1, 2013): 1091–1101.

[15] This slogan is a preliminary formulation that receives some necessary qualification below, but for now it serves well enough.

[16] Dewey, *Democracy and Education*, 28.

[17] The tradition's genealogy is usually traced back to the influence of the later Wittgenstein. For a fully developed, influential theory of meaning that treats the *act* of meaning as explanatorily primary, and symbolic meaning as explanatorily derivative, see the work of Paul Grice, who argued that the meaning of utterance types should be explained in terms of what speakers who token those utterance types on particular occasions generally mean thereby to do. See the essays collected in Paul Grice, *Studies in the Way of Words* (Harvard University Press, 1991), especially "Utterer's Meaning, Sentence-Meaning, and Word-Meaning".

[18] The point here is that "mean," even though a verb, doesn't, across its full range of usage, pick out a *practice*. "Interpret," by contrast, mostly does. (There are deviant cases, but they're rare.) As the body text suggests, in *some* cases in which our use of "mean" isn't practical in the relevant sense (isn't used, that is, to denote the activity or behavior of an agent)—as when, for example, we say of some string of language that it means such-and-such—it may be possible to explain this non-practical "mean" in terms of an agential mean: Grice's project of explaining what he calls "timeless meaning" in terms of "occasion meaning" is an example of that sort of attempt. But these attempts aren't viable in cases where we speak of the meaning of objects that aren't partially constituted by the role they play in systems of overt communication (as, for example, words are). In cases of so-called natural meaning—like when we discuss the meaning of the iridium deposits found at the K-Pg layer in the geological record—there's no relevantly associated *agential* meaning on whose basis we may explain the non-practical meaning in question. The likeness between natural and nonnatural non-agential meanings, and, therefore, the basis for capturing them under the same general concept, is, instead, the potential they present to *interpreting* (rather than merely *meaning*) agents. This is one basic reason why the practice of interpretation, rather than meaning, makes for a better foundation for a comprehensive account of meaning. I discuss natural meaning in greater detail in §3.2.

[19] When I speak of Hubbard's interpretation of the broken lock, I don't, of course, mean to imply that that interpretation is undertaken in isolation from considerations about other aspects of the crime scene, like the fibers on Swann's shoes, or the fact that the garden looks undisturbed by human travel—both of which are, of course, relevant to Hubbard's concluding that **(3)** rather than **(2)**. Naming the broken lock as the primary object of interpretation in this case just picks out a phase in Hubbard's thought process: the phase in which Hubbard adopts an interpretive perspective in which the broken lock is the structural focal point. From this perspective, he considers the other features of the crime scene primarily in light of what they indicate about the significance of the broken lock.

[20] For my purposes (and, I think, for most others) agency is best thought of as coming in degrees. The standard I proposed in the body text is a relatively undemanding one. It doesn't require, for example, that an organism be able to represent to itself, or to evaluate, the goals (*as* goals) that motivate/explain/rationalize its behavior. Not all agents who meet this standard will necessarily be interpreters, for reasons that will become clear as this section progresses: interpretation requires that an organism have certain cognitive capacities *and* certain cognitive limitations. But all interpreters necessarily meet at least this standard of agency.

[21] John Dewey, "Beliefs and Existences," in *The Influence of Darwin on Philosophy: And Other Essays in Contemporary Thought* (H. Holt and Company, 1910), 171; John Dewey, *Experience and Nature* (Courier Dover Publications, 1958), 28.

[22] Here I use "object" in its grammatical rather than its metaphysical sense. I use it, that is, to indicate the structural role it plays in the practice of interpretation (reflected in the grammar of sentences in which "interpret" appears as a verb), rather than to pick out a particular kind of entity. Utterances, actions, and

situations, for example, may not be regarded as *objects* in the metaphysical sense of that term, but they can all serve comfortably as objects of interpretation.

[23] The grammatical properties of "inspire" are somewhat different than those of "interpret": treated as a functional process, the argument of an inspirational (as opposed to an interpretive) structure is usually the subject of the verb (rather than its direct object), or, in passive constructions, a prepositional object.

[24] Corresponding to the requirement that an interpretation ought to be fruitful is the important idea, arguably necessary for the very possibility of overt communication, that instances of agential meaning out to be *relevant*. See Paul Grice, "Logic and Conversation," in *Studies in the Way of Words* (Harvard University Press, 1991), 22–40, and Dan Sperber and Deirdre Wilson, *Relevance: Communication and Cognition* (Wiley, 1996).

[25] Translation is another contrast case here. Translations aim at preserving meaning from one system of overt communication to another; they don't, *qua* translation, have to be fruitful. An English-to-French translation may be useless to someone who speaks only German, but that doesn't make it a worse *translation*. Put another way, interpretation involves a *value asymmetry* between its input and output that translation does not. All things equal, an interpreter ought to be better off with respect to some of her extra-interpretive interests on the basis of having made a successful interpretation: an object's *meaning* has some value to the interpreter that the uninterpreted object of interpretation doesn't afford. (Compare the value of an utterance made in a language you speak with the value of an utterance made in a language you don't.) Translation, a process defined *in terms of interpretation* (because it's defined in terms of meaning), doesn't have this requirement. Rather than taking us from an object of interpretation to a meaning, translation takes us from one (special kind of) object of interpretation to another. When humans translate, this is almost always accomplished *by means of* interpretation, but that's not necessary. Google Translate, for example, makes successful translations without engaging in any interpretation at all.

[26] Some readers might worry that the example involving **(5)** and **(6)** trades on a conflation of natural meaning and linguistic meaning, so that **(6)** reads as a "better" interpretation of Alice's utterance than **(5)** because we implicitly treat Alice's utterance as a *linguistic* object, and *qua* linguistic object, **(6)** is at least a lot closer to its meaning than **(5)**. I don't think the example does trade on that sort of ambiguity—**(6)** is better than **(5)** without requiring the kind of qualification employed in the explanation above, because **(6)** is equally valid and more fruitful with respect to the evaluative context—but it's easy enough to produce examples in which object-type ambiguities aren't possible. Consider, for example, the ways that Holmes is generally better at interpreting physical traces than Watson (even in cases where Watson's interpretations are licensed).

[27] Geertz, "Thick Description: Toward an Interpretive Theory of Culture," 5.

[28] David Fleming, "To Each His Own," *ESPN.Com*, July 10, 2012, http://espn.go.com/espn/magazine/archives/news/story?page=magazine-20100322-article34. Quoted in Matthew B. Crawford, *The World Beyond Your Head: On Becoming an Individual in an Age of Distraction* (New York: Farrar, Straus and Giroux, 2015), 82.

[29] See, for example, the account of perception Russell offers in Bertrand Russell, *The Problems of Philosophy* (Cosimo, Inc., 2007). Almost no contemporary philosophers of perception subscribe to sense-data accounts, but there does remain vigorous debate about the extent that perception involves a kind of inference. (Even for the inferentialists, though, the inferences in question are unconscious.) For a brief summary of the current state of this debate, and an argument against regarding visual perception as cognitively mediated (even unconsciously) see Nico Orlandi, *The Innocent Eye: Why Vision Is Not a Cognitive Process* (Oxford: Oxford University Press, 2014). My own purposes are unaffected by however

interpretive (in my sense) perception turns out to be. (For Dewey, *perception*—which he thought of as both active and intelligent—was interpretive, but bare sensation wasn't.) I am reasonably confident, however, that standard cases of basic sense perception are less interpretive than standard cases of basic language processing.

[30] Martinich has Quine and Davidson primarily in mind here, and goes on to dismiss his initial worry, perhaps dubiously, by claiming that Quine and Davidson have effectively shown that "the difference between 'understanding' and 'interpretation' is a matter of usage, not meaning"; that, for understanding, interpretation, *and* translation, "the cognitive activity is the same in each case"; and, finally, that "understanding, interpretation, and translation are the same thing." A. P. Martinich, "Interpretation and Translation," in *The Philosophy of Language*, ed. A. P. Martinich, 5 edition (New York: Oxford University Press, 2006), 541–43.

[31] Dewey, *Democracy and Education*, 29.

[32] Perhaps it's worth re-emphasizing: my point here doesn't rest on any categorical claim that visual perception *isn't* interpretive. Even readers who think my use of visual perception as an illustrative contrast class involves a misconstrual of the nature of visual perception have, in the process of forming their objection, likely appreciated the point of the illustration.

[33] Richard Wollheim, "Seeing-as, Seeing-in, and Pictorial Representation," in *Art and Its Objects* (Cambridge University Press, 1980), 216.

[34] Elisabeth Camp, "Showing, Telling and Seeing. Metaphor and 'Poetic' Language," *Baltic International Yearbook of Cognition, Logic and Communication* 3, no. 1 (January 1, 2007): 4.

[35] I borrow the nutrition analogy from Gary Hatfield, who uses it to illustrate his relationalist theory of color. See Gary Hatfield, "Objectivity and Subjectivity Revisited: Colour as a Psychobiological Property," in *Colour Perception: Mind and the Physical World*, ed. Rainer Mausfeld and Dieter Heyer (Oxford University Press, 2003), 187-202.

[36] As will become clear in subsequent discussion (§3.2), I use the term "affordance" in sympathy with James J. Gibson's coinage in *The Ecological Approach to Visual Perception* (Psychology Press, 1986), though I ultimately favor an understanding of meaning that diverges from Gibson's in important respects.

[37] Paul Grice, "Meaning," in *Studies in the Way of Words* (Harvard University Press, 1991), 213.

[38] Paul Grice, "Utterer's Meaning and Intentions," in *Studies in the Way of Words* (Harvard University Press, 1991), 101, 108.

[39] Raymond W. Gibbs and Herbert L. Colston, *Interpreting Figurative Meaning* (Cambridge University Press, 2012), 6.

[40] My point isn't that readers familiar with the figurative language literature won't ultimately be able to make sense of this passage in one way or another—even, perhaps, the sense Gibbs and Colston intended. The point is how readily experienced readers impute polysemous interpretations of "means" and "meaning" to passages like these (but not always, I suspect, with a high degree of coordination), and the apparent ease with which writers exploit that willingness. The result is, at best, a degree of communicative obscurity; and at worst, genuine confusion at the level of individual thought.

[41] See especially Grice, "Meaning"; Paul Grice, "Meaning Revisited," in *Studies in the Way of Words* (Harvard University Press, 1991).

[42] Grice, "Meaning." If you think *American Beauty* is a good movie (and I know, many people do), please forgive my self-indulgent example. Maybe you're right—in which case I invite you to substitute in a movie title that you think makes the example work. Please write the title here: _____.

[43] Dorit Bar-On, "Expressive Communication and Continuity Skepticism," *The Journal of Philosophy* 110, no. 6 (June 1, 2013): 293–330.

[44] Ruth Garrett Millikan, *Varieties of Meaning: The 2002 Jean Nicod Lectures* (Cambridge, Mass: The MIT Press, 2004).

[45] Grice, "Meaning Revisited," 291–92.

[46] Carol E. Cleland, "Prediction and Explanation in Historical Natural Science," *The British Journal for the Philosophy of Science* 62, no. 3 (September 1, 2011): 552.

[47] Several of Cleland's papers include engaging case-study length accounts of this discovery. For book-length accounts see Walter Alvarez, *"T. Rex" and the Crater of Doom*, (Princeton, New Jersey: Princeton University Press, 2015); James Lawrence Powell, *Night Comes to the Cretaceous: Dinosaur Extinction and the Transformation of Modern Geology* (New York: St. Martin's Press, 1998); and Lisa Randall, *Dark Matter and the Dinosaurs: The Astounding Interconnectedness of the Universe* (New York, NY: Ecco, 2015). (Randall narrates the discovery in the process of presenting her hypothesis that dark matter in the Milky Way is responsible for the apparent periodicity of comet strikes, because it periodically disturbs objects in the Oort Cloud.) For a detailed account of the significance of the co-occurrence of iridium, shocked minerals, and Ni-rich spindels, see Peter Schulte et al., "The Chicxulub Asteroid Impact and Mass Extinction at the Cretaceous-Paleogene Boundary," *Science* 327, no. 5970 (March 5, 2010): 1214–18.

[48] For discussions that treat Dretske's semantic theory of information as a candidate analysis for Grice's notion of natural meaning, see, e.g., Frederick Adams, "Causal Contents," in *Dretske and His Critics*, ed. Brian P. McLaughlin (Basil Blackwell, 1991); and Robert K. Shope, *The Nature of Meaningfulness: Representing, Powers, and Meaning* (Rowman & Littlefield, 1999). Dretske himself notes that "meaning" in Grice's *natural* sense "is close (if not equivalent) to the ordinary sense of 'information.'" Fred I. Dretske, *Knowledge and the Flow of Information* (MIT Press, 1983), 242.

[49] Dretske, *Knowledge and the Flow of Information*, 7.

[50] Fred I. Dretske, "Précis of Knowledge and the Flow of Information," *Behavioral and Brain Sciences* 6, no. 01 (1983): 56. Dretske concedes that there are circumstances in which it isn't possible to formulate a determinate propositional description of the information a signal carries about its source—for instance, when the conditional probability of each of the possible states of the source, given the signal, is greater than zero (and, therefore, less than one). (If any of the possible states (say, F) at the source s has a conditional probability of zero given the signal, then we can at least determinately describe its content as "s isn't F"; complementarily, if any of the possible states (say, G) has a conditional probability of 1 given the signal, we have "s is G." But if no possible state is definitively ruled out, neither of these strategies will be viable. In cases like these Dretske recommends characterizing the information in terms of the state s is "probably" in, but this is a workaround, since a specification of the probabilities that s is in a given state isn't, strictly speaking, a description of a state of s that obtains with probability 1, conditional on the state of r.) But this slight complication, which Dretske plausibly argues isn't a weakness of his account, is ancillary to the line of argument I'm pursuing in the body text. For Dretske's discussion of the issue see *Knowledge and the Flow of Information*, 68-69.

[51] James Jerome Gibson, *The Ecological Approach to Visual Perception* (Psychology Press, 1986), xiii.

[52] Ibid., 36.

[53] Ibid., 38.

[54] Ibid., 39.

[55] Ibid., 127.

[56] Dretske, *Knowledge and the Flow of Information*, 7.

[57] Keep in mind, per the discussion in §2.5, that interpretive effort needn't always be strenuously thoughtful, and shouldn't be taken to rule out certain sorts of automatic processing.

[58] Dretske, *Knowledge and the Flow of Information*, 72.

[59] "A signal from a distant, and rapidly receding, galaxy is not devoid of information simply because we are incapable of determining what the relevant conditional probabilities are, and hence what information (or *how much* information) the signal contains. For the amount of information contained in the signal depends, *not* on the conditional probabilities that *we can independently verify*, but on the conditional probabilities themselves." Ibid., 56.

[60] Ibid., 72.

[61] The question of how to make precise sense of the conditional probabilities Dretske's theory relies upon is vexed. As we've seen, Dretske repeatedly emphasizes the *objective* status of semantic information carried by a signal, viz. "It cannot be stressed too strongly…that the set of conditional probabilities defining the equivocation between source and receiver are meant to be *objective* features of the communication system. We may not *know* what these conditional probabilities are, we may not be *able* to determine what they are, but this is irrelevant." (55) However, as many of Dretske's illustrations make clear, and as many of his critics have noted, these conditional probabilities are, in an important sense, framework-relative. They depend, for example, on which aspects of the system in question are classified as the source and its possible states, which count as channel conditions (which are not relevant to calculating the amount of information produced at the source), and so on. Some critics have contended that the seeming arbitrariness of framework selection renders the conclusions about knowledge that Dretske derives from his theory of information trivial (see, e.g., Gilbert Harman, "Knowledge and the Relativity of Information," *Behavioral and Brain Sciences* 6, no. 1 (1983), 72); others (Millikan is an example) suggest that probabilistic analysis on the basis of natural law prevents information from being understood as ever applying to *individuals* (which are not the subjects of natural laws). Whether or not it makes sense to think of the conditional probability of a meteor strike given the observed composition of the K-Pg layer as anything less than 1 (presuming the strike to have actually occurred), or even whether it makes sense to talk of such a "conditional probability" *at all* depends on the proper way for selecting the relevant framework, along with the way that criticisms like those mentioned above are either credited or defused. These issues, however, are ancillary to the point I make in the body text, which is that differentiating natural meaning from information in the ways the CI account prescribes provides a conceptual framework within which questions like these can be meaningfully asked and addressed.

[62] Millikan, *Varieties of Meaning*, 33.

[63] Ibid., 33–34. In fact, Millikan argues, "Not even a *high* probability is always required for an attribution of natural meaning." (36)

[64] Gibson, *The Ecological Approach to Visual Perception*, 129.

[65] Ibid., 140, 138.

[66] Ibid., 134.

[67] Ibid., 42.

[68] Ibid., 63.

[69] Ursula K. Le Guin, *A Wizard of Earthsea* (Houghton Mifflin Harcourt, 2012), 55.

[70] Galileo to Mark Welser, May 4, 1612, in Galileo Galilei, *Discoveries and Opinions of Galileo*, trans. Stillman Drake (New York: Anchor, 1957), 92.

[71] William Shakespeare, Stephen Orgel, and A. R. Braunmuller, *Romeo and Juliet*, ed. Peter Holland (New York, New York: Penguin Classics, 2016), 39.

[72] Paul Horwich, *Meaning* (New York: Clarendon Press, 1999), 1.

[73] Elisabeth Camp, "Logical Concepts and Associative Characterizations," in *The Conceptual Mind: New Directions in the Study of Concepts*, ed. Eric Margolis and Stephen Laurence (Cambridge, Massachusetts: The MIT Press, 2015), 592.

[74] David Lewis, "General Semantics," *Synthese* 22, no. 1/2 (1970): 22.

[75] William G. Lycan, *Philosophy of Language: A Contemporary Introduction (Routledge Contemporary Introductions to Philosophy)* (Routledge, 2008), 3.

[76] Thomas Hobbes, "The Elements of Philosophy," in *The Philosophy of Hobbes in Extracts and Notes Collated from His Writings*, ed. Frederick Woodbridge (H. W. Wilson Company, 1903), 15.

[77] John Locke, *An Essay Concerning Human Understanding*, excerpted in A. P. Martinich and David Sosa, eds., *The Philosophy of Language*, 6th ed. (Oxford University Press, USA, 2012), 657.

[78] Gottlob Frege, "On Sense and Reference," in *Meaning*, ed. Mark Richard, Wiley Blackwell Readings in Philosophy (Wiley-Blackwell, 2003).

[79] Willard V. Quine, "Quantifiers and Propositional Attitudes," *The Journal of Philosophy* 53, no. 5 (1956): 180.

[80] Donald Davidson, "Truth and Meaning," in *Inquiries into Truth and Interpretation*, 2 edition (Oxford: Clarendon Press, 2001), 90. The remark about abstractness and (especially) obscure identity conditions makes it relatively clear that Davidson is trying to distinguish himself from Quine, here, but Quine also argued for the uselessness of meaning-entities—at least with respect to names, the case in which positing entities seems most attractive. See Willard V. Quine, "On What There Is," *The Review of Metaphysics* 2, no. 1 (1948): 21–38.

[81] Lycan, *Philosophy of Language*, 2.

[82] OK, it isn't actually *that* hard. But we can grant Lycan (and ourselves) some reasonable *ceteris paribus*-spirited assumptions—like neither **(10)** nor **(11)** is an instance of employing a pre-established secret code, and the doxastic context isn't curiosity about whether a particular inscriber is a competent user of the English language. I'll also ask you to put aside worries about whether **(10)** and **(11)** treated as tokens, as they appear in my text and Lycan's, are counter-examples to the point I've tried to make on their basis. Complexities like those are mostly beside the point, particularly because, contrary to what I've suggested so far, Lycan himself doesn't even necessarily need them. When he calls **(10)** and **(11)** "meaningless," he treats them as candidates for being abstract linguistic expression-types, and, as I'm in the process of arguing, the "meaning" of such types isn't properly cashed out in terms of the full, undiscriminated suite of interpretive affordances inscriptive tokens (i.e., tokens in virtue of shape only) **(10)** or **(11)** would have in any possible context of interpretation.

[83] As Davidson puts the point: "Nothing said so far limits first meaning to language; what has been characterized is (roughly) Grice's non-natural meaning, which applies to any sign or signal with an intended interpretation. What should be added if we want to restrict first meaning to linguistic meaning? The usual answer would, I think, be that in the case of language the hearer shares a complex system or theory with the speaker, a system which makes possible the articulation of logical relations between utterances, and explains the ability to interpret novel utterances in an organized way." Donald Davidson, "A Nice Derangement of Epitaphs," in *Truth, Language, and History*, 1 edition (Oxford : New York: Clarendon Press, 2005), 93.

[84] Gottlob Frege, *The Foundations of Arithmetic: A Logico-Mathematical Enquiry into the Concept of Number*, trans. J. L. Austin, 2nd Revised edition (Northwestern University Press, 1980), 71. I follow Zoltan Gendler Szabo (and standard German) in translating "Satz" as "sentence", rather than as "proposition" (J.L.

Austin uses "proposition" in his standard translation). See Zoltan Gendler Szabo, "Compositionality," in *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta, Summer 2017 (Metaphysics Research Lab, Stanford University, 2017), https://plato.stanford.edu/archives/sum2017/entries/compositionality/.

[85] Paul Grice, "Utterer's Meaning, Sentence-Meaning, and Word-Meaning," in *Studies in the Way of Words* (Harvard University Press, 1991), 118. When I say "roughly" in the passage above, I mean "roughly"! Grice's distinction between an utterance's "total signification" and its conventional meaning is, in fact, finer-grained than the distinction between its full suite of interpretive affordances and its conventional meaning, because even "total signification" is limited, for Grice, to an utterance's nonnatural meaning, whereas its full suite of interpretive affordances is not so limited. But this detail is unimportant to my point about the inadequacy of restrictivism about linguistic meaning, which Grice's finer distinction demonstrates just as well. (As we've already noticed, for Grice, even nonnatural meaning is broader than, and explanatorily prior to, linguistic meaning.)

[86] Hobbes, "The Elements of Philosophy," 16–17.

[87] Hilary Putnam, "The Meaning of 'Meaning,'" in *Philosophical Papers: Volume 2, Mind, Language and Reality* (Cambridge University Press, 1979), 227. The way I've framed things in this paragraph runs together at least two disputes—the dispute over "psychologism" (as Frege called it), and the dispute over whether content is "internally" or "externally" determined, but the distinction isn't essential for the paragraph's general framing purposes, and Putnam's slogan is, I think, reasonably applied to both disputes (and was so intended).

[88] Hilary Putnam, *Representation and Reality*, Revised ed. edition (Cambridge, Mass.: A Bradford Book, 1991), 19.

[89] Ibid., 22–23.

[90] Ibid., 34–36.

[91] Hilary Putnam, "Brains in a Vat," in *Reason, Truth and History* (Cambridge University Press, 1981), 1.

[92] For brief case studies, see Putnam, *Representation and Reality*, 13-17 (electrons), and Kwame Anthony Appiah, *Thinking It Through: An Introduction to Contemporary Philosophy* (Oxford University Press, 2003), 130–36 (genes).

[93] Nancy J. Nersessian, *Creating Scientific Concepts* (Cambridge, Mass.: A Bradford Book, 2010), ix, 5.

[94] Hobbes, "The Elements of Philosophy," 13, 16.

[95] Ludwig Wittgenstein, *Philosophical Investigations*, ed. P. M. S. Hacker and Joachim Schulte, 4th ed. (Wiley-Blackwell, 2009). The first recommendation of the "language game" methodology (as a therapy for the confusion caused by imagining that every word has a "meaning") appears in §5; the private language argument is contained in §§243-315.

[96] Dewey, *Experience and Nature*, 166, 170.

[97] See, e.g., Merlin Donald, *Origins of the Modern Mind: Three Stages in the Evolution of Culture and Cognition* (Harvard University Press, 1993); Kim Sterelny, *Thought in a Hostile World: The Evolution of Human Cognition*, 1 edition (Malden, MA: Wiley-Blackwell, 2003); Michael Tomasello, *Origins of Human Communication* (Cambridge, Mass.: A Bradford Book, 2010).

[98] Dewey, *Experience and Nature*, 179.

[99] Geertz, "Thick Description: Toward an Interpretive Theory of Culture," 5.

[100] Ibid., 12.

[101] Fodor and Pylyshyn, *Minds without Meanings*, 58.

[102] Ibid., 65. If you're wondering how anyone can be wrong more than "all of the time," don't ask me.

[103] Ibid., 76–77.

ETHICAL MEANING

A Comprehensive-Interpretive Account of Meaning in and of Life

Introduction

In the mid-1960s, Nina Simone had a conflict with her husband and manager, Andrew Stroud, over the trajectory of her career. Simone wanted to use her music to address the black struggle for civil rights, but Stroud disdained these activities, arguing (probably correctly) that they limited Simone's commercial appeal. Their daughter Lisa described the difference like this: "My father had a strategic plan in terms of how mom's career was going to go. He wanted her to be able to win all the awards and to become the huge star that he knew she could be…But she wanted something more. There was something missing in her—some *meaning*."[1]

This distinctively ethical sense of "meaning" is a familiar and important part of our everyday ethical discourse. We invoke it when we make major life decisions, like which career path to pursue, or when we attempt to describe the special sort of value we assign to marriage or child-rearing or friendship. In more extreme circumstances, the conviction that life has meaning has given some people the will and courage to endure almost unimaginable suffering, even when all hope of personal happiness seems lost.[2] Conversely, when we're afflicted with depression, even material comfort, professional success, and a conventional family life are no safeguard against despair: life, and the relationships and activities it comprises, seem *meaningless*. And many of us, like Simone, suffer more or less recurrent bouts of anxiety about the meaning (or meaninglessness) of or in our lives. But what is it, exactly, that we worry is missing?

64

The answer I offer in this chapter aims to help us achieve a better understanding of our curiosity, rather than its through satisfaction. My goal, in other words, is to clarify the *concept* of meaning as it functions in ethical contexts (henceforward, "ethical meaning"), rather than to argue for a particular *conception* of ethical meaning. I depart from much of the recent philosophical literature on this topic by regarding the way "meaning" functions *across* discursive contexts (the linguistic, natural, and literary senses of "meaning," for example) as centrally significant to understanding the way it functions in ethical contexts: understanding ethical meaning requires understanding meaning in general.

It's easy to see the attraction of treating meaning's different senses separately. We shouldn't, for example, expect that a theory of semantics that neatly explains the compositionality and generativity of language will provide the antidote to what William James called "the nightmare or suicidal view of life";[3] nor should we expect even a deep understanding of life's meaning to answer questions about, say, reference in intensional contexts. In fact, attempts to assimilate meaning's senses risk impoverishing one or more of its dimensions to accommodate another. Consider the triviality of a position that Thaddeus Metz canvasses (but does not endorse) in his recent survey of theories of meaning in life—that statements about life's meaning are not "well-formed propositions" because "to be meaningful is just to be a symbol, and since life cannot be a symbol, life is not the sort of thing that can be meaningful."[4] We shouldn't rule out the possibility that our thought and talk about ethical meaning is confused, but we're right to be unimpressed by the pedantic diagnosis that the confusion arises from an inability to distinguish lives from words.

Nonetheless, I believe that a substantive, articulable integrity underlies our use of "meaning" across contexts, and that appreciating the character of that integrity helps us better understand meaning in its more specialized senses—particularly, for my current purpose, its ethical sense. In the previous chapter, I proposed that we understand meaning, in general, according to its role in the structure of *interpretation.* Whether the context is ethical, aesthetic, or linguistic, **meanings are the outputs of acts of interpretation**.[5] Reflecting on the distinctive characteristics of interpretation in general, and then considering the way these function in ethical contexts, usefully clarifies what we mean when we use "meaning" in its ethical sense, and provides natural answers to several of the central questions that arise from that usage.

My attempt to substantiate that claim is an effort in three parts. First, I'll try to develop a richer sense of the challenge a theory of ethical meaning should meet. What, broadly, is the kind of work we expect the concept of meaning to do in everyday and theoretical ethical discourse? And what are the central philosophical questions those discursive responsibilities raise? A good theory of ethical meaning ought to explain how and why many of these responsibilities are coherently assigned to a unitary concept and, if there are contexts in which our use of the concept is confused or inapt, it ought to identify the source of the problem. Second, I'll briefly review the comprehensive account of meaning that, I contend, provides the foundation for an understanding of ethical meaning capable of meeting the challenges presented in Part 1. (Following my usage in Chapter 1, I'll refer to my proposal for a comprehensive account of meaning as the "CI account"—short for "comprehensive-interpretive"—and the approach to understanding narrower senses of meaning on its basis the "CI approach.") Finally, I'll apply the general features of

meaning adumbrated in Part 2 to the questions about ethical meaning raised in Part 1, detailing

which features of meaning in everyday and theoretical ethical discourses the CI approach

vindicates (and how), along with proposing and motivating some of the revisions to discursive

practice taking that approach involves.

§1 Meaning in Ethical Discourse

A good philosophical account of ethical meaning should be faithful to the way we employ the

concept, in everyday discourse, to frame ethical issues. This faithfulness criterion, however,

doesn't require the supposition that there is a single, rigorous, coherent set of necessary and

sufficient conditions implicit in our "meaning" talk, waiting to be uncovered: as with most words

with even moderate currency in our talk and thought, this isn't the case. Some degree of well-

motivated stipulation is inevitable. Nor should the faithfulness criterion forbid us from proposing

revisions to some of what *is* largely implicit in our "meaning" talk. But, if an account of ethical

meaning is going to have any bearing on the concerns we articulate in its name, it needs to keep

faith with a reasonable quorum of the central features implicit in our everyday discourse. Lisa

Stroud's diagnosis of her parents' conflict, quoted at the outset, illustrates several of these.

First, we generally speak of meaning in life as something *valuable* or desirable: something,

all things equal, worth wanting; something whose absence is cause for regret or discontent.

Second, meaning in life is *elusive*—the kind of thing that goes missing easily enough. People often

worry that it's lost, or was never there; finding or making it is a significant challenge; possessing it

is an accomplishment. Third, ethical meaning has a kind of *depth*, in the sense of depth that

stands in contrast to more "superficial" values or desires, like fame, material wealth, "base"

pleasures, or other outward markers of status or success. In William James's words, considering

meaning in life involves turning our attention away from "the buzzing and jigging and vibration

of small interests and excitements that form the tissue of our ordinary consciousness," toward

"the profounder bass-note of life."[6]

In addition to its role in everyday ethical discourse, meaning is also a tool and, somewhat

less frequently, a target in ethical theory. Charles Taylor, for example, contends that all cultures

are entitled to a presumption of equal worth if they "have provided a horizon of meaning for large

numbers of human beings, of diverse characters and temperaments, over a long period of time."[7]

In a similar vein, Kwame Anthony Appiah invokes meaning to explain the ethical importance of

identity—a person's sense of who she is. Neither moral considerations nor natural constraints are

sufficient, he thinks, to give our lives a definite shape: "even when we have taken these things into

account, we know that each human life starts out with many possibilities…everybody has, or

should have, a variety of decisions to make in shaping a life."[8] We need, then, some additional

source of reasons to make these choices meaningful: "One thing identity provides is another

source of value, one that helps us make our way among those options. To adopt an identity, to

make it mine, is to see it as structuring my way through life."[9] Should we lose touch with this

source of value, even if by an excess rather than a lack of power, life, Appiah contends, would

become "meaningless."[10]

The use to which Appiah puts meaning in the discussion above accords well with the

idea, commonly espoused in the literature that makes ethical meaning its primary target of

inquiry, that "meaning" names a *distinctive* realm of value, reducible to neither duty nor self-interest, narrowly construed. Metz, for example, objects to purpose-centered conceptions of the meaning of life because "they fail to differentiate the meaning of life from other normative categories," and he takes one of the philosophical questions definitive of a conceptual investigation of life's meaning to be: "What should an agent strive for besides obtaining happiness and fulfilling obligations?"[11] Wolf finds in "meaning" a vocabulary for articulating a richer conception of practical reason than those, prevalent in contemporary philosophy, that regard it as exhausted by considerations of prudence and morality: in her influential treatment, she casts meaning as "a third sort of value a life can possess…not reducible to or subsumable under either happiness or morality."[12] Robert Nozick's work represents an interesting variation on the distinctiveness theme. In his parlance, meaning and value are two distinct facets of "worth," so that it makes sense to ask if there might "be a conflict between the meaning and value of a person's life," or to wonder if there's "any meaning even to a valuable life."[13]

Appiah's discussion evokes a second theme that has recently preoccupied philosophers of ethical meaning—the question of whether meaning is, in some sense, "subjective" or "objective".[14] On one hand, the sort of meaning that Appiah claims identities generate for their bearers is, in an important way, particular to *them*: for Mr. Stevens (the narrator of Kazuo Ishiguro's *Remains of the Day*), for example, becoming a great butler matters not only because of the objective importance of the service he provides, but because this was his father's profession[15] On the other hand, as we saw above, identity generates meaning in ethical contexts by providing the kind of structure of external constraint within which rational choice is possible. Is the meaning that Mr.

Stevens' professional achievements confer on his life diminished by the fact that his master, whose judgment and influence Stevens reveres, turns out to have been a Nazi dupe? And to what extent does Stevens' own attitude toward that question determine its answer?

Consider a strangely ghoulish complementary case. Addressing a group of college students at the end of the 19[th] Century on the question "Is Life Worth Living?", William James suggested that the reflective life, so susceptible to the "skepticism and unreality that too much grubbing in the abstract roots of things will breed," leads us "to the edge of a slope, at the bottom of which lie pessimism and the nightmare or suicidal view of life."[1] In his attempt to provide whatever remedy reflection can offer for the diseases it breeds, James recalls a formative experience from his medical education:

> Consider a poor dog whom they are vivisecting in a laboratory. He lies strapped on a board and shrieking at his executioners, and to his own dark consciousness is literally in a sort of hell. He cannot see a single redeeming ray in the whole business; and yet all these diabolical-seeming events are often controlled by human intentions with which, if his poor benighted mind could only be made to catch a glimpse of them, all that is heroic in him would religiously acquiesce. Healing truth, relief to future sufferings of beast and man, are to be bought by them. It may be genuinely a process of redemption. Lying on his back on the board there he may be performing a function incalculably higher than any that prosperous canine life admits of; and yet, of the whole performance, this function is the one portion that must remain absolutely beyond his ken.[2]

Should we be consoled by the suggestion that we might share this dog's plight? Do the good consequences of his sufferings make them, or the extinguishing life they harrow, meaningful? Or must meaning be appreciated to be had? And would either of these versions of meaning, should they be realized, make life worth living?

---

[1] James, "Is Life Worth Living?," 39.
[2] Ibid., 58.

We can find an illustrative theoretical contrast on the objectivity/subjectivity issue in Susan Wolf's and Richard Taylor's accounts of ethical meaning. Taylor proposes, as a thought experiment, that we reimagine the Sisyphus myth, so that, in addition to being consigned to roll his boulder up a hill, again and again, for all eternity, Sisyphus is also blessed (or, depending on your appraisal of the experiment, cursed) with an injection that gives him an insatiable desire to do just this, and causes him to find the activity endlessly engaging.[16] If we're tempted to dismiss the potential meaningfulness of this scenario, Taylor argues, we should think more carefully. If we imagine that our own lives possessed or could possess some meaning significantly different in kind from what Sisyphus, fulfilled, finds in his rock rolling, we're deceived: the only difference is that "whereas Sisyphus himself returns to push the stone up again, we leave this to our children." Properly appreciated, though, this is no cause for despair. As long as we realize plans, however ephemeral, that engage our interests, no further questions need be asked: "The day [is] sufficient to itself, and so [is] the life…The meaning of life is from within us, it is not bestowed from without, and it far exceeds in both its beauty and permanence any heaven of which men have ever dreamed."[17]

Though Wolf joins Taylor in recognizing a subjective component of ethical meaning—a meaningful life, she thinks, is one its subject must "find fulfilling"—she appraises Taylor's "Sisyphus Fulfilled" scenario differently. On Wolf's view, ethical meaning has an objective as well as subjective component: an activity that inspires our passion and engages our interest only makes our lives meaningful if it *also* involves "participating in or contributing to something whose value is *independent* of oneself."[18] No matter how engaging, fulfilling, and worthwhile an agent might

find a given activity, if its object is not "worthy of love"—if it possesses no value whose source can be rightly described as "outside the agent"—then it doesn't confer meaning on the agent's life. Perhaps, having received his injection, Fulfilled Sisyphus is "better off" (though Wolf questions even this, pointing out that the badness of being deluded might be worse than the purely hedonistic benefits the injection confers), but, according to Wolf, his life has been made no more meaningful. Neither are the lives of more familiar characters meaningful, like the "personal assistant to a Hollywood star…seduced by the glitter and fame that surround her into thinking that catering to her employer's every whim is a matter of national significance," or the ambitious young attorney who sees "his ardent defense of an unscrupulous corporate client as a noble expression of justice in action." Their activities, however subjectively engaging, fail to meet Wolf's standard of objective worth.[19]

This brief survey of the way ethical meaning has been used and theorized in recent philosophical work adds to the list of questions raised by our discussion of meaning in ordinary ethical discourse. We want an account of meaning that, in addition to accounting for the desirability, elusiveness, and depth the concept evokes in everyday ethical discourse, helps us to clarify and respond to questions like these: Is the *way* that ethical meaning is valuable distinct from other sorts of value, like those that rationalize and/or motivate moral and egoistic activity? Or, as Nozick thought, is meaning distinct even from value itself—a second component of a more general category of worth? Insofar as meaning is distinctive, what is its character? How does it relate to other normative categories relevant to ethics, like well-being, self-interest, moral duty, and pleasure? Can it be possessed simply on the basis of certain subjective states, or does it require

some objective component? By what process can we come to possess meaning in our lives? Do we

see it, *find* it, or *make* it?

I claimed, above, that reflecting on meaning as it functions *across* discursive contexts—

not only in its ethical dimension, but in, for example, its linguistic or literary or natural modes as

well—provides insight into the questions just posed. Specifically, an account of meaning couched

in terms of its distinctive role in the practice of interpretation succeeds in comprehending

meaning across this broad range of discursive contexts, and a careful examination of the nature of

interpretation as a practice discloses features of meaning that, applied to ethical contexts, provide

a strong foundation for meeting the challenges developed in this section. I briefly review this

comprehensive-interpretive (CI) account of meaning (presented in more detail in chapter 1)

below.

§2 The Comprehensive-Interpretive Account of Meaning

To review: according to the CI account, meaning is united across its varieties by the distinctive

role it plays in the general form of interpretation. Interpretation is a functionally structured,

mentally mediated activity, undertaken by an organism situated within an evaluative context,

subject to norms of fruitfulness and validity.

(1) **Functionally Structured:** interpretation is *functionally structured* in the bare sense that it

can be usefully described as taking an input and issuing an output. I'll refer to the input of

an interpretation as its *object* and the output (if the interpretation goes well) as a *meaning*.

Simplifying a particularly famous case, Luis and Walter Alvarez interpreted the high

73

concentration of iridium in the K-Pg layer of the geological record (object of

interpretation) to mean that a massive asteroid struck Earth ~65 million years ago

(meaning).

(2) **Mental Mediation:** interpretation is *mentally mediated*, rather than automatic. Items are

objects of interpretation rather than, say, objects of perception when they are, in Clifford

Geertz's formulation, "on their surface enigmatical": their import is not immediately

apparent.[20] Interpreters attend to objects of interpretation *as media*, possessed of features

independent of the meanings they disclose, in a way that perceivers (we now think)

generally do not attend to the "sense data" posited by certain indirect theories of

perception. Interpretation stereotypically involves a kind of conscious effort that basic

sensory perception, for example, doesn't.

(3) **Undertaken by Organisms Situated in an Evaluative Context:** interpretation is

undertaken by agents for the purpose of guiding action, affection, or belief. The cognitive,

affective, or practical goals that rationalize an act of interpretation, along with the

complex of constraints and potentialities presented by the interpreter's environment,

constitute the interpretation's *evaluative context*. For example: both the positions of trees

in the forest and the footprints on the forest floor carry information, in Fred Dretske's

sense, but only the latter are particularly meaningful in the evaluative context of a hunter

tracking prey.

(4) **Subject to Norms of Validity and Fruitfulness:** a good interpretation should be both

valid and fruitful. Interpretation should be *fruitful* in the sense that its output ought to be

relevant with respect to the evaluative context in which it's undertaken. (We rightly call certain objects or states of affairs "meaningless" when they afford no insight into the task at hand.) Second, regardless of their fruitfulness, interpretations should be *valid*: their objects must relate to their outputs in a way that licenses concluding the latter on the basis of the former. If the iridium in the K-Pg layer were not causally related to an asteroid strike, for example, then it wouldn't be rightly interpreted as meaning that the asteroid struck—even if the asteroid did, in fact, strike just when the interpreters conclude it did.

(5) **Clarification:** An object has the meanings it does relative to potential rather than actual acts of interpretation. Just as a substance can be nutritious for an organism whether or not it's actually ingested and metabolized, an object can be meaningful with respect to a potential interpreter in an evaluative context whether or not it is actually interpreted.[21]

§3 The CI Account Applied

In the first section of this essay, I posed several questions about meaning in ethical contexts. What explains the characteristic *depth* and *elusiveness* of ethical meaning? What, if anything, is *distinctive* about the sort of value that meaningfulness confers on life? How does meaning relate to other sources of motivation, like self-interest and moral duty? Is ethical meaning subjective or objective, or does it involve components of both? What's happened when we lose the sense that our lives are meaningful? And how should we describe the process by which we might come to

experience our lives as meaningful?—do we simply *see* the meaning in our lives, or must we find

it, or make it?

Having summarized, in §2, a proposal for understanding meaning *across* its contexts of

use, I want now to consider what that comprehensive account implies about how we should

approach these questions about meaning in its ethical dimension. First, though, I'd like to develop

a terminological distinction that the CI account puts us in position to make.

§3.1 Meaning in Life and Meaning of Life

Some philosophers use the phrases "the meaning of life" and "meaning in life" roughly

interchangeably, but for others, the distinction is important. Susan Wolf, for example, favors the

"meaning in" construction as a way of distancing her work from the relatively grandiose questions

stereotypically evoked by the "meaning of" construction. As Stephen Macedo puts it, Wolf's

subject "is not the question of the ultimate meaning of human life: whether humans are part of a

larger narrative or higher purpose," nor does she aim "to fend off existential dread." Instead she

hopes to explain what we seek when our aim is find meaning *in* life: "Is it distinctive, or deducible

to other aims and conceptions? Is it a helpful category for thinking about good lives that are

worth living?"[22] The structural machinery introduced by the CI account of meaning puts us in a

position to make this distinction more precise, in a way that will be useful in the discussion that

follows.

According to the CI account, ethical meaning arises when we engage in interpretive

practice whose objects are distinctively *ethical*—related, essentially, to the project of living well. In

the simplest case, ethical meaning arises when we treat our own lives as objects of interpretation (situated, therefore, within some larger evaluative context). These cases are best captured by the "meaning *of*" construction, which accords with Macedo's usage, above. However, as we've already noticed, ethical meaning seems to arise on other occasions as well: it's often natural to characterize as ethical the meaning we ascribe to important relationships, activities, or mementoes, for example.[23] In these cases, our lives themselves aren't the *objects* of interpretation. Instead, they provide the evaluative context within which our interpretation of a given object takes place: we interpret a relationship, activity, or memento according to how it does or might contribute to, or hinder, our effort to live well.[24] For cases like these, we should employ the "meaning *in*" construction. Interpretations of meaning in a life, of course, will often be made in light of a certain conception of the meaning *of* that life, but this needn't be the case: a life can provide the evaluative context for the interpretation of other objects without being itself regarded as meaningful, in the "meaning of" sense. (I'll have more to say on this topic below, in my discussion of Nozick's worry that ethical meaning generates an infinite regress.) I'll use the phrase "ethical meaning" as general, capturing both meaning of life and meaning in life.

With these terminological preliminaries complete, let's turn our attention to the questions I posed in this section's introduction.

§3.2 Depth and Elusiveness

In the discussion, in §1, of Nina Simone's conflict with her husband over the direction her career should take, we noticed that ethical meaning is often regarded as *elusive*—something easily lost,

whose secure possession is a genuine achievement—and also characterized by a particular sort of *depth*, in contrast to more stereotypically superficial desires or ambitions, like fame, wealth, and base pleasures. The CI account enables us to explain these features of ethical meaning in tandem, by the fact that meaning, across contexts, is mentally mediated.

Meaning, that is, arises when interpretation is required to appreciate an object's import, or value—when its worth, to speak metaphorically, lies somewhere beneath the surface. Stereotypically superficial goods like wealth, social status, and pleasure wear their desirability on their sleeves: absent some reflective effort, we're often motivated to pursue them thoughtlessly, automatically. Nor does their appreciation or enjoyment seem to require significant interpretive effort: pleasure, for example, needn't be regarded as being situated within some larger evaluative context to be desired, nor even be consciously represented to be enjoyed. In fact, when we *do* apply critical reflection to our superficial desires and satisfactions—when we represent them to ourselves consciously, and ask of our pleasure, or wealth, or social status, what real *meaning* it has—we risk becoming alienated from their apparent value, and they may cease, at least temporarily, to motivate or satisfy us.[25]

Treating something as an object of ethical interpretation, then, involves an expectation that it should possess depth. When we ask, for example, if our high social status or wealth is really *meaningful*, we demand that it justify itself by something more than its superficial attractions and satisfactions; we want to recognize it as having some positive import in a larger evaluative context. Does the wealth I've earned reflect the positive contribution I've made to my society, or has it come at the expense of others? And have I, or will I, use it in a way that matters?

78

Is my high social status a reflection of who I really am—or am I just a fraud? Do I really like or respect the people whose admiration I command? And do they *really* like me, or are they just faking, too?

It should be clear enough that the characteristic elusiveness of ethical meaning is intimately related to its depth. Interpretation regards surfaces *as* surfaces, whose objects are situated in larger evaluative contexts, and possess internal and relational properties that require effort to uncover and comprehend. When, psychologically, we distinguish between appearance and reality, medium and meaning, a psychological precariousness is activated that isn't present in immediate, automatic enjoyments. It's a precariousness with multiple dimensions, which I'll discuss in more detail below, when I address the questions of how we can come to possess and to lose ethical meaning.

Before leaving the topics of depth and elusiveness, however, I want to note a revisionary implication of the way the CI account explains these features of ethical meaning. Understood in contrast to the kind of depth ethical meaning has, *superficiality* isn't necessarily pejorative— despite the fact that, in everyday ethical contexts, it's hard to use the term without pejorative force. A good might motivate us automatically, requiring no interpretive effort for its enjoyment, yet still be of great ethical worth. Sheer pleasure, for example, might be a significant component of the good life, regardless of the way it's produced, or its role in some broader evaluative context. If deep ambitions and satisfactions really are of greater ethical value than superficial ones, we should produce a substantive explanation for that fact. From the CI perspective, that value asymmetry isn't built in to the meanings of the terms.[26]

§3.3 Desirability and Distinctiveness

As much of the language above indicates (for instance "goods" and "apparent goods"), I've

proceeded, to this point, under the assumption that ethical meaning is something desirable and

valuable—an assumption largely reflected in our everyday discourse. We also noticed that many

philosophers have regarded the kind of value meaning confers on life to be, in some important

sense, *distinctive*: in most cases, distinguished from reasons of self-interest, on one side, and

moral duty, on the other—but, in at least one interesting case, distinct from value altogether,

standing alongside it as one of two dimensions of ethical "worth". The CI account of meaning

suggests that we revise both assumptions—value/desirability and distinctiveness. The revisions,

elaborated below, better serve the way we use "meaning" in everyday discourse, preserving what's

intuitively important in both assumptions, but forestalling confusions that arise from the

oversimplifications that, unrevised, they effect.

§3.3.1 Distinctiveness

Let's begin by taking a closer look at the versions of distinctiveness favored by Susan Wolf and

Robert Nozick. According to Wolf, meaning constitutes its own "category or dimension of value"

because it arises when we respond (with some measure of success) to *reasons of love*, rather than

reasons of self-interest or moral duty. "When I visit my brother in the hospital," she writes,

> or help my friend move, or stay up all night sewing my daughter a Halloween costume, I
> act neither for egoistic reasons nor for moral ones. I do not believe that it is better *for me*
> that I spend a depressing hour in a drab, cramped room, seeing my brother irritable and
> in pain, that I risk back injury trying to get my friend's sofa safely down two flights of

stairs, or that I forego hours of much-wanted sleep to make sure that the wings will stand out at a good angle from the butterfly costume my daughter wants to wear in the next day's parade. But neither do I believe myself duty-bound to perform these acts, or fool myself into thinking that by doing them I do what will be best for the world. I act neither out of self-interest nor out of duty or any other sort of impersonal or impartial reason. Rather, I act out of love.[27]

A proper understanding of meaning, Wolf thinks, ameliorates the distortion generated by dualist conceptions of practical reason, which require us to recast (often unconvincingly) concerns like the ones illustrated above as either self-interested or morally required—or, worse, to abandon them altogether.

While Wolf attempts to rescue meaning from the sort of grandiose existential pretensions with which it's often popularly associated, Nozick embraces its traditional role in articulating those concerns. "From the point of view of the universe," he asks, "is there any importance to ethical behavior, is there any meaning even to a valuable life?"[28] This question makes at least preliminary sense in Nozick's idiom, unlike in Wolf's, because Nozick treats meaning and value as independent qualities a life can have, rather than regarding meaning as a distinctive kind of ethical value. A second contrast between Nozick's approach and Wolf's—one that constitutes a kinship between Nozick's account and the one I advocate—is that Nozick thinks that ethical meaning ought to be understood in a way that makes it continuous with the way other sorts of entities mean: one species of a larger genus.

That latter consideration provides the basis for Nozick's distinction between the meaning and value of life. To ask of anything what its meaning is, Nozick claims, "is to ask how it is connected, perhaps in specified ways, to other things."[29] Specifically—in concurrence with the CI

account—Nozick contends that recognizing something's meaning involves seeing its importance within a wider context.[30] Applying this general idea to ethical discourse, he stipulates that "the value of a person's life attaches to it within its limits, while the meaning of his life attaches to it as centered in the wider value context beyond its limits." A life's meaning and its value stand alongside one another as the two independent components of its overall *worth*: "Value, we now see, was only part of the picture—meaning is the other."[31]

Though each of these conceptions of the distinctiveness of ethical meaning has its attractions, neither is fully satisfying. Wolf is right to draw our attention to reasons of love, and to urge us to resist their assimilation to reasons of self-interest or moral obligation. But restricting the domain of ethical meaning to activities that arise from these reasons alone unnecessarily narrows the concept in a way that fails to do justice to the role it plays in ethical discourse, and creates an unmotivated discontinuity between ethical meaning and meaning in general. When we interpret our lives as possessing import with respect to a larger evaluative context, for example, our moral projects frequently take on a central significance.[32] Our most meaningful relationships, moreover, matter not only because of the love they (sometimes) sustain, but also because of the duties they engender. Abraham Lincoln's life illustrates both points. It would be strange to deny that the fight against slavery in which Lincoln ultimately found his calling was a central source of meaning in his life, or that his engagement in that effort was motivated by a profound sense of moral duty—a conviction that, in acting as he did, he did what would be best for the world.[33] Lincoln's relationship with his wife, Mary, highlights a narrowness in the other direction. He often conceived of their marriage in terms of duty rather than love, and, though the quality and

content of the relationship's meaning might certainly have been contingent on the balance it ultimately struck between those two forces, the bare fact of its meaningfulness was not. Even if Lincoln's marriage were wholly loveless, it would be implausible to deny that it played an ethically meaningful role in his life. In short: our moral projects are too central to the meaning we find in life, and the value of our meaningful relationships too thoroughly bound up in the duties they involve to make the segregation between meaning, happiness, and morality that Wolf proposes a sensible rendering of the role meaning plays in ethical discourse.

Nozick's proposal, as we've noticed, shares with the CI approach the general idea that something's meaning is constituted, at least in part, by its import in a wider context. But the subsequent suggestion that we regard meaning, *alongside* value, as a distinctive category of "worth" creates unnecessary confusions. For instance, Nozick asks "How meaningful is it to achieve value?"—as if, absent confirmation that achieving value is meaningful, we might be rightly unmotivated to achieve value. (He worries that, "unless it can be shown how our lives and existence can have meaning…ethics and value themselves will seem *meaningless*.")[34] But the CI account suggests, sensibly, that worries like these are misplaced. Meaning isn't the ultimate measure of mattering—some independent dimension on which value must register to generate reasons, or to count, motivationally. Instead, "meaning" is vocabulary for a certain way of recognizing or appreciating value: without the presumption of a shared *evaluative* context, the demand to know something's "meaning" is incomprehensible.

Relatedly, Nozick's way of conceiving ethical meaning's distinctiveness creates an unnecessarily demanding standard for its possibility (and, complementarily, a too radical account

of meaning's elusiveness). Almost any claim to ethical meaning, Nozick argues, can be undercut simply by adopting an ever-wider perspective, from which the purported meaning can be called into question: "However widely we connect and link, however far our web of meaningfulness extends, we can imagine drawing a boundary around all that, standing outside looking at the totality of it, and asking 'but what is the meaning of that, what does that mean?'"[35] Meaning, in other words, is like importance: "to be important for something which itself is unimportant is for these purposes to be unimportant"; similarly, having meaning within a context that cannot itself be shown to be meaningful is to have no meaning at all: "When the concern is the meaning of our life or existence…we want meaning all the way down."[36] Given this problem of infinite regress, the only scenario that Nozick can imagine in which ethical meaning is possible is one that involves connecting with a limitless being, like God. (Limitlessness, Nozick thinks, makes it possible for something to be its own meaning—though he doesn't explain how or why the issue of meaning should arise at all with respect to a limitless being.)[37]

While the CI approach doesn't deny the possibility that ethical meaning is, ultimately, achievable only by way of connection with some infinite being, it *does* deny that we can arrive at this conclusion simply by reflecting on the concept of meaning, as Nozick does. The regress that arises from regarding meaning as independent from value, satisfying only if it goes "all the way down" (that is, has a source that is itself meaningful), isn't generated by the CI approach, which requires only that meaning be understood in some value-laden, rather than some necessarily *meaningful* context. If the end or ends in whose context a life has meaning have intrinsic value, that meaning isn't undermined if those ends themselves have no *further* meaning with respect to

84

some *wider* context.[38] In many cases, from the CI perspective, the sort of endlessly iterative

demands for meaning that Nozick describes are pathological rather than genuinely challenging:

the "problems" they pose require dissolution rather than solution. (Nonetheless, the CI approach

offers a satisfying explanation for *why* such demands for meaning at least appear to be endlessly

iterable—one I'll give when I discuss the particular ways in which meaning can be lost, below.)

§3.3.2 Desirability

My focus, so far, has been on the supposed distinctiveness of the value of ethical meaning, but the

first of the two questions I raised at the outset of this section—the question of whether ethical

meaning is necessarily valuable or desirable at all—remains. My analysis, in §1, of the role of

meaning in ordinary discourse, might tempt us to build the desirability of ethical meaning right

into the concept: implicit in much of our ethical discourse, I noted, is the idea that meaning is

something worth wanting, something whose absence is cause for regret or discontent. And this

apparent feature of everyday discourse is largely reflected in the theoretical literature on ethical

meaning: Metz, for instance, limits his survey of philosophical work on the meaning of life with

"writings that take a meaningful life to be one desirable facet of a person's existence."[39] In this

case, however, the broader perspective afforded by treating ethical meaning as one mode of a

larger phenomenon is salutary. By introducing considerations beyond meaning's role in ordinary

ethical discourse (the requirement, that is, that ethical meaning be understood in accordance with

the features of meaning in general), the CI approach highlights aspects of ethical meaning which,

though, on reflection, *do* play an important role in ethical discourse, are easily overlooked when we make ordinary discourse our only guide.

In particular, the CI approach recommends a qualified revision to the assumption of desirability: that there is an important sense in which ethical meaning is not *necessarily* something good or desirable or valuable. Remember that, according to the CI account, an object's meanings are its *interpretive* affordances. The import uncovered by interpretation, though its appreciation should be *useful* to a doxastic, affective, or practical agent, needn't be all-things-considered positive with respect to the evaluative context in which it arises. Dark clouds, for example, may be meaningful but unwelcome to a group of people enjoying a picnic.

In the ethical context, an important consequence of this observation is that the meaningful aspects of a person's life (meaning in) won't all be sources of satisfaction, accomplishment, or fulfillment: a failing marriage, from the CI perspective, may be as meaningful as a successful one; a relationship with a cold or estranged parent as meaningful as a relationship with a loving parent. Similarly, if lives themselves take on meaning with respect to a broader evaluative context (meaning of), this meaning may not be positive. It may be that leading a meaningless life necessarily counts against a person's having lived well (though that's not a claim I commit to here); but, from the CI perspective, having led a meaningful life doesn't *necessarily* count in favor of a person's having lived well.

Accommodating this observation requires less revision to everyday discourse than might initially appear, and makes sense of meaning-related phenomena that might otherwise be confused or overlooked. While it's true that we often speak of meaning and meaningfulness as if

these are desirable or valuable qualities for a life to possess, the presumption that the meaning

desired or appreciated is positive might be better understood as a pragmatic rather than semantic

phenomenon. Here, "meaning" *is* like "importance" or "mattering": the fact that we comfortably

articulate our hopes to achieve these without stipulating that it be "in a good way" doesn't prevent

us from recognizing that not everything that matters or has importance is for the better rather

than the worse. Our reluctance to call the lives of people who have perpetrated great evil

"meaningless" reflects as much; and our hesitation to deny that their lives possess a kind of

meaning doesn't involve ethical admiration, or indicate to us that they led lives worth living. It's

worth noticing, in this regard, that the existential suffering and sense of alienation that have, in

some cases (like Lincoln's), eventually provided the occasion for moral heroism or artistic

brilliance, have, in others, been a source of extraordinary evil. Humans don't always respond

constructively to the painful, desperate feeling that their lives lack meaning, and the fact that

harmful behaviors can appear to provide relief from the affliction when positive outlets are

unavailable makes sense in light of the fact that ethical meaning does not entail ethical goodness.[40]

§3.4 Subjectivity & Objectivity

In §1, we met two characters—Sisyphus Fulfilled (who, by the grace of an injection, experienced

as deeply meaningful a task that, from an external perspective, seems paradigmatically pointless

and drudgerous), and William James's Dog (whose terrible sufferings bring about valuable

consequences he can't understand)—that, respectively, raise questions about the role of

subjectivity and objectivity in ethical meaning. Is the deep sense of fulfillment that Sisyphus

derives from his endless stone rolling, though a wholly subjective state, enough to confer meaning on his existence? Conversely, though James's Dog experiences his role in human medical research as nothing more than being in "a kind of hell," is his excruciating suffering somehow redeemed or made meaningful by the objective goods it's involved in bringing about?

These cases frame the issue of the subjectivity or objectivity of ethical meaning in terms of the importance (or irrelevance) of a certain kind of subjective experience. We can call a theory of ethical meaning *strongly subjective* if it regards a life or activity as meaningful to the extent that it involves or gives rise to some relevant subjective state—say fulfillment, rapt interest, or love (understanding each of these as purely subjective states). A theory of ethical meaning will be *strongly objective*, conversely, if it regards the subjective states of the liver or the actor as irrelevant to the meaning of the life or activity in question. Theories that treat both subjective and objective components as important, so that a life or activity must give rise to relevant subjective states *and* possess certain objective features, can called *hybrid* accounts, and be characterized as both *weakly* subjective and *weakly* objective.

Neither Sisyphus Fulfilled (whose existence, despite its apparent pointlessness from an external perspective, might be regarded as meaningful according to a strongly subjective account) nor William James's Dog (whose suffering, on a strongly objective account, might be claimed to confer meaning on the dog's life in spite of his total obliviousness and indifference to the valuable consequences that might issue from it) make either extreme position look intuitively attractive. Wolf, for example, takes our intuitive reluctance to suppose that Sisyphus's paradigmatically meaningless existence can be made meaningful merely by altering his subjective states to

recommend a hybrid position like hers, according to which meaning arises "when subjective

attraction meets objective attractiveness."[41] Even when subjectively fulfilling, Sisyphus's stone-

rolling isn't meaningful because the sense of fulfillment it gives him isn't *merited*: his activity has

no "objective" value—no value independent of the positive feelings it gives him—and doesn't,

therefore, connect him with anything larger than himself. The suffering of James's Dog, by

contrast, isn't a source of meaning in his life because, regardless of the objectively valuable

consequences it might have, it neither engages his interest nor provides him any sense of

fulfillment. (These absences disqualify independently valuable activities from counting as sources

of meaning even when the agents who undertake them *do* recognize their value. An "alienated

housewife," Wolf says, may know that her efforts have valuable consequences, but if she

nonetheless finds the relevant activities unfulfilling, they won't count as a source of ethical

meaning.)[42]

The CI approach validates Wolf's rejection of strongly subjective approaches to ethical

meaning, and not merely because they give potentially counterintuitive interpretations of certain

scenarios (like Sisyphus Fulfilled). Recall, from §2, that interpretation has criteria of quality: in

particular, a good interpretation (one that uncovers a genuine *meaning* of its object) must be

valid—licensed by some relevant sort of relationship between an object of interpretation and a

purported meaning. The fact that interpretations (unlike, say, instances of inspiration) can be

invalid implies that an object's being interpreted as possessing a certain meaning does not, on its

own, entail that the object has that meaning: *being meaningful* and *being experienced as*

meaningful are distinct phenomena, and, generally, each can occur in the absence of the other. If,

on the other hand, the strong subjectivist wants to claim that activities and other objects are ethically meaningful just because they issue in some positive subjective attitude that itself involves no reference to the activity's meaningfulness (so that the pleasure or sense of fulfillment one gets from the object or activity is just brute, rather than an instance of experiencing the object or activity *as meaningful*, or as valuable for any reason other than the pleasure it provides), then, given the considerations raised in §3.2, we're left wondering why the concept "meaning" applies at all. Activities or objects that are brutely pleasurable don't require interpretation to motivate or gratify: they don't, therefore, fall within the purview of "meaning" understood comprehensively, nor do they provide us with any clear reason to employ that concept in ethical discourse.

The same considerations that lead us to join Wolf in rejecting strongly subjective conceptions of ethical meaning, however, should lead us to question the role she gives subjective experience in her account. Why should a life or activity be required to give rise to a particular sort of positive subjective experience (Wolf's preferred candidate is "fulfillment") in order to be ethically meaningful? From the CI perspective, no such requirement should be made: an object can have a meaning without that meaning necessarily being appreciated, and, at least on the conceptual level, a life or activity can have import in a wider evaluative context independent of the fulfillment (or any other positive subjective experience) it might provide to the liver or actor. The subjective aspect of Wolf's account, then, appears (partially, at least) to confuse experiencing a life or an activity *as* meaningful with that life or activity *being* meaningful.

Should we regard this consequence as a drawback of the CI approach to ethical meaning? Philosophers who have introspected to arrive at a pretheoretical concept of ethical meaning

90

frequently note (or assume) that meaningless existence is characterized by boredom, alienation, and disengagement.[43] If subjective states like these are sufficient to undercut ethical meaning, then it stands to reason that their rough opposites (engrossed interest, fulfillment, pride) should be required for meaning. Though the CI approach denies that requirement, it can explain the characteristic *association* between these subjective states and meaning/meaninglessness. Instead of being *causes* or *constituents* of meaningfulness or meaninglessness, the CI approach regards the associated subjective states as *symptoms*: loss of interest, boredom, and alienation are natural responses to activities we come to regard as meaningless; complementarily, activities we regard as meaningful naturally inspire our interest, and, when we engage in them successfully, give us good reason to feel fulfilled or proud. In other words, we generally find certain activities and relationships fulfilling *because* they're meaningful; they aren't meaningful because we find them fulfilling.

In fact, we should prefer the symptom understanding of associated subjective responses to the causal/constituent understanding for reasons independent of its accordance with a comprehensive account of meaning. It isn't, after all, uncommon for people whose work is generally regarded as deeply meaningful to have had ambivalent or even wholly alienated feelings about that work. Nina Simone, for example, became disillusioned with the efforts that, earlier in her life, had seemed so meaningful: "There aren't any civil rights," she told an interviewer who, in 1993, invited her to reflect on the role of her music in the Civil Rights Movement; "There is no reason to sing those songs, nothing is happening, there is no Civil Rights Movement— everybody's gone."[44] Kafka famously asked that all his work be destroyed at his death, and burned

the large majority of his work himself.[45] In neither case should we regard the meaningfulness of the work or lives of these artists as fluctuating with their feelings of fulfillment and alienation. In fact, it's hard to imagine either artist's life and work having the meanings they did *without* the feelings of alienation and discontent that sometimes haunted them: intellectual alienation, as we noticed above, isn't always a pathology; it's no accident that work of extraordinary imagination and expressive power often comes from people who find fulfillment elusive and alienation familiar.

These considerations relate to a more general and important point: the symptom understanding of the association of meaningfulness with characteristic subjective responses does better justice to depression as a pathology than the causal/constituent understanding does. Depressives find it difficult or impossible to have characteristic positive subjective responses (like interest, fulfillment, and pride) to experiences, activities, and relationships that naturally and justifiably generate those responses in people who aren't afflicted by depression. But it would be wrong to think that this affliction condemns depressed people to meaningless existences. Though depression, especially in extreme cases, might make leading a meaningful existence more difficult (positive attitudes certainly make the work needed to maintain good relationships and productive activities easier to undertake), many depressed people find the courage and resolve to pursue meaningful projects and build meaningful relationships *in spite* of their hardship. In these cases, though depression might deprive its sufferers of the important subjective goods that often accompany meaningful existence, it doesn't render their activities, relationships, or lives meaningless.

What, then, should we think about people like Wolf's alienated housewife (her other examples are an assembly line worker and a conscripted soldier), who recognize the value of their work but remain subjectively disengaged or unfulfilled? In some of these cases, it might be wrong to diagnose their negative subjective responses as symptoms of a sense of meaninglessness: they recognize what the CI account suggests is the meaning of their activities—their value interpreted in light of a larger evaluative context. But it also seems wrong to regard their feelings of alienation as pathological or unmerited, as we might in the case of a person whose capacity for positive affect has been enervated by depression. There are two important points to be made in such cases. First, negative subjective responses like alienation, boredom, and lack of interest aren't *only* merited by the recognition that an activity or existence is meaningless. We might, for example, rightly feel alienated if we're being exploited, or if, unfairly, we've been denied the opportunity to shape our own existence—even if the exploitation or coercion involves engaging in activities we recognize are meaningful. Second, if exploitation or coercion *does* give rise to alienation, the exploited or coerced party may not see the activities they undertake in those conditions as sources of meaning in or meaning of *their* lives, because understanding life as an *ethical* (as opposed to, say, a biological) object might involve emphasizing our autonomous activities over ways that we're coerced by others. The things we're coerced to do, according to this line of thinking, aren't features of our ethical selves: only the actions we take freely are. On this interpretation, however, the subjective response remains a symptom of an objective condition—being coerced or exploited. What's in question is what gets included in relevant object of interpretation (in the case of meaning of) or evaluative context (in the case of meaning in).

93

We might also rely on this line of reasoning to deny that the horrible suffering that James's Dog endures on the operating table confers meaning on his life. Whatever valuable consequences might ultimately issue from the research that causes his agony, they'll have been achieved by coercion and exploitation, rather than being the product of the dog's will. Hence, while they may be relational properties of his biological life, they won't count as features of his life in the distinctively ethical sense. But arguments like this draw our attention to a kind of inappropriateness about posing questions about the ethical meaning of the lives of many nonhuman animals. Presumably, dogs don't engage in any attempt to interpret their lives within some broader evaluative context, nor do they treat their lives as an interpretive frame within which certain objects or experiences take on some special sort of significance.[46] In this sense, it's inapt to talk about the "meaning" of a dog's life. If, however, dogs' lives can sensibly be treated as distinctively ethical objects by interpreters who *do* engage in such practices, then there may indeed be an important sense in which the suffering of James's Dog *does* confer meaning on its life—a kind of meaningfulness not altogether unlike the kind that might provide a mourning parent some degree of solace when she learns that the organs of her dead child were used to save the lives of other children.[47]

Ultimately, then, the CI approach leads us to reject the suggestion that ethical meaning necessarily involves the enjoyment of any particular subjective state; but it does offer an explanation for the strong intuitive association between ethical meaning and certain characteristic subjective states. (The same goes for meaninglessness.) Before moving on, however, we should note that the question of the importance or irrelevance of subjective states isn't the *only* way to

frame the subjective-objective distinction. We might also call a conception of ethical meaning "subjective" to the extent that it allows for *subject relativity*: the possibility that what's ethically meaningful for one person isn't for another.

Recall Mr. Stevens, from *The Remains of the Day*, for whom achievement in the profession of domestic service was especially meaningful in part because Stevens' father was an accomplished butler. In this scenario, at least some of the meaning Stevens derives from being an excellent butler is subject relative—responsive to particular features of Stevens's identity. (Achievement in bultering wouldn't be a similar source of meaning for Stevens's master, Lord Darlington, whose heritage is importantly different.) This sort of subject relativity doesn't entail that meaning is strongly or weakly subjective in the sense defined above, and the CI approach gives us no reason to abjure it. Given that particular lives differ from one another (and, therefore, provide different objects and contexts of interpretation), we shouldn't be surprised if they give rise to different meanings.

§3.5 Meaning Lost & Found

At the beginning of this essay, I suggested that a conceptual account of ethical meaning might only clarify, rather than decisively answer, some of the urgent ethical and metaphysical questions we use "meaning" to pose. Particularly, I cautioned that even a clearer, more coherent understanding of "meaning" in ethical contexts wouldn't necessarily provide us with easy-to-follow instructions for making our lives meaningful, nor protect us infallibly against the possibility of meaninglessness. Nonetheless, better theory should lead to more graceful practice,

95

and the conceptual framework developed above generates important insights into the processes by which we can come to lose and to find meaning in our own lives. I'll conclude with some remarks on that topic.

§3.5.1 Lost

The discussion of subjectivity and objectivity emphasized an important distinction, between *having* and *appreciating* ethical meaning. Just as we can come, mistakenly, to believe that activities, relationships, and experiences are meaningful in ways that they really aren't, so can our lives be meaningful in ways we don't necessarily appreciate. When we discuss the processes by which meaning can be *lost*, then, we need to distinguish between modes of loss rarely disambiguated in everyday discourse. In one case, our lives or activities may lose meaning that they genuinely had—our children grow up and no longer need our care as they once did; a terrible accident robs a concert pianist of her dexterity; the way of life to which we've adapted our skills and talents slowly dies.[48] In the other case, we lose our *sense* of meaningfulness. Activities, relationships, and experiences that once seemed meaningful (or, at least, gave rise to the subjective states that characteristically attend meaningfulness, like engaged interest and fulfillment), come, instead, to seem meaningless; we experience, in varying degrees, despair, alienation, fatigue, emptiness.

Though these two modes of meaning loss are conceptually distinct, they frequently happen in tandem: a loss of meaning in the "having" sense (henceforward, simply "loss of meaning") is often all too keenly felt by the dispossessed (and, hence, lost in the "appreciation"

sense as well)—especially when the loss is sudden. In other cases, however, loss of meaning and

loss of appreciation happen out of rhythm. We might be beset by a sudden sense of

meaninglessness even when our external circumstances don't seem relevantly altered. In her

memoire *Eat, Pray, Love*, for example, Elizabeth Gilbert recalls sobbing on the bathroom floor in

her beautiful suburban home, her husband asleep in the next room, suddenly having noticed that

she'd "reached a state of hopelessness and life-threatening despair":

> We'd only just bought this house a year ago. Hadn't I wanted this nice house? Hadn't I
> loved it? So why was I haunting its halls every night now, howling like Medea? Wasn't I
> proud of all we'd accumulated—the prestigious home in the Hudson Valley, the
> apartment in Manhattan, the friends and the picnics and the parties?...I had actively
> participated in every moment of the creation of this life—so why did I feel like none of it
> resembled me?[49]

Cases like these can be especially terrifying, because they mean either that we've lost our ability to

appreciate what's meaningful, or that what had always seemed meaningful to us really wasn't.

When we suffer a loss of appreciation for no obvious or immediate reason, therefore, it's not only

the meaningfulness of our present and future that seems under threat—the supposed meaning of

past events and satisfactions becomes precarious as well. We're tempted to reinterpret the past in

light of the disillusionment of the present, finding there a record of fraudulence and futility.

Absent a substantive account of the nature of value and the way it relates to human life,

the CI approach to ethical meaning has little more to say about loss of meaning in the "having"

sense than has already been presented. It does, however, offer some additional insight into loss of

appreciation. Specifically, it suggests that appreciation can deteriorate in two ways (or, be absent

for two reasons). In one scenario, an interpreter may confidently inhabit an evaluative

perspective, but see no relevant role for her life to play. George Bailey becomes convinced of something like this in *It's a Wonderful Life*, when he concludes that he's "worth more dead than alive." He doesn't suffer from existential doubts about the general possibility of life's meaning—he would be fulfilled if he could support his children, make his wife happy, and contribute to his community.[50] Instead, he loses faith in his ability to accomplish those things.

In the second scenario, rather than coming to doubt the import of her life from the perspective of a confidently inhabited evaluative context, an interpreter becomes unmoored from any evaluative context at all. Here, the problem is not that she finds her particular life meaningless, but that she cannot access a perspective from which *anything* appears meaningful. She suffers from what Camus identified as the feeling of absurdity: "the divorce between a man and his life, the actor and his setting."[51] As David Foster Wallace described the experience,

> It's like worse than anything... It's worse than any kind of physical injury, or any kind of—it may be what in the old days was called a spiritual crisis or whatever. It's just feeling as though the entire, every axiom of your life turned out to be false, and there was actually nothing, and you were nothing, and it was all a delusion. And you were better than everyone else because you saw that it was a delusion, and yet you were worse because you couldn't function.[52]

When people lose their religious faith, experience a drastic philosophical reorientation, or even achieve, at long last, some cherished goal, they become vulnerable to this second sort of appreciation dissolution.[53] In each case, the evaluative context within which they've previously interpreted meaning suddenly falls away.[54]

§3.5.2 Found

In "Notes on a Balinese Cockfight," Clifford Geertz observed that

98

It is in large part *because* the marginal disutility of loss is so great at the higher levels of betting that to engage in such betting is to lay one's public self, allusively and metaphorically, through the medium of one's cock, on the line. And though to a Benthamite this might seem merely to increase the irrationality of the enterprise that much further, to the Balinese what it mainly increases is the meaningfulness of it all. And as (to follow Weber rather than Bentham) the imposition of meaning on life is the major end and primary condition of human existence, that access of significance more than compensates for the economic costs involved.[55]

Viktor Frankl's professional experience treating suicidal patients and his personal experiences in Auschwitz and Dachau led him to a similar conviction, "that man's main concern is not to gain pleasure or to avoid pain but rather to see a meaning in his life."[56]

These remarks don't just serve to emphasize the importance of the topic of this final section—the process by which lives can come to possess meaning—they also raise, by the subtle difference in the way they're phrased, an important question. Geertz describes the process by which human life comes to possess meaning as an *imposition*—more a contrivance than a discovery. Frankl's vision-based metaphor, by contrast, suggests that the process is better understood in terms of appreciation than in terms of creation. Is the meaning of our lives *out there*, waiting to be uncovered and faithfully fulfilled? Or is it, rather, a product of our own creation—something altogether new that we bring into the world by the strength of our imagination and will?

We can consider questions like these with greater clarity by extending the stipulative conceptual structure I began developing in the previous section, which distinguished between *having* and *appreciating* meaning. Let's add that a person that a person *possesses* ethical meaning when her life has meaning *and* she appreciates that fact. Further, let's understand *appreciation* in

two dimensions: one cognitive, depending on whether we *recognize* the meaning our lives have; the other affective, depending on whether we *enjoy* the meaning of our lives (if, that is, we experience the positive subjective responses, like fulfillment, engaged interest, and pride that are appropriately generated by living a positively meaningful existence).[57]

We can organize our options for responding to the question raised by the contrast between Geertz's and Frankl's terminology by distinguishing between three conceptual metaphors for describing the process of coming to possess ethical meaning—*seeing*, *finding*, and *making*— arranged successively along a dimension that emphasizes *receptivity* at one pole and *activity* at the other. (When we describe the process of coming to possess ethical meaning as *seeing* the meaning in or of our lives, we emphasize the receptive aspect most strongly; when we talk of *making* meaning, we emphasize the active aspect; *finding* meaning picks out a position intermediate between these two.) If the meaning of or in our lives needs only to be seen or found, then it's already, in some important sense, *had*: our challenge is to come to appreciate it. The seeing metaphor suggests that what's necessary to achieve appreciation is the development of the right sort of receptivity. The meaning of your life is right before your eyes, if only you could see it; the change required is one of orientation, attention, or focus. Finding, by contrast, suggests that appreciation requires a kind of *searching* activity. The meaning we seek exists, but it's hidden, or at a great distance. The idea that meaning needs to be found evokes the familiar idea of a quest. Making, finally, casts *having* rather than *appreciating* ethical meaning the primary concern. The challenge is neither to open yourself to life's meaning, nor to discover and fulfill it; instead, you

have to make it for yourself. Where finding requires a *searching* activity, making requires *creative*

activity.

Which of these three models for coming to possess meaning does the CI approach favor?

The natural suggestion is the intermediate position, *finding*. We might first observe that the

finding metaphor does better justice to the requirement that interpretation be *mediated* than does

the seeing metaphor. (In fact, when I described the mediated character of interpretation above, I

used the relative immediacy of direct perception as an illustrative contrast.) Next, we can notice

that the *validity* criterion of quality interpretation mitigates against the creation metaphor. While

interpretation, like searching, may require creative thinking, it isn't a *fundamentally* creative

enterprise. Put in terms of functional structure, the output of an interpretation is constrained by

its input in a way that the output of pure invention is not.

However natural, though, this line of thinking is too simple. We should, instead, resist the

question it attempts to answer. *Seeing*, *finding*, and *making* are *all* useful models for the processes

by which we can come to possess meaningful lives. Instead of ranking them, the CI approach

should help us better understand them. Each model, as we noticed above, treats the challenge of

coming to possess meaning differently—but, depending on her situation, a person might find

herself facing any of these challenges. It's important, then, to understand each strategy, and to be

able to recognize the circumstances in which each is appropriately deployed.

§3.5.2.1 Seeing

Cultivating the right sort of receptivity is important when the fundamental features of a

meaningful existence are already in place, but our ability to appreciate them is inhibited. The

vision metaphor extends nicely to describe the possible sources of that sort of inhibition: we may

be deceived by an illusion, or have our focus misdirected, or be blinded by fear, anger, or anxiety.

I mentioned, above, that George Bailey's malady is at least largely of this kind: he's so

preoccupied, first, by the image of another life he could have had, and, second, by a crippling fear

of failure, that he fails to appreciate the deep meaning his life already has. What he needs to

undertake to achieve that appreciation isn't a quest, but, in the first place, a reorientation of his

evaluative perspective—one that foregrounds the value of his family, friends, and community;

and, in the second case, greater sensitivity to his own value. His guardian angel, Clarence, helps

him accomplish the reorientation by showing him what the world would really be like without

him—using the desperate sense of loss George is already experiencing to awaken him to the value

of what he'd previously taken for granted. George's friends and family do as much as Clarence to

help him accomplish the second task, when, at the movie's conclusion, they forestall the failure

that so preoccupied him by pooling their resources to assist him.

In a real-life example, Richard Feynman gave a moving description of the process in his

advice to a former student, who had written to Feynman in despair, having lost faith in his ability

to make a significant contribution to fundamental physics. "You say you are a nameless man,"

Feynman wrote,

> You are not to your wife and to your child. You will not long remain so to your
> immediate colleagues if you can answer their simple questions when they come into your
> office. You are not nameless to me. Do not remain nameless to yourself—it is too sad a

way to be. Know your place in the world and evaluate yourself fairly, not in terms of the naïve ideas of your own youth, nor in terms of what you erroneously imagine you're teacher's ideals are.[58]

Feynman doesn't promise his student that a great discovery is around the corner; instead, he invites him to shift his perspective, and to appreciate the value his life already has.

§3.5.2.2 Finding

In other cases, however, becoming properly receptive isn't enough, and a person needs to conduct an active search in order to come to appreciate the meaning of her life. In these cases, where *finding* is necessary, a person's life has meaning, but it's hidden, distant, or otherwise alien. On the awful night Elizabeth Gilbert spent crying uncontrollably on her bathroom floor, she hadn't wholly lost her sense of self, or her faith in the meaning of her existence. But she felt trapped in a lived structure in which she could no longer recognize herself. Hard as she tried, no effort at reorientation could reconcile Gilbert to the suburban family life she had mistakenly committed herself to. She needed to make radical changes to the structure of her everyday life before she could clearly recognize and enjoy its true meaning. In her case, these active changes included the end of her marriage, the sale of her house, a new spiritual practice, and a year of travel and writing.

Gilbert's progress from alienation to possession, nonetheless, fits the finding paradigm more neatly than the making paradigm: her active process was more search than creation. This is most apparent from the perspective of Gilbert's subjective experience, which is characterized by a strong sense of destiny: she embarks on her quest in part because a medicine man in Bali, reading

her palm, predicts her return. ("I'm the kind of person who, when a ninth-generation Indonesian medicine man tells you that you're destined to move to Bali and live with him for four months, thinks you should make every effort to do that," she explains.)[59] Even for those inclined to skepticism about the prophetic powers of medicine men, however, there's a more prosaic sense of destiny by which to understand Gilbert's search—one that can be discovered or realized in something like the way we describe people as finding their vocations. Gilbert's journey, insofar as it fits the finding model, was guided by her sense of self—her already existing talents, tastes, and dreams—rather than being an effort to remake herself. Late one night she looked around her suburban bathroom and finally saw, clearly, that the life she was living didn't resemble the person living it. She left and, drawing on the talents, passions, experiences, and social connections she already had, found one that did.

§3.5.2.3 Making

For people whose critical speculations bring them to a point from which there seems to be *no* evaluative context within which their lives might have meaning, or *no* distinctive, coherent ethical self to be interpreted, the model of *making* meaning—wringing it from the void by an act of sheer will—is an attractive last resort. Something like this appealed to Wallace, during certain phases, as a solution to his troubles: "Hyperc[onsciousness] makes life meaningless," he wrote in a personal notebook, "but what of will to construct OWN meaning? Not the world that gives us meaning but vice versa?"[60]

From the CI perspective, as we noted above, the idea of *making* meaning in this radical sense is incoherent: interpretation is not a *fundamentally* creative act. We uncover the meanings of the objects we interpret; we don't invent them. There is, nonetheless, a way of understanding the idea of making ethical meaning friendly to the CI account. For, although we don't directly make the *meanings* of our lives or their constituents, we often *do* play a significant role in making our lives, along with the relationships, experiences, and accomplishments they comprise. Some marriages are successful because each member in the relationship has, by some miracle, found a soulmate in the other. Other marriages are successful because the people in them work to *become* soulmates: each person changes major parts of who she is to become a better partner for the other.

This last observation highlights an important truth: the finding and making paradigms are really two sides of the same coin. (Few if any successful relationships, after all, are wholly characterized by either the finding or making model. They generally involve some balance between becoming a different person than you were, and finding someone compatible enough with the person you already are to motivate the hard-won changes.) Lives aren't constituted by their meanings in the same way words are: for a word to be a word, it must be meaningful; for it to be the word that it is, it must have the particular meaning it does.[61] Because there is this sense in which a life's identity conditions are less strictly bound to its meanings than a word's, lives have *potential* meanings in a way that words don't. In some cases, we realize these potential meanings by recognizing them, and then treating them as a guide for living: Gilbert is told her destiny by a medicine man, and she undertakes a journey to fulfill it. In other cases, our lives come to possess

the meanings they do by creative effort: we make ourselves, as Nietzsche suggested, as a work of art.[62] Ultimately, however, these are complementary aspects of the same holistic process.

The making aspect of the process, however, is especially salient when we consider ethical meaning in the context of suffering and misfortune. Shenk described Lincoln as having "*forged* meaning from his affliction," and his choice of words was well-advised: it would do our misfortunes too much credit to attribute to *them* the value that, with sufficient courage and imagination, we can produce in response to them. Here, Viktor Frankl himself is an obvious example. Though Frankl drew on his experiences in the Holocaust in his therapeutic practice and writing, those meaningful accomplishments were *made from*, rather than *found in* those experiences. (This isn't to suggest that the trials Frankl faced were not themselves meaningful— but the positive ends toward which he redirected his sufferings were of his own making, rather than merely being contained, however hidden, in the hardships he bore.)

Shenk's Lincoln is an inspiring illustration of the way that seeing, finding, and making ethical meaning work in tandem. "In mythical stories," Shenk writes,

> a character undertakes a journey, receiving at every step totems that, at the time, have no clear value but at the end turn out to provide the essential tools for a final struggle. We can see this in Lincoln's journey. In the first stage, he asked the big questions. Why am I here? What is the point? Without the sense of essential purpose he learned by asking these questions, he may not have had the bedrock vision that governed his great work. In the second stage, he developed diligence and discipline, working for the sake of work, learning how to survive and engage. Without the discipline of his middle years, he would not have had the fortitude to endure the disappointments that his great work entailed. In the third stage, he was not just working but doing the work he felt made to do, not only surviving but living for a vital purpose. Yet he constantly faced the same essential challenges that had been presented to him throughout. All through his career fighting the extension of slavery, and all through his presidency, he faced painful fear and doubt— indeed, he faced it on an awful scale. But he repeatedly returned to a sense of purpose;

from this purpose he put his head down to work at the mundane tasks of his job; and with his head down, he glanced up, often enough, at the chance to effect something meaningful and lasting. We justly look upon the transcendence of his final days with admiration, noticing the amazing balance between earthly works and self-dissolution. But even then, he was a product of his journey.[63]

Conclusion

William James wondered if life was "a real fight, in which something is eternally gained for the universe by success," or if, instead, "it is no better than a game of private theatricals from which one may withdraw at will."[64] Understanding ethical meaning in light of the comprehensive-interpretive account of meaning I've proposed above doesn't settle James' question, or tell us whether the dichotomy he poses is a false one. It doesn't tell us whether the meaning in or of our lives has a natural or supernatural source, or whether meaning is possessed by knowing God, having and raising children, or standing in awe of natural beauty. It doesn't even tell us if our lives have meaning at all. (It doesn't itself, that is, justify or vindicate our tendency to treat our lives as ethical objects with import in larger contexts of value, or as themselves a context of value within which experiences, relationships, and other objects take on a particular sort of significance.) Answering these questions, according to the CI approach, requires a substantive account of value and its metaphysics.

What the CI approach does provide is a clearer idea of the object of the anxieties and ambitions we articulate when we use "meaning" in ethical discourse. It vindicates our sense that questions about the meaning in and of our lives draw us beyond "the buzzing and jigging and vibration of small interests and excitements that form the tissue of our ordinary consciousness,"[65] challenging us to question the routines that give us comfort, and to confront the possibility that

we've missed or lost something of profound importance. It also, however, cautions us against supposing that the concept of ethical meaning itself requires us to adopt especially grandiose standards for its possession (like requiring that we commune with some infinite, absolutely meaningful being). Ethical meaning is as prosaic or extravagant as the genuine value that animates it.

The CI approach reminds us that meaning in and of our lives shouldn't be conflated with their appreciation, either in its cognitive or affective dimensions. In doing this, I hope it offers those of us who suffer from depression some measure of solace, and a reason to persevere in times of despair. It warns us that not all meaningful lives are good lives, and that our desire to lead a meaningful life can inspire acts of terrible evil as well as acts of great moral courage. Finally, even if a clearer concept of ethical meaning doesn't, on its own, provide us detailed instructions for how to lead or appreciate a meaningful life, it does provide the resources for a general description of the dynamics of those achievements. If living meaningfully (and appreciating the meaning in and of the lives we lead) is possible at all, it requires a cultivated sensitivity to what, among the welter of dreams, impulses, and ambitions that impose themselves on us, is genuinely valuable; it requires the energy and the courage to search for the things that matter when they don't come conveniently; and it requires the imagination, conviction, and discipline to bring new value into the world, even in the face of suffering and misfortune.

---

[1] Liz Garbus, *What Happened, Miss Simone?*, Documentary, Biography, Music, (2015).

[2] Viktor E. Frankl, *Man's Search for Meaning*, 1 edition (Beacon Press, 2006), 113.

[3] William James, "Is Life Worth Living?," in *The Will to Believe: And Other Essays in Popular Philosophy* (Longmans, Green, and Company, 1896), 39.

[4] Thaddeus Metz, "Recent Work on the Meaning of Life," *Ethics* 112, no. 4 (July 2002): 801, doi:10.1086/340462.

[5] This slogan needs some qualification (which it received in the previous chapter), but for now it serves well enough.

[6] James, "Is Life Worth Living?," 32.

[7] Charles Taylor et al., "The Politics of Recognition," in *Multiculturalism: Examining the Politics of Recognition*, ed. Amy Gutmann, Expanded paperback edition (Princeton, N.J: Princeton University Press, 1994), 72. Here and elsewhere, Taylor is probably using "meaning" broadly, in a way that includes both ethical meaning and other varieties—consonant with the metaphorical way he construes "language" earlier in the essay, as "covering not only the words we speak, but also other modes of expression whereby we define ourselves, including the 'languages' of art, of gesture, or love, and the like" (32).

[8] Kwame Anthony Appiah, *The Ethics of Identity* (Princeton University Press, 2007), xii.

[9] Ibid., 24.

[10] To illustrate his point, Appiah invites us to imagine a scenario in which, "through some sort of instantaneous genetic engineering," we were able to alter any aspect of our capacities, desires, and appearances—"you could have Michael Jordan's fade-away shot, Mozart's musicality." Rather than describing a utopia, Appiah thinks, this would be "a kind of hell. There would be no reason to choose any of these options…life would be meaningless." Ibid., 19. Elisabeth Camp makes a similar point in an essay on Wordsworth and the narrative theory of the self, asserting that a person who limited herself to exclusively impersonal reasons could only live a life that, if it were "practically feasible at all," would be characterized by "meaningless and existential despair." Elisabeth Camp, "Wordsworth's Prelude, Poetic Autobiography, and Narrative Constructions of the Self," *Nonsite.org*, accessed April 21, 2015, http://nonsite.org/issues/issue-3/wordsworth%e2%80%99s-prelude-poetic-autobiography-and-narrative-constructions-of-the-self.

[11] Metz, "Recent Work on the Meaning of Life," 802–3.

[12] Susan Wolf et al., *Meaning in Life and Why It Matters* (Princeton, N.J.: Princeton University Press, 2012), 3.

[13] Robert Nozick, *Philosophical Explanations*, Fifth Printing (Belknap Press of Harvard University Press, 1983), 612, 570.

[14] There are a variety of ways to cash out this distinction as it applies to ethical meaning, and some of the cases presented below are ambiguous with respect to which version of the distinction is at issue. I'll attend to those clarifications in a subsequent section, when I address the questions that, for the moment, I'm simply concerned with posing.

[15] Appiah, *The Ethics of Identity*, 11, 111.

[16] Richard Taylor, *Good and Evil: A New Direction : A Forceful Attack on the Rationalistic Tradition in Ethics* (Prometheus Books, 1970), 259.

[17] Ibid., 266, 263, 268. Interpretively speaking, Taylor is actually a complicated case. Wolf casts him, as I do in the body text, as her subjectivist foil. On her paraphrasing, Taylor thinks the injection transforms

Sisyphus's life "from horribly unfortunate to exceptionally good" (Wolf, 17). This is a loose but generally fair representation of the view he records in *Good and Evil*, where the Sisyphus thought experiment first appeared, and from which I drew the quotations above. But in a subsequent treatment of the same issue ("The Meaning of Human Existence," collected in several Intro readers), Taylor reversed course. There he argues, for example, that nuns who chant prayers all day lead paradigmatically meaningless lives (assuming their religious beliefs to be mistaken), regardless of the sense of spiritual fulfillment these rituals bring them, or the spiritual and moral concerns that motivate them. For dialectical convenience, I follow Wolf in focusing on Taylor's earlier, more strongly subjectivist take on the issue, but it's worth noting that accurately rendering Taylor's considered view, accounting for all his writings on the subject, would need to be more complex (and, given that the second essay makes no mention of the first, potentially indeterminate).

[18] Wolf et al., *Meaning in Life and Why It Matters*, 20, 25, 11.

[19] Ibid., 44. It should be noted that Wolf's central point doesn't hang on the cogency of these particular examples: she seems happy enough to admit that someone might be able to demonstrate that the objective value of the imagined assistant's or lawyer's efforts are greater than she gives them credit for being.

[20] Clifford Geertz, "Thick Description: Toward an Interpretive Theory of Culture," in *The Interpretation of Cultures: Selected Essays* (Basic Books, 1973), 5.

[21] I borrow the nutrition analogy from Gary Hatfield, who uses it to illustrate his relationalist theory of color. See Gary Hatfield, "Objectivity and Subjectivity Revisited: Colour as a Psychobiological Property," in *Colour Perception: Mind and the Physical World*, ed. Rainer Mausfeld and Dieter Heyer (Oxford University Press, 2003), 187--202.

[22] Stephen Macedo, "Introduction," in *Meaning in Life and Why It Matters* (Princeton, N.J.: Princeton University Press, 2012), xi.

[23] This list of example object kinds isn't exhaustive. Many sorts of entities admit of ethical interpretation.

[24] One of the most moving examples of the ethical interpretation of an object I know of is recounted by Viktor Frankl in the Preface of the 1992 edition of *Man's Search for Meaning*, in which he explains the circumstances that led to his eventual internment in a concentration camp: "The reader may ask me why I did not try to escape what was in store for me after Hitler had occupied Austria. Let me answer by recalling the following story. Shortly before the United States entered World War II, I received an invitation to come to the American Consulate in Vienna to pick up my immigration visa. My old parents were overjoyed because they expected that I would soon be allowed to leave Austria. I suddenly hesitated, however. The question beset me: could I really afford to leave my parents alone to face their fate, to be sent, sooner or later, to a concentration camp, or even to a so-called extermination camp? Where did my responsibility lie? Should I foster my brain child, logotherapy, by emigrating to fertile soil where I could write my books? Or should I concentrate on my duties as a real child, the child of my parents who had to do whatever he could to protect them? I pondered the problem this way and that but could not arrive at a solution; this was the type of dilemma that made one wish for 'a hint from Heaven,' as the phrase goes.

"It was then that I noticed a piece of marble lying on a table at home. When I asked my father about it, he explained that he had found it on the site where the National Socialists had burned down the largest Viennese synagogue. He had taken the piece home because it was a part of the tablets on which the Ten Commandments were inscribed. One gilded Hebrew letter was engraved on the piece; my father explained that this letter stood for one of the Commandments. Eagerly I asked, 'Which one is it?' He answered, 'Honor thy father and thy mother that thy days may be long upon the land.' At that moment I decided to stay with my father and my mother upon the land, and to let the American visa lapse." (Frankl,

Man's Search for Meaning, xv-xvi.) Frankl lost his wife, who was pregnant, and both of his parents to the Holocaust. He was, however, present to comfort his father as he died, providing him with morphine to ease his pain in his final days.

[25] This sort of *intellectual alienation* may be salutary or regrettable, depending on the case. Understood from the CI perspective, intellectual alienation occurs when our cognitive attitude toward an object (where "object" is construed liberally, as usual) shifts from one that treats it as possessing immediate, intrinsic import, to one that treats it as doubtful or enigmatical on its surface, and thereby calls into question its meaning. In the former attitude, the appearance/reality distinction isn't active; in the second, it is. Intellectual alienation can be helpful, for example, in cases in which we've been *fetishizing* the objects of our concern—when we've conflated a medium and its meaning, failing to differentiate between a symbol and what it represents, or a tool and what it can be used to bring about. Scrooge, for example, is the archetype of a person who suffers from a wealth fetish, having come to value money for its own sake, rather than for what it can be used to do—and the intellectual alienation from that motivational orientation he eventually experiences changes his life, and the lives of others, for the better. In other cases, however, intellectual alienation can be a pathology. People who have developed hostility to a certain sport, for example, often invite others to adopt a perspective according to which the internal goal of the game should have to justify itself by demonstrating some external utility—as if, say, baseball were analogous to carpentry: "What's the point of grown adults trying to hit a ball with a stick while others wildly chase after it, as if it's full of tasty jelly beans instead of a bunch of worthless rubber bands?" If a person who once thoroughly enjoyed baseball were convinced to feel ridiculous by this sort of spurious demand for meaning, that would probably be regrettable.

[26] Relatedly, the particular kinds of (apparent) goods I've chosen as examples of stereotypically superficial goods are not themselves essential to my point; the sorts of apparent goods that motivate and satisfy superficially can vary from person to person, and culture to culture. This isn't to say that some sorts of goods don't more readily lend themselves to superficial motivation and enjoyment. It seems much easier, for instance, to see pleasure as desirable or valuable absent being placed in some larger evaluative context than it does to see suffering in this way—which is probably one reason why, in the case of suffering, we're far more frequently inclined to wonder about its meaning.

[27] Wolf et al., *Meaning in Life and Why It Matters*, 4. Reasons of love, for Wolf, aren't confined to cases like the ones illustrated in the quote above, where our love is directed toward another human being. They can also arise from "nonpersonal pursuits" that engage our passion, like making art or even doing philosophy—provided that these objects possess some genuine value that "lies outside of oneself" (Wolf, 5). I'll have more to say about the kind of independent value Wolf thinks meaning-generating objects and activities need to have when I discuss the question of objectivity, below.

[28] Nozick, *Philosophical Explanations*, 570.

[29] Ibid., 594.

[30] Ibid., 599, 604.

[31] Ibid., 612–13.

[32] It should be noted that Wolf doesn't want to deny that projects that make life meaningful can also have great moral worth: "meaning and an interest in meaning," she thinks, "are likely more often than not to complement and reinforce moral concerns." Nonetheless, for morally valuable projects to count as meaningful, she requires that they be pursued out of love, rather than merely from a sense of duty. Though the distinction between these two sources of motivation isn't wholly clear it's at least partly based on the intuitive requirement that love is characterized by the enjoyment of certain positive "feelings and attitudes,"

111

which Wolf variously characterizes as "passionate engagement," "fulfillment," (14) and "high quality pleasure" (27). From the CI perspective, however, a morally important project may generate meaning even if it doesn't produce such feelings—an issue that will be addressed in more detail in the discussion of subjectivity and objectivity.

[33] In an analysis of Lincoln's struggle with depression, Joshua Wolf Shenk has argued convincingly that the extreme psychological suffering Lincoln endured, which, in his early adulthood, sometimes left him overwhelmed to the point of incapacitation by the difficulty of "understanding the world and his place in it," nonetheless prepared him, later in life, to move beyond "self-centered concern with his own suffering…to see and grapple with the suffering around him." Lincoln "forged meaning from his affliction so that it became not merely an obstacle to be overcome, but a factor in his good life." Joshua Wolf Shenk, *Lincoln's Melancholy: How Depression Challenged a President and Fueled His Greatness* (Boston: Houghton Mifflin Harcourt, 2005), 64, 8.

[34] Nozick, *Philosophical Explanations*, 570.

[35] Ibid., 597.

[36] Ibid., 599.

[37] Ibid., 602.

[38] Thomas Nagel makes a similar point (though not with direct reference to Nozick) in Thomas Nagel, "The Absurd," *Journal of Philosophy* 68, no. 20 (1971): 717.

[39] Metz, "Recent Work on the Meaning of Life," 781.

[40] Aspects of the psychology of mass shootings suggest that they can be understood, at least in part, as perverted responses to a desire for meaning. Consider, for example, the tendencies of shooters to stage their crimes in especially meaningful locations (for example, schools or offices that they associate with significant failures or rejections); to identify obsessively with particular narrative portrayals of violence (*Natural Born Killers* and *The Dark Knight Rises* are prominent examples); or to associate their actions with highly salient political movements despite having had little or no formal involvement with the causes to which they claim to pledge allegiance. Each tendency reflects, in different ways, the killer's ambition to understand his actions symbolically, within some larger evaluative context. Nor do the killers always see their own actions unambiguously, as contributing positively to some (from an external perspective, perverse) set of values. James Michael Holmes, for example, reports having attempted to call a mental health hotline just before the massacre he committed at the midnight showing of *The Dark Knight Rises* in Aurora, CO, hoping someone might talk him out it. At Holmes's trial, the prosecution used this claim as evidence that Holmes knew right from wrong. (Sadie Gurman, "Theater Gunman Says He Called a Crisis Line," *The New Orleans Times-Picayune*, June 4, 2015.)

[41] Wolf et al., *Meaning in Life and Why It Matters*, 9.

[42] Ibid., 21.

[43] See Metz (2002) for a selective survey of philosophical work that takes boredom as necessarily undercutting or detracting from ethical meaning—an assumption he says few have thought to question. Metz, "Recent Work on the Meaning of Life," 791, 794.

[44] Garbus, *What Happened, Miss Simone?*

[45] Elif Batuman, "Kafka's Last Trial," *The New York Times*, September 22, 2010, http://www.nytimes.com/2010/09/26/magazine/26kafka-t.html.

[46] This point relates to another, which also mitigates against counting the suffering of James's Dog as a source of meaning in or of its life. Returning to the idea that meanings, in general, are interpretive affordances, it's worth noting that, although an object's meanings are independent of the way it is *actually*

interpreted on a particular occasion (a piece of physical evidence might afford a key insight that would crack the case, even if the detective who examines it is too careless to notice), they must, in some broad sense, be *accessible* to the interpreter in question. Grass, for example, affords nourishment to cows, but not to humans. (This is one reason why the total meaning of an object isn't necessarily identical to the totality of information it carries—the latter quantity, as Dretske understands it, being totally independent of the features of any potential interpreter.) Putting aside the question of the extent to which dogs engage in interpretation in general, it seems unlikely that the medical knowledge won by his agony will be accessible in the right way to the dog. Hence, from the dog's perspective, it isn't a source of meaning in or of his life.

[47] Whether we ultimately class this sort of meaning as *ethical* depends, as with the cases of exploitation, on the details of a precise account of what makes an object or an evaluative context distinctively ethical—an interesting issue beyond the scope of this paper.

[48] Mr. Stevens suffers this last misfortune though, as with much else, he lacks the courage to admit it to himself. By the end of *Remains*, when Stevens is busy practicing his slang in the hopes of pleasing his new American master, butlering is a good ways into its descent from an important and respected social role to a decorative piece of nostalgia.

[49] Elizabeth Gilbert, *Eat, Pray, Love: One Woman's Search for Everything Across Italy, India and Indonesia* (New York: Riverhead Books, 2007), 15, 11.

[50] This is an admittedly oversimplified account of the movie's narrative. In fact, in addition to George Bailey's loss of faith in himself, there's a second (and chronologically prior) theme related to his inability to appreciate the meaning his life *does* have, even before the disastrous series of events precipitated by his Uncle Billy's loss of $8,000. George's small-town life running the Building and Loan his father started is far from the life he dreamed of for himself—a fact thrown into relief by the glamorous life of his younger brother, a heroic World War II Navy pilot. The transformation in George's attitude that takes place between his suicidal crisis, early in the film, and the deep sense of fulfillment and gratitude he enjoys at the movie's conclusion, isn't *merely* constituted by a better appreciation of his own worth, as I suggest above; it also involves a shift in his evaluative perspective—one that allows him to genuinely appreciate what he's always had, rather than pining after what he doesn't. (This is a phenomenon I'll discuss below, in the "found" section.) Correcting the simplification makes the example no less apt for the point at issue: even if George's evaluative perspective shifts over the course of the movie, his primary problem is never that he's become unmoored from any evaluative perspective at all. Both early in the movie, when he envies his brother's freedom and success, and later in the movie, when he comes to doubt his own value even for the limited role he inhabits, his discontent arises from his negative appraisal of the import of his life in the context of a suite of values he implicitly accepts.

[51] Albert Camus, *The Myth of Sisyphus: And Other Essays*, Reissue edition (Vintage, 1991), 6.

[52] David Lipsky, *Although Of Course You End Up Becoming Yourself: A Road Trip with David Foster Wallace*, 1 edition (New York: Broadway Books, 2010), 67–68.

[53] John Stuart Mill—perhaps too philosophical a thinker for his own psychological health—precipitated a crisis of this kind simply by *imagining* accomplishing all of his cherished goals. As he recalled in his *Autobiography*: "…it occurred to me to put the question directly to myself: 'Suppose that all your objects in life were realized; that all the changes in institution and opinions which you are looking forward to, could be completely effected at this very instant: would this be a great joy and happiness to you?' And an irrepressible self-consciousness distinctly answered, 'No!' At this my heart sank within me: the whole foundation on which my life was constructed fell down." (John Stuart Mill, *The Autobiography of John Stuart Mill* (Rockville, MD: Arc Manor, 2008), 63–64.)

[54] This second sort of threat to meaning appreciation explains Nozick's intuition, discussed above, that it is *always* possible to raise questions about something's meaning. The repeated question, "But what does *that* mean?" applied to each new perspective from which the previous question was answered, *seems* endlessly iterable because it's pragmatically ambiguous between two readings. According to the CI account, asked literally, the question *presupposes* some shared evaluative context. (If, in some unexceptional situation, you ask after the meaning of some everyday object, but I can't discern your larger purpose, I'll be at a loss to answer your question—and not because I lack some particular understanding of the object.) But the question can also be used, pragmatically, to make that usual presupposition the target of inquiry—so that what I'm really asking is if there's any evaluative context within which the object has a meaning at all. The contrast can be roughly illustrated by imagining two different observers asking about the meaning of a surprising move in a game of chess being conducted by two grandmasters. When a chess commentator asks herself (or a partner) "what does that mean?" she wants the move's import in the context of a game she understands. How did the move fit into the player's larger plan for winning the game? And did the move in fact strengthen or weaken the player's position? Etc. Now imagine, instead, an observer who doesn't know the rules of chess, and is trying to learn them by watching the game played. His "what does that mean?" asks a different question—perhaps it means, for example, that the king and the rook can, under certain circumstances, move in tandem. Finally, imagine an observer who's uncertain whether what he's watching is a purposeful activity at all. By asking "what does that mean?" he introduces by presupposition the hypothesis that the activity *is* purposeful, and thus that there is some larger evaluative context within which to understand it. Nozick is right to think that this last version of the question can always be sensibly posed, but its force is far less substantive than the version asked by the chess commentator. In the last case, for example, the presupposition may fail, making the literal question nonsensical, even if its illocutionary effect is comprehensible.

[55] Clifford Geertz, "Notes on a Balinese Cockfight," in *The Interpretation of Cultures: Selected Essays* (Basic Books, 1973), 434.

[56] Frankl, *Man's Search for Meaning*, 113.

[57] Unreflective people might occasionally enjoy the meaning of their lives without recognizing it, and depressives with hard-earned intellectual discipline can recognize the meaning in their lives even at times when they find themselves unable to enjoy it. Though little rides on this aspect of my stipulative vocabulary, I'm reluctant to deny that a person in either of these situations possesses the meaning in or of her life. In both cases, however, it seems natural to think that something desirable is missing.

[58] Richard P. Feynman and Timothy Ferris, *Perfectly Reasonable Deviations from the Beaten Track: The Letters of Richard P. Feynman*, ed. Michelle Feynman (New York: Basic Books, 2005), 201.

[59] Gilbert, *Eat, Pray, Love*, 28.

[60] D. T. Max, *Every Love Story Is a Ghost Story: A Life of David Foster Wallace*, 1St Edition edition (New York: Viking, 2012), 209.

[61] Here I'm simplifying a complicated philosophical issue in the interest of dialectical efficiency. The question of a word's identity conditions, for example, is a difficult one. In fact, I think that what Hillary Putnam said of meanings themselves—that they have "an identity through time but no essence"—is true of words. (Hilary Putnam, *Representation and Reality*, Revised ed. edition (Cambridge, Mass.: A Bradford Book, 1991), 11.) In this sense, they're more like lives than my simplification in the body text suggests. A good example is the word "freedom" in the context of American political history. Eric Foner's extraordinary work tracing the way that word's meaning shifted, over the course of American history, in responses to the political challenges it was recruited to address, would be unjustly trivialized if we imagined that the

"freedoms" whose succession he documents were nothing more than a set of homonyms. (See Eric Foner, *The Story of American Freedom* (New York: W. W. Norton & Company, 1999).) Nonetheless, the fact that the contrast I draw in the body text isn't absolute makes it no less genuine. Words *are* more robustly constituted by their meanings than lives are. Unlike lives, words are *symbols*: things made for the purpose of bearing meaning, which are what they are to the extent that they serve that purpose.

[62] Friedrich Nietzsche, *The Gay Science: With a Prelude in Rhymes and an Appendix of Songs*, trans. Walter Kaufmann, 1 edition (New York: Vintage, 1974), 232.

[63] Shenk, *Lincoln's Melancholy*, 215.

[64] James, "Is Life Worth Living?," 62.

[65] Ibid., 32.

IRONY AS DYNAMIC REVERSAL

A Comprehensive Account of Irony

§1 Irony's Integrity

Irony's ambit in contemporary philosophy is nearly as broad as its history is long, and its history begins with Socrates. Today, irony preoccupies philosophers of language and psycholinguists, who work to integrate irony into broader theories of mind, language, and communication; ethicists, who hope to elucidate the role irony plays in the way we understand our experiences and ourselves; and ancient philosophers, who wonder whether the "irony" Socrates practiced was a kind of deep pedagogy, or just a disingenuous rhetorical trick. Thanks to Richard Rorty, "ironism" is even the name for (an appropriately anti-metaphysical) approach to metaphysics— one that forswears any appeal to truths vouchsafed by mind-independent reality.[1] This remarkable versatility makes irony interesting, but difficult to understand.

The difficulty is compounded when we consider how variously we use the word "irony" in everyday discourse. Individual utterances can be ironic, but so can entire works of literature. Irony can characterize situations as well as linguistic performances. We worry, increasingly, that our culture has become "too ironic," but when we discuss art and literature we often use "irony" as a success term, like "beautiful" or "elegant": consider *New Yorker* critic James Wood's distaste for David Foster Wallace's fiction, in which, Wood complains, "irony is starved to sarcasm, and sympathy to voyeurism. It is literally impossible for the reader to enter the story."[2]

You might, therefore, suspect that if "irony" names even a family resemblance concept, it's a family like the 19th Century Hapsburgs: big, incestuous, and plagued by a host of illegitimate

pretenders. You'd be in good company. Most contemporary treatments of irony make early

mention of its troublesome promiscuity. Wayne Booth worries that "irony has come to stand for

so many things that we are in danger of losing it as a useful term altogether,"[3] and Jonathan Lear

laments that "by now 'irony' has been used for pretty much everything."[4] Dan Sperber and

Deirdre Wilson warn against assuming that irony is a natural kind. If it isn't, they think "it is

possible that verbal irony and its associated attitude have about as much claim on our attention as

black bile and the atrabilious temperament."[5]

Irony isn't the protagonist of an Edith Wharton novel, though, so maybe it isn't necessary

to worry ourselves too strenuously about its integrity. A sensible response to the sort of

apparently promiscuous usage detailed above might just be to divide the disciplinary terrain,

making progress where it can be had, rather than distracting ourselves with the task of

comprehending a whole whose unity is questionable to begin with. In his investigation into the

ethics of ironic experience, for example, Lear doesn't have much time for detailed disputes about

whether one- or two-staged models of linguistic comprehension best capture the psychological

dynamics of producing and understanding the sort of one-off sarcastic remarks adolescents

typically employ; Sperber and Wilson, meanwhile, can be forgiven if they don't think a careful

interpretation of Kierkegaard is the most promising avenue to achieving a better understanding of

the psychological "natural kinds" that interest them.

But I think irony's diverse theorists have more to learn from one another than the current

state of discourse suggests. We *should* care about irony's integrity. First of all, in addition to its

broad currency in discourse, irony is strangely notorious for a mere figure of speech. Echoing the

complaints of Socrates's interlocutors, Theophrastus characterized an ironist as someone who "never can be got to do anything, or to commit himself in speech so that he is forced to take sides in an active discussion."[6] Much more recently, Christie Wampole blamed irony for "a vacuity and vapidity of the individual and collective psyche,"[7] and Maryan Ronagh and Lawrence Souder have identified the increasingly common practice of publishing ironic research papers in science journals as a violation of research ethics whose dangers have been severely underrated.[8] Interesting as they are, though, ethical, aesthetic, and cultural discussions of irony too frequently rely on vague or inadequate definitions of irony (if any definition is offered at all). The careful thought and the increasing body of empirical research that philosophers of language and psycholinguists have dedicated to irony should help to clarify questions about its ethics and aesthetics.

Second, the problem of apparently promiscuous usage doesn't just plague the interdisciplinary study of irony—it arises *within* the empirical study of irony as a purely psycholinguistic concern. Theorists have become increasingly divided on the range of linguistic phenomena that a theory of irony ought to cover. For example, should instances of extended discourse (like Jonathan Swift's "A Modest Proposal") be treated as cases of verbal irony, or is verbal irony properly restricted to irony at or below the level of individual utterances? Should researchers treat sarcasm, hyperbole, and understatement as "types" of irony, as Raymond Gibbs did in a recent study,[9] or is Deirdre Wilson right to protest that "merg[ing such a variety of utterances] into a general category of 'irony' detracts from, rather than enhances, our understanding of how irony works"?[10] Taking a broader perspective helps to answer questions

118

like these. Understanding the common character of irony as it arises across media puts us in a

position to make principled distinctions between purported instances of verbal irony. In fact,

careful attention to the way that the various kinds of irony (situational, literary, and verbal, for

example) cohere reveals significant shortcomings in some of the leading accounts of verbal irony.

None recognize, for example, the fact that ironists can be sincere.

In this chapter, I offer an account of irony as *dynamic reversal*, according to which irony

characterizes the class of interpretive structures whose inputs (or objects of interpretation, or

media) interact disharmoniously with their outputs (or *meanings*), so that one reverses or

undermines the other. (The idea of an interpretive structure draws on the theory of interpretation

developed in Chapter 2.) The account departs from prevailing psycholinguistic theories of irony

by identifying irony with a formally characterized class of interpretive structures, rather than with

any particular psychological mechanism. This methodological shift has a number of advantages

over traditional mechanistic approaches. First, it does better justice to the increasing body of

empirical research that shows irony can be effected and interpreted by more than one

psychological mechanism. (As Gibbs has recently concluded, "ironic talk can serve multiple

communicative purposes, each requiring different psychological mechanisms.")[11] Second, it

provides a principled basis for discriminating ironic from unironic usage for the purpose of

designing and interpreting experiments on the psychology of irony production and

interpretation, whereas the distinctions offered by mechanistic approaches too often beg the

question. Third, the interpretive approach to irony does a better job accounting for the

underlying character that unites irony's various kinds, rather than disfiguring some to fit the

model of others. Finally, the interpretive approach provides the basis for a richer understanding of the full range of irony's uses, including sincere deployments of irony overlooked or ruled out by approaches that restrict irony to productions and comprehensions achieved by means of a single psychological mechanism. In short, the dynamic reversal account gives better answers to questions about irony's *mechanism* (how does irony work?), its *use* (what do we use irony to do?), and its *continuity* (how are different sorts of irony—verbal irony, dramatic irony, and situational irony, for example—related to one another?).

By providing more complete answers to these questions, the dynamic reversal account offers a basis for more productive interdisciplinary exchange, and puts us in a better position to address the ethical, political, and aesthetic concerns irony raises. Its aim isn't to provide an occasion for pedantry, by stating a single rule for deciding, definitively, whether a particular instance deserves to be called "ironic"—according to the dynamic reversal account irony involves a combination of characteristic features, each of which comes in degrees. Instead, the account should help us better understand both our temptation and reluctance to include borderline cases. More importantly, it should help us recognize irony's potential both to enrich and to enervate human experience and communication—to decide when an irony is elegantly or poorly constructed, rightly or wrongly deployed.

I'll begin with a critical review of recent theories of verbal irony: Grice's standard implicature view, Sperber and Wilson's echoic account, and the pretense theory defended by Herbert Clark and Richard Gerrig, and recently updated by Gregory Currie. In the process of highlighting the shortcomings of each of these accounts, I introduce, successively, the positive

commitments that characterize dynamic reversal: twofold interpretation, dynamism, and reversal. I conclude with a discussion of the ways the dynamic reversal account facilitates empirical inquiry, encourages interdisciplinary exchange, and provides a foundation for addressing ethical and aesthetic questions about irony.

§2 Verbal Irony

§2.1 The Standard Gricean Model

Philosophers who achieve too dominant an influence over a particular field of philosophical inquiry risk sharing something like the storyline of the Scarecrow in *The Wizard of Oz*—a strawman that, for most of the movie, everyone tries to light on fire. When that plotline finally gets old, the revisionists appear to tell us that the scarecrow had a brain all along.

In contemporary discussion of verbal irony, Paul Grice is the scarecrow. Like most theories of irony, Grice's begins with the old idea that ironists mean the opposite of what they say; his distinctive contribution was to integrate this basic idea into a broader theory of pragmatics. According to Grice, irony is a species of conversational implicature, the phenomenon in which speakers, exploiting the shared expectations that govern conversation as a rational activity, communicate more than they literally say.[12] On Grice's picture, and ironist flouts (violates in a way intended to be transparent to their audiences) the conversational maxim that instructs speakers not to say what they believe to be false; and her audience, in order to preserve the assumption that the ironist is nonetheless cooperating in the conversation, concludes that she

must not really have intended to communicate what she's said. In fact, she must really have meant to communicate its opposite. For example, when Billy Bones says to Jim

(1) Silver's a fine friend!

when Jim and Billy know in common that Silver has betrayed Billy, Jim will recognize that Billy has transparently said something he doesn't believe, and conclude that he must therefore have been trying to communicate something other than what he has literally said. "This must be some obviously related proposition;" Grice explains, "the most obviously related proposition is the contradictory of the one he purports to be putting forward."[13] The basic process Grice posits involves an initial, transparent insincerity, signaling an operation of meaning reversal.

Dan Sperber and Deirdre Wilson have capably identified the weaknesses of this standard Gricean view. First, Grice's account covers only assertions, but non-assertoric speech acts can be used ironically as well. For example—

(2) Did you remember to water the garden? [uttered in the middle of a downpour]
(3) Thanks for holding the door! [when someone has just let a door close in your face]

In neither case does the speaker say something she believes is false, so Grice's insincerity condition is, at best, too limited. **(2)** also calls the reversal condition into question: if questions have contradictories or opposites, the plausible candidates ("Didn't you remember to water the plants?" "Did you forget to water the plants?") are not what an ironic speaker of **(2)** really means.

Second, Sperber and Wilson point out that Grice's account makes it difficult to see why anyone would ever choose to speak ironically. As Grice construes irony's mechanism, the contextual information an ironist's audience must possess for the irony to be successfully

recognized *includes* the proposition that the ironist is supposedly attempting to communicate. So, for instance, in **(1)**, Jim has to know already that Billy thinks Silver is a bad friend in order to initiate the reasoning process whose output is the recognition that Billy is using **(1)** to communicate that he thinks Silver is a bad friend. Thus, Sperber and Wilson conclude, "the standard approach… would make every ironical utterance uninformative, both on the level of what is said and on the level of what is implicated."[14]

Before examining the revisions to the basic Gricean picture that have been proposed to remedy these difficulties, let's take a moment to review the standards for a good theory of verbal irony that the discussion above implies. First, a good account of *what irony is* should provide (or, at least, not preclude) an explanation of how irony works—I've called this the **mechanism question**. Second, a good account should provide (or, at least, not preclude) an explanation for what makes irony useful—the **usefulness question.** Finally, in line with the considerations I mentioned at the outset, a good theory of verbal irony should render perspicuous (or, at least, not obscure) its *continuity* with other recognizable types of irony (e.g., dramatic and situational irony)—this is the **continuity question**.[15] With these standards in hand, let's turn our attention to the contemporary accounts that try to meet them.

§2.2 Echoic Mention

Sperber and Wilson suggest replacing the standard meaning-inversion view with an "echoic mention" theory of irony.[16] They contend, as we've seen, that meaning substitution theories preclude good answers to both the usefulness and the mechanism questions. So they claim,

instead, that irony should be understood as an *echoic mention*, rather than as a *use* of the proposition ironically uttered. An ironist mentions, by echoing, a piece of propositional content, in order to express a dissociative, derogatory attitude toward it.

Grasping Sperber and Wilson's idea of echoic mention requires that we expand our usual understanding of both echoing and mentioning. Nonetheless, the basic sort of phenomenon they have in mind is familiar enough, as in the following exchange between A and B:

(4) A: Where can I buy pretzels at this time of night?
    B: Where can you buy pretzels? At this time of night? At Barney's, of course.[17]

B clearly "echoes" A's question. And when she utters the questions she does, she isn't *using* them in the way A is. Instead, B is "implicitly mentioning" A's question: she utters the content in order to refer to it, rather than to use it. The relevant technical details of the special kind of mentioning involved here are (1) the mentioning is implicit (that is, B does not set off the questions she utters as an act of quotation by employing an explicit semantic marker, as she would have, for instance, if she had instead uttered, "Where can you buy pretzels, you ask?"); and (2) the mention is of a *proposition* rather than a linguistic expression, so B refers to (but does not use) the *meaning* of the questions mentioned, rather than their linguistic features. (It's not like the kind of mention made here: "'Where can I buy pretzels?' is a question composed of five words.") So, in an echoic mention, we refer to, without using, the proposition we implicitly mention.

Of course, **(4)** isn't itself an instance of *irony*—irony arises, on Sperber and Wilson's view, when the speaker expresses a "dissociative" attitude toward the echoed content, "mak[ing] it clear that he rejects it as ludicrously false, inappropriate, or irrelevant."[18] They provide an illustration

from *Pride and Prejudice*, in which Darcy and Elizabeth are discussing Wickham, whom Elizabeth pities and admires, but Darcy knows to be a cad:

(5) "You take an eager interest in that gentleman's concerns," said Darcy in a less tranquil tone, and a heightened colour.
"Who that knows what his misfortunes have been, can help feeling an interest in him?"
"His misfortunes!" repeated Darcy contemptuously, "yes, his misfortunes have been very great indeed."[19]

On Sperber and Wilson's analysis, Darcy echoes Elizabeth's rhetorical question, indicating by the attitude he adopts toward its meaning that he regards it as "ludicrously false," and thereby implicates that Wickham "has not been the victim of misfortunes."[20]

Not all ironic utterances are as immediately amenable to an echoic explanation, of course: Sperber and Wilson concede that irony is frequently effected without any direct echo of an immediately preceding utterance. Instead, they argue for an extended sense of "echo"—"some are immediate echoes, and others delayed; some have their source in actual utterances, others in thoughts or opinions; some have a real source, others an imagined one; some are traceable back to a particular individual, whereas others have a vaguer origin."[21] Though this extension has drawn complaints from some commentators, [22] it's worth accepting Sperber and Wilson's terminology— at least provisionally—to notice the way their proposal remedies some of the problems they identified in Grice's account.

First, Sperber and Wilson employ a different, broader version of the insincerity condition than Grice does. Instead of requiring an ironist to say (or make as if to say) something she believes to be false, an ironist *dissociates* herself from the propositional content she alludes to, in order to ridicule it. This alteration remedies the defect of limiting irony to assertions. Second,

instead of viewing irony as necessarily involving meaning inversion, they emphasize irony's expressive function: ironists are primarily concerned with expressing an attitude toward a given proposition, rather than communicating propositional content (opposite meaning or otherwise). This reorientation effectively addresses the uninformativeness objection. Because, according to echo theory, ironists comment on mentioned utterances, rather than simply using their converse meaning, the illocutionary force of an ironic utterance is no longer obviated by the conditions necessary for its successful interpretation. Finally, an expressivist account ameliorates problems associated with traditional approaches to "figurative meaning"—how can interpreters be expected to disambiguate the countless potential meanings a single utterance can have if figurative meanings are allowed?—and with integrating irony into Grice's more general theory of implicature, which generally involves a speaker implicating something *in addition* to what she's said, rather than *instead of* something she's said. On the echoic mention account, Sperber notes approvingly, irony "involve[s] only one meaning, the literal one."[23]

Nonetheless, the echoic theory of irony fails to capture several of irony's interesting features—especially when we turn our attention to subtler, more complex cases. In particular, it rules out the interplay between multiple possible interpretations (including sincere interpretations) characteristic of a certain class of complex ironies, and it fails to account for the important fact that irony offers its users opportunities to modulate the transparency of their intentions to achieve a variety of important illocutionary effects. These aspects of irony, I suggest, can be better explained by recognizing the "twofold" nature of ironic interpretation, according to which interpreters attend to the features of both a medium and its meaning.

126

§2.2.1 Interplay with Sincerity

Consider **(5')**, a reimagined version of **(5)** in which there's a sense, though different from

Elizabeth's, in which Darcy *really does* think of Wickham as having been a victim of misfortune.

Perhaps Darcy has an opinion something like Socrates's, according to which committing injustice

(which Darcy knows Wickham to have done) is the worst thing that can happen to a person. Or,

suppose he's just now seen Wickham fall of his horse into a pile of manure, and has this in mind

when he responds to Elizabeth. In cases like these, there's a plausible interpretation on which

Darcy *does* straightforwardly endorse the literal meaning of the proposition he's uttered. This sort

of case is especially vivid when we imagine a second, knowing audience member: suppose, for

example, that Bingley is also present for the exchange, and similarly witnessed Wickham's

mishap.

In cases like **(5')**, awareness of the presence of the sincere interpretation of Darcy's

utterance is essential to a full appreciation of Darcy's irony: because of it, Darcy can appear to

agree with Elizabeth while, in the very same breath, he completely subverts her meaning.

(Elizabeth's suggestion is that the misfortunes Wickham has suffered are grave, and that he

should, therefore, be regarded with compassion and respect; Darcy's is that Wickham's

misfortune is buffoonish, and makes him a proper object of ridicule. Elizabeth's comment carries

the additional implication that Darcy is to blame for Wickham's misfortunes, while Darcy's

suggests that Wickham has brought them upon himself.)

If achieving this sort of (appearance of) double meaning is a central feature of Darcy's

irony in **(5')**, the strict criterion of dissociation that Sperber and Wilson advocate is too simplistic

to capture it (or others like it). For, while the irony does rely, on one hand, on Darcy adopting a

certain kind of dissociative attitude toward the way *Elizabeth* means the utterance, it equally plays

upon another sense in which he does, or can, or pretends to *identify* with his utterance in the

straightforward way. So a *merely* dissociative mention analysis doesn't capture the full force of the

irony. Furthermore, echo's inability to render the complexity of cases like **(5')** obscures an

otherwise clear point of continuity between verbal ironies like Darcy's and paradigmatic cases of

literary and dramatic irony, in which a character in a novel or play makes a sincere statement

whose deeper significance is, nonetheless, grasped only by the novel's reader, or the play's

audience.

Notice also that, in **(5')**, it remains up to Darcy how overt he wishes to make his

subversion. If he delivers the line with a tone of sympathy and compassion, Elizabeth may miss

the subversive character of Darcy's utterance altogether, and be deceived into believing that Darcy

does, after all, agree with her. If he delivers the line relatively tonelessly, or with an air of studied

inscrutability, she may suspect that he is somehow not in sympathy with the thrust of her

rhetorical question (as a fully "sincere" interpretation of his utterance would suggest), but that he

is further hinting at something that he does not want, for some reason, to make fully explicit. If he

delivers the line in a heavily sarcastic or "dripping" tone, Elizabeth will immediately recognize

that she is being openly mocked, and will also probably be less sensitive to the possibility that

there is some deeper meaning, relative to which Darcy's utterance might be regarded as sincere.[24]

This fact—that irony affords its users the possibility of modulating the transparency of the full range of their communicative intentions—is a central feature of some of irony's most important, most interesting uses (and the special risks that attend these).

§2.2.2 Modulating Transparency

During a period in my life when I occasionally ran across the Ben Franklin Bridge (which connects Philadelphia, PA to Camden, NJ) I took an interest in a billboard along the route. The billboard displayed a stereotypical advertisement for a law firm specializing in personal injury, but the ad was (apparently) defaced by a giant, spray-painted message that read, "WE HATE THEM!!!" Careful inspection of the lettering, however, disclosed too precise a uniformity of script to have been produced by a human hand: the "defacement" was itself a part of the ad. I thought the ad was therefore likely to be attempting a kind of irony—speaking loosely, at least. The audience was intended to recognize that the sign had not, in fact, been defaced, and that the feigned negativity was a playful way to capture the viewers' attention, and to emphasize the firm's success at defeating now-embittered opponents. I wondered how effective the ad was.

Then, one day, I noticed the billboard had been modified. Beneath the declaration of hatred, a new, spray-painted signature had appeared, claiming credit for the vandalism. It read: "♡, The insurance companies!" The firm had apparently concluded that their ironic intent wasn't sufficiently recognizable, and that too many viewers were associating the apparent defacement as a manifestation of widespread dislike for unscrupulous attorneys, rather than a clever allusion to

the firm's winning ways. But I had the strong impression that the added signature had vitiated whatever ironic effect the sign once had.

Presume, for the moment, that my reaction was apt. Why would the addition of the signature line have spoiled the ad's irony? Well, the addition's (intended) effect was to weaken the surface plausibility of the misleading appearance by which the original irony was produced: the added signature makes it much clearer that no vandalism has actually occurred, and signals much more blatantly the ad's intended message, that the firm regularly takes on and defeats the insurance companies. So the suggestion is that even if an ironist produces a misleading appearance with the ultimate intention that it be so recognized (hence, transparent), the ironic effect is strengthened if the misleading appearance retains some fairly robust degree of plausibility, and suffers in the converse case. Presumably, this is why the advertisement's designers began with a subtler version of the ad. A skillful ironist strikes just the right balance between misleadingness and transparency—keeping us always on the edge of our interpretive seats, but never toppling us for good. In other circumstances, she modulates the transparency of her irony to include some potential audience members in her ironic intent, but to exclude others.

The dissociative echo model of irony production and comprehension does a poor job of accounting for this intuitive feature of irony.[25] According to the echoic version of the insincerity condition, the creation of a misleading appearance has no role to play in the mechanism by which irony is produced and interpreted. Instead, irony's mechanism is thoroughly overt: "When a whole utterance such as (1) is interpretively used [read: mentioned], the question of whether the speaker has obeyed a maxim, norm or convention of literal truthfulness should not arise"—more

general cases of mention (or, in Wilson's more recent work, "interpretive use"), like free indirect discourse or nonironic echoes (as in **(4)**), aren't in any way enhanced by modulating their transparency, shading toward deception. With mention, there should be no discrepancy between what a speaker appears to be doing and what she *is* doing; at most, there's ambiguity (I may fail to notice the interpretive mode and thereby misinterpret the speaker). But this sort of ambiguity seems only to hobble, rather than enhance, mention/interpretive use.

The deficiency is particularly stark when we compare echo theory to the competitor theory I'll turn to next: pretense. The pretense theory's version of the insincerity condition, according to which a speaker *pretends* to execute a speech act, rather than echoing its content, provides a natural explanation for the strengthening effects of "garden path" style ironies, in which a speaker's ironic intentions become transparent gradually: whatever delights pretense affords are made possible by the pretenders' (players and audience) ability to treat them, in certain ways, as real. It stands to reason, then, that these delights might in many cases be amplified by a temporary lack of awareness that any pretending is going on at all. Furthermore, unlike echoing, mentioning, or expressing a dissociative attitude, pretense remains itself in both opaque and transparent contexts: Alec Leamas (had he really existed) could have pretend to betray English intelligence to the East Germans (ideally, opaque); and Richard Burton could pretend to be Leamas in the movie *The Spy Who Came in from the Cold* (transparent). Pretenders, therefore, have both the ability and the motivation to modulate the transparency of their pretense in the way characteristic of much ironic use. Echoers and dissociators have neither.

These two related deficiencies of echo theory indicate that the *general* mechanism by which irony operates is better explained as a distinctive outgrowth, or exploitation, of the "twofold" character of interpretation, rather than as an allusive echo. This commitment, at any rate, is the first plank of the positive account of irony I offer in this chapter: *irony arises in twofold interpretive situations*.

§2.2.3 "Twofold" Interpretation

Though I discussed twofoldness in chapter 1, when I presented my general theory of interpretation, I'll rehearse (and expand) some of the basic points here. I borrow the term "twofold" from Richard Wollheim, who introduces it in his essay "Seeing-as, Seeing-in, and Pictorial Representation." There, Wollheim argues against E.H. Gombrich's contention that, in order to properly see a scene represented in a painting, we need to suspend our awareness of the features of the painting *as a representational medium*. In contrast, Wollheim contends that the kind of seeing required to perceive scenes in pictorial representation has what he calls a "twofold" character:

> That the seeing appropriate to representations permits simultaneous attention to what is represented and to the representation, to the object and to the medium, and therefore instantiates seeing-in rather than seeing-as, follows from a stronger thesis which is true of representations. The stronger thesis is that, if I look at a representation as a representation, then it is not just permitted to, but required of, me that I attend simultaneously to object and medium.[26]

Though Wollheim uses "representation" in the passage above as an abbreviation for "pictorial representation," his basic insight can be applied to a more general set of interpretive situations (including many cases of linguistic interpretation)—and it's only in in situations like these that

132

intentionally produced ironies can arise. Consider what the discussion above showed was necessary to fully appreciate an irony like **(5')**: the interpreter needed to notice that a single utterance admitted of two possible interpretations, and that these related to one another in an interesting way. She needed, that is, to attend not only to the utterance's straightforward communicative uses, but also to features that are particular to it *as a medium* of communication: in this case, the fact that it is the sort of thing that admits of different interpretations depending on the context in which it is employed.

Sperber and Wilson's mention theory advanced our understanding of verbal irony so significantly because it captured an important part of this insight. As Sperber had it: "Mention theory is based on an extension of the logical notion of mention. When an expression is *mentioned*—as opposed to being *used*—it refers to itself."[27] But mention theory encounters difficulty because the mechanism by which it achieves this insight is categorical, and so closes off the possibility of simultaneously attending to the utterances in the usual way. Yet it's just the interplay between these two ways of regarding the utterance that constitutes experiencing its irony. Twofold interpretation improves on mention because it not only allows for, but requires this sensitivity to the interplay between medium and meaning.

Nonetheless, we need to be careful not to *equate* twofold interpretation with irony. Consider the way Wollheim qualifies the role of dissociation in his account of "seeing in":

> I have, however, spoken of 'relative dissociation', and advisedly. For the artist who (as we have seen) exploits twofoldness to build up analogies and correspondences between the medium and the object of representation cannot be thought content to leave the two visual experiences in such a way that one merely floats above the other. He must be

concerned to return one experience to the other. Indeed he constantly seeks an ever more intimate *rapport* between the two experiences...[28]

By its talk of a *rapport* between the experience of medium and the experience of object, the passage suggests the way that twofoldness opens up the possibility of such experiences taking on an ironic quality: namely when, instead of establishing a rapport, the experience of the medium and the experience of the object (or, in my vernacular, meaning) seem to be at odds, or even to undermine one another.[29]

Here's an example that illustrates both the way that an irony can arise from a twofold interpretive situation, and the fact that twofoldness alone is a necessary but insufficient condition for irony. I suggested above that irony arises from twofold interpretive situations when, instead of having a *rapport*, a representation and its object interact disharmoniously, so that one undermines the other. This is the kind of experience we encounter in

(6) David Hockney's *Picture Emphasizing Stillness* [figure 1]

which tries, in part, to take itself as its object.

Figure 1: David Hockney's "Picture Emphasizing Stillness". The text inserted between the leopard and the conversing humans reads "THEY ARE PERFECTLY SAFE THIS IS A STILL".

Between the leopard and the unsuspecting human, Hockney has inserted a string of text that

reads, "They are perfectly safe: this is a still." The effect is to disrupt the usual (relative) harmony

of the twofold perceptual experience by creating confusion about the boundary between the

medium and the object of the pictorial representation. By explicitly inviting the viewer to attend

to an aspect of the representation that, in normal circumstances, would constitute a fact about the

medium of the representation, rather than its object (namely, "stillness," which we would

normally regard as true of the painting *as* medium, but *not* true of the state of affairs it depicts),

Hockney calls into question what the painting's *actual* object is. If it were, in fact, the state of

affairs that we'd "see in" the picture had Hockney not inserted the text that he did, then he neither

would nor could coherently remark, within the picture, on its stillness—because this undermines

the effect. It doesn't undermine the effect merely by drawing our attention to the picture's

stillness; this, as we've learned from Wollheim, is a fact about the medium to which we *already*

necessarily attend in the kind of twofold perception appropriate to viewing pictorial

representations. What undermines the effect is the *overt invitation* to attend to an aspect of the

representation that, were the artist actually attempting to depict a state of affairs in which a

leopard is pouncing on two unsuspecting humans, would be part of the medium rather than the

object of representation. This upsets our conception of the (implied) artist's intentions, which,

according to Wollheim, ground the norm that is distinctive to perceiving pictorial representation.

Because we can't square our conception of the artist's intention (given the presence of the textual

comment) with what otherwise would have struck us as the picture's object of representation,

we're led to recast our interpretation of what the picture's object of representation really is. As

Alexander Sturgis reads the painting, Hockney's goal is to dramatize, or actually to *depict*, the

ironic predicament painters face when they depict motion in a static medium:

> Depicting a figure falling in space, although inevitably suggesting motion at the same
> time, can also draw attention to the stillness of the depiction itself – and indeed to the
> limitations of the still image in conveying movement. The faster the movement the more
> aware one is of its freezing, for we never actually see plummeting objects suspended in
> midair. The point was neatly made by David Hockney in one of the paintings he made in
> the early 1960s...[30]

Let's take account. Despite depicting a commentary on the enterprise of depicting motion in a

still-life, if **(6)** is, in some sense, a "pretend" painting, it is no less a real painting for that. **(6)**

requires the same twofold attention to medium and object it ironizes, or the irony would dissolve. Unless we simultaneously hold in mind the object of representation (and, in this case, that includes what would be *its* object, since *it* is also a representation) with the medium of representation (which, among other functions, is what signals to us the artist's intentions), we cannot *see in* the painting what Sturgis does.

§2.3 Pretense Theory

I mentioned pretense in my analysis of each of the last two examples, which suggests that an account of irony's mechanism couched in pretense might succeed where dissociative echo falls short. Though the two theories are rivals, they share several significant features. Like echo theory, pretense theories are generally expressivist (so that ironists *primarily* express attitudes, rather than communicate content); also like echo, pretense theories regard ironists as adopting a dissociative stance toward the content of their utterances. The two diverge in their account of the manner in which this dissociative attitude is effected, along the lines that their names suggest. According to pretense theory, in its simplest form, ironists *pretend* to use the propositions they utter, rather than echoing them.

Herbert Clark and Richard Gerrig, who first presented pretense theory as an explicit rival to mention theory, take their inspiration from H.W. Fowler's treatment, in *A Dictionary of Modern English Usage*:

> Irony is a form of utterance that postulates a double audience, consisting of one party that hearing shall hear and shall not understand, and another party that, when more is meant than meets the ear, is aware both of that more and of the outsiders' incomprehension. [It] may be defined as the use of words intended to convey one meaning to the uninitiated

part of the audience and another to the initiated, the delight of it lying in the secret intimacy set up between the latter and the speaker.[31]

According to Clark and Gerrig, a person (S) speaks ironically to an audience (A) by pretending to be someone else (S') who *sincerely* uses the proposition S uttered to an audience (A'), who understands the utterance literally. S''s utterance should be "in one way or another, patently uninformed or injudicious," and S intends A to recognize this fact, along with A''s risible failure to recognize the inadequacy of S''s statement. Hence S expresses to A his derogatory attitude to the proposition he utters ironically.[32]

This way of understanding irony certainly does a better job accounting for **(5')**. When Darcy ironically agrees with Elizabeth's assessment of Wickham's misfortunes, Clark and Gerrig would say he only *pretends* to agree—a fact that a sufficiently astute and well-informed audience member (as we imagined Bingly being) would appreciate, recognizing that Darcy is in fact subverting the assertion he pretends to agree with, and hence recovering the irony.

Additionally, Clark and Gerrig make an explicit appeal to the continuity desideratum to argue for the superiority of pretense to echoic mention, noting that "the rhetorical device of irony… is just one of several types of irony." (They follow Fowler in listing dramatic irony and irony of fate as the others.) Echo, they contend, "doesn't allow for the resemblance among the three types of irony," but their pretense theory does: all three kinds (verbal, dramatic, and irony of fate) are supposedly linked by Fowler's stipulation of the double-audience. For example, the dramatic irony in *Oedipus Rex* consists in Oedipus's saying things with an intelligible but

superficial meaning to his fellow characters, but with a deeper meaning related to his impending

doom to the theater audience.[33]

Gregory Currie has proposed an improvement on Clark and Gerrig's version of pretense

theory, replacing their emphasis on a "double audience" with the simpler idea that an ironist

pretends to occupy a more limited perspective than she in fact does. Irony is then generated by

the contrast between the pretended and the actual perspective:

> The pretence one engages in with irony is partly one of behavior; one pretends to be
> doing something which one is not doing: speaking seriously and assertively, seriously
> asking a question, seriously expressing a distaste. But the pretence that is fundamental to
> irony is not a pretence of doing; it's a pretence of being. In pretending to assert or
> whatever, one pretends to be a certain kind of person – a person with a restricted or
> otherwise defective view of the world or some part of it.[34]

On Currie's picture, an ironist's utterance should "be an indication that he or she is pretending to

have a limited or otherwise defective perspective," which targets that perspective (or one that

closely resembles it) as unreasonable.[35] The mechanical aspect of this picture evokes my own

analysis of **(5')**, in which I described Darcy as "affecting an attitude of mildly unsympathetic

inscrutability, meant less to implicate to Elizabeth that she is being contradicted, and more to

suggest that there is a broader perspective on the situation (Darcy's) according to which

Elizabeth's rhetorical question may be taken to mean something quite different than she intends."

Like Clark and Gerrig, Currie counts pretense's ability to answer to the continuity

question as among its virtues. Though he calls his theory "modest," he credits it with a wider

reach than Sperber and Wilson allow theirs (at least, explicitly): when Currie declares that "irony

is pretense" in the title of his essay, he means the definition to capture more than just verbal

ironies. "Indeed," he points out, "irony does not even need language; I may stagger back in a parody of horrified distaste when confronted by an austerely elegant Sung vase, ironically expressing my rejection of your ludicrously demanding aesthetic standards."[36] Though Currie doesn't promise that a pretense of limited perspective can capture irony in *all* of its forms (situational irony in particular he regards as a "bloated" category), he explicitly endorses continuity as a desideratum: "an account like my own will gain in strength if it can be shown that it sheds some light on these other kinds [of irony]."[37]

I agree that pretense theory does a better job characterizing the mechanism of many verbal ironies than echo does; I agree that it gives a better account of continuity as well. Nonetheless, pretense's answers to both the mechanism and continuity questions remain incomplete. As it stands, pretense theory yields watered-down accounts of dramatic and situational irony which, because they overlook one of irony's central features, are far too inclusive; and it fails to capture an interesting and important form of verbal irony, in which no pretense is necessary. In the process of illustrating each of these points, I'll introduce the two final features of my own account of irony: dynamism and reversal.

§3 Situational and Dramatic Irony

This is Currie's attempt to characterize dramatic irony in terms of pretense of perspective:

> Electra mourns over the ashes she thinks are those of Orestes; Malvolio's hope is based on a letter we know is fake. Through credulousness, wishful thinking, or merely because they don't know the facts, characters in fiction often have a limited perspective, highlighted by the more inclusive view granted to the audience. So 'dramatic irony' turns out to be simply the fictional representation of the kind of contrast between perspectives that one pretends to be the victim of when one speaks ironically.[38]

Just as dramatic irony is derivative, in this way, of "irony proper" (the name Currie gives to cases covered directly by his pretense theory), situational ironies are derivative of dramatic ironies: they're just actual situations that, were they represented in a performance or literature, would count as dramatic ironies. They "mirror" the situations portrayed in dramatic irony, "except that there need be no onlooker aware of the disparity at the time." Alleged situational ironies whose incongruity does not "implicate" any limited perspective are only ironic in a "bloated" sense that Currie seems to think we are better off discarding.[39]

Neither of these attempts to characterize dramatic and situational irony does its subject justice. Consider first Currie's account of situational irony, since it's especially feeble. In situations rightly described as ironic, according to Currie, "if we had known differently, we would have acted differently," and this fact implicates a contrast between the actors' limited perspective and a theoretical broader perspective that would be available to a knowing audience, were this situation being portrayed in a play or a novel. Examples of situational irony Currie gives are robbing a bank whose vaults were emptied the day before, or quitting your job the day before you were to receive a promotion. ("The representation of any of this in performance of literature would count as dramatic irony.") [40] This account makes it sound as if any instantiation of Robert Burns's old dictum that "the best laid schemes o' mice an' men / Gang aft agley" should count as a situational irony, and that any fictional portrayal of such a situation in which the reader is aware, in advance, that something is amiss, should count as dramatic irony. But there's more to irony than that.

A contrast between two famous wedding mishaps illustrates the point. Though both cases are drawn from literature (broadly construed), I'll treat them, for the moment, merely as hypothetical situations: things that could have happened. (I address the additional complexities of literary and dramatic irony later.) The first wedding is from Alanis Morissette's infamous song, "Ironic":

(7) *It's like rain / on your wedding day!*

The second is Oedipus's unfortunate wedding, in which

(8) having already unwittingly killed his father, Oedipus fulfills the Delphic oracle's prophecy by marrying his mother—despite having spent his life trying to avoid precisely this fate.[41]

Oedipus and the unfortunate couple imagined in (7) share the property that Currie says characterizes situational (and dramatic) irony: in both cases, had the principals known better, they would have acted differently. In both cases, we can posit an interpreter inhabiting a broader perspective, who could thereby recognize that both Oedipus and the couple in (7) are on course to fall victim to the Burns Dictum. But I think many readers will join me in recognizing that (8) has a far more robust ironic quality than (8)—many people, in fact, deny that (7) has any genuine claim to irony at all.[42] Pinpointing the discrepancy that Currie's treatment ignores brings another of irony's key features into relief.

Oedipus's situation is robustly ironic in a way that mere rain on a wedding day isn't because Oedipus *brings about his own undoing by the actions he chooses on the basis of his intention to evade that very undoing.* A **fatalist** reading of *Oedipus* reminds us that we are the playthings of the gods, and that any attempt to avoid our fate, no matter how well-conceived and

142

thoughtfully executed, is ultimately doomed. A **moralistic** reading reminds us that we'll be punished for our crimes.[43] But an *ironic* reading of the play reminds us that we are the authors of our own undoing—a theme to which the play returns again and again. Had Oedipus never sought knowledge of his identity and fate at the Delphic Oracle, he'd never have learned about the terrible prophecy concerning himself and his family. Had he never made the explicit effort to avoid that fate, he'd never have fled his adopted home and set out for Thebes; neither would he have met and killed his true father on that journey, nor defeated the Sphinx and won his true mother's hand in marriage. If, once king, he hadn't made it his personal mission to unravel the previous king's murder… well, you get the picture. The play's ironies are thickest when, for example, Oedipus promises personally to punish the murderer of Thebes's previous king, once his identity has been discovered. (Oedipus himself, of course, is the murderer, though he's unaware of that fact when he vows retribution.)

Unlike in **(8)**, couple **(7)**'s desire to avoid rain on their wedding day probably wasn't a significant causal factor in their scheduling their wedding for a day on which, unfortunately, it happened to rain. So, in **(7)**, all we have is (i) a general scenario that generates or is characterized by certain norms, goals, or expectations (e.g. when we imagine a stereotypical wedding, perhaps the sun is shining and the weather is pleasant; when couples get married, they generally hope for nice weather, etc.), (ii) which expectations (etc.) happen to get violated in the (purportedly) ironic course of events. (We can add, for Currie's sake, that had couple **(7)** known in advance that it would rain on the day on which they happened to have in fact scheduled their wedding, they might not have scheduled it for that day—but the point seems tellingly otiose.) In **(8)**, by contrast,

we have more than an incongruity between an expectation or an intention and an outcome: the *reversal* Oedipus suffers is *dynamic*, so that the ultimate incongruity between his intention (which plays roughly the structural role of the *representation* in Wollheim's twofold interpretation) and its outcome (which plays the part of *object*) is *brought about* by their interaction: Oedipus forms the intention to *avoid* his fate, and precisely by the actions he takes on the basis of that intention he ultimately fulfills it. This is possible because Oedipus's intention to avoid his fate describes or represents a certain picture of his reality (the one Oedipus would like to bring about) while simultaneously being a *feature of* his actual reality, which, through Oedipus, can interact causally with other features of that actual reality. Because our practical perspectives are severely limited, these interactions aren't always in harmony with the state of affairs the intention represents.

Currie's account of dramatic irony is deficient in a way analogous to his account of situational irony. According to Currie, dramatic irony occurs when dramatic tension is created by a contrast between the wider perspective enjoyed by a play's audience and the limited perspective possessed by the fictional character, which is encompassed by the audience's wider perspective. But, again, this account captures too much. The contrast between a reader's wider perspective and a fictional character's more limited one grounds a great variety of significant literary devices, not merely dramatic irony. Foreshadowing often exploits this contrast, for instance. So do suspense-generating scenes in horror movies, in which we, the audience, know that the monster is in the room, but the protagonist doesn't. Here the contrast between the character's perspective and the audience's generates real dramatic tension, but such scenes are not necessarily (or even usually) instances of dramatic irony.

144

Compare these to a case of genuine literary irony: Mark Twain's portrayal of Pap (Huck's drunkard father) in *Adventures of Huckleberry Finn*. In the scene below, Huck, who is essentially being held hostage by his abusive father, recounts a long rant he has heard Pap deliver upon Pap's return from a trip into town:

(9) *I got the things all up to the cabin, and then it was about dark. While I was cooking supper the old man took a swig or two and got sort of warmed up, and went to ripping again. He had been drunk over in town, and laid in the gutter all night, and he was a sight to look at. A body would a thought he was Adam -- he was just all mud. Whenever his liquor begun to work he most always went for the govment. This time he says:*

"Call this a govment! why, just look at it and see what it's like. Here's the law a-standing ready to take a man's son away from him—a man's own son, which he has had all the trouble and all the anxiety and all the expense of raising. Yes, just as that man has got that son raised at last, and ready to go to work and begin to do suthin' for him *and give* him *a rest, the law up and goes for him. And they call that govment! That ain't all, nuther. The law backs that old Judge Thatcher up and helps him to keep me out o' my property. Here's what the law does: The law takes a man worth six thousand dollars and up'ards, and jams him into an old trap of a cabin like this, and lets him go round in clothes that ain't fitten for a hog. They call that govment! A man can't get his rights in a govment like this. Sometimes I've a mighty notion to just leave the country for good and all. Yes, and I told 'em so; I told old Thatcher so to his face. Lots of 'em heard me, and can tell what I said. Says I, for two cents I'd leave the blamed country and never come a-near it agin. Them's the very words. I says look at my hat -- if you call it a hat -- but the lid raises up and the rest of it goes down till it's below my chin, and then it ain't rightly a hat at all, but more like my head was shoved up through a jint o' stove-pipe. Look at it, says I—such a hat for me to wear—one of the wealthiest men in this town if I could git my rights.*

"Oh, yes, this is a wonderful govment, wonderful. Why, looky here. There was a free nigger there from Ohio—a mulatter, most as white as a white man. He had the whitest shirt on you ever see, too, and the shiniest hat; and there ain't a man in that town that's got as fine clothes as what he had; and he had a gold watch and chain, and a silver-headed cane – the awfulest old gray-headed nabob in the State. And what do you think? They said he was a p'fessor in a college, and could talk all kinds of languages, and knowed everything. And that ain't the wust. They said he could vote when he was at home. Well, that let me out. Thinks I, what is the country a-coming to? It was 'lection day, and I was just about to go and vote myself if I warn't too drunk to get there; but when they told me there was a State in this country where they'd let that nigger vote, I drawed out. I says I'll never vote agin. Them's the very words I said; they all heard me; and the country may rot for all me—I'll never vote agin as long as I live. And to see the cool way of that nigger—why, he wouldn't a

145

give me the road if I hadn't shoved him out o' the way. I says to the people, why ain't this nigger put up at auction and sold?—that's what I want to know. And what do you reckon they said? Why, they said he couldn't be sold till he'd been in the State six months, and he hadn't been there that long yet. There, now—that's a specimen. They call that a govment that can't sell a free nigger till he's been in the State six months. Here's a govment that calls itself a govment, and lets on to be a govment, and thinks it is a govment, and yet's got to set stock-still for six whole months before it can take a hold of a prowling, thieving, infernal, white-shirted free nigger, and—"

Pap was agoing on so he never noticed where his old limber legs was taking him to, so he went head over heels over the tub of salt pork and barked both shins, and the rest of his speech was all the hottest kind of language—mostly hove at the nigger and the govment, though he give the tub some, too, all along, here and there. He hopped around the cabin considerable, first on one leg and then on the other, holding first one shin and then the other one, and at last he let out with his left foot all of a sudden and fetched the tub a rattling kick. But it warn't good judgment, because that was the boot that had a couple of his toes leaking out of the front end of it; so now he raised a howl that fairly made a body's hair raise, and down he went in the dirt, and rolled there, and held his toes; and the cussing he done then laid over anything he had ever done previous. He said so his own self afterwards. He had heard old Sowberry Hagan in his best days, and he said it laid over him, too; but I reckon that was sort of piling it on, maybe.[44]

If it were our task to make the irony in **(9)** clear to a group of high school English students, how should we proceed? We might first draw their attention to the fact that, over the course of his rant, Pap again and again *says* one thing, while simultaneously (in so doing, even) *shows* its opposite. He blathers on about being a responsible father while in the act of delivering a drunken rant to his son, whom he has kidnapped and is holding prisoner. He rails against the government's effectiveness while simultaneously making it clear to the reader that the laws that protect his son from him are just and necessary. He rages against the injustice of a black man's being allowed to vote while simultaneously (though unwittingly) demonstrating that the black man in question is a far more responsible and thoughtful citizen than Pap is. He attempts to exact revenge on a tub of salt pork by dealing it a mighty kick, but succeeds only in injuring himself. Though these events have a more comical quality than **(9)**, they share its essentially ironic form:

146

it's not just that Pap doesn't know he's a worthless citizen and a disgraceful father, and we do; it's

that his very protestations against these characterizations provide us with strong reasons to accept

them. His assertions have propositional content – they make claims about the world (in

Wittgenstein's vocabulary, they *say*) – but they themselves are also facts in the world, whose full

import is not exhausted by their propositional content (they *show*). And here, what they show

undermines what they say.

We should also notice that, to explain this primary irony, we haven't yet had to address

the fact that **(9)** is a scene in a work of fiction. It would be just as ironic, in the sense covered so

far, if this scene were real, rather than fictional. Nor have we needed to appeal to any notion of

pretense in our explanation. But the scene *is* of course fictional, and for this reason (among

others) the irony here is literary. Because Pap is a fictional character, we encounter his actions and

utterances with an awareness that he and they are the inventions of an author, and, given that

awareness, we have access to a second (higher-order) interpretive level, on which we attempt to

discern the *author's* communicative intentions, as distinct from the character's.[45] We therefore

regard the irony described in the paragraph above as *intended*. But that doesn't constitute some

additional irony, over and above the one noted in the previous paragraph. Instead, it's a particular

way of regarding the irony we've already identified.

Nonetheless, recognizing **(9)** as the invention of an author directs our attention to

important issues that the analysis above hasn't yet addressed. Because **(9)** is the invention of an

author, intentionality functions on multiple levels. On top of the intentions I've already appealed

to in my analyses of **(8)** and **(9)**—Oedipus's intention to evade his fate, and Pap's various

communicative intentions—we also have Twain's higher-order intentions to consider. From this perspective, the conclusions we readers are apt to draw about Pap on the basis of his rant (those that undermine what he *says* in the course of the rant) look like sayings in their own right; they are the conclusions Twain *intends* us to draw by the way he's crafted the scene. I'll call ironies that are intentionally authored in this way *artificial ironies* to distinguish them from mere situational ironies (as **(8)** would be, if it had actually happened).[46] Though situational ironies are generated, in part, by the intentions of their protagonists, the irony itself is unintentional. But in artificial ironies, the irony itself is intended: they therefore pose a version of the mechanism question that situational ironies don't. How does Twain *intentionally manufacture* the irony in **(9)**?

Now appealing to pretense makes sense. It's no stretch to describe fiction writers as pretending to adopt a more limited perspective than they have; and the affinity to pretense is even clearer in plays, where actors actually play the characters and utter the lines.[47] In the Twain example, the kind of mechanistic use of pretense I've been discussing is, in a superficial but still important sense, doubled, because Pap's rant is recounted to us second-hand, through Huck's narration. We can consider Pap's remarks as akin to pretended assertions whose author is Twain, but they come to us retold by Huck (whose assertions about Pap's assertions can also be interpreted as pretended assertions whose author is Twain). So Twain has two layers of pretense at his disposal, and therefore a second order of perspective discrepancies he can exploit. He uses this double pretense masterfully throughout the novel, often to create or heighten ironies. **(9)** is no exception: the ironic reversals suffered by Pap are so obvious, and progress so rapidly, each more blatant than the last, that they'd be at risk of the kind of vitiating effect I described in §2.2.2,

were they not filtered by Huck's straightforward, guileless retelling. Because Huck himself seems either unaware of or unconcerned with the ironies of Pap's rant (and takes seriously issues the reader finds absurd, like Pap's boasts about the relative legendariness of his string of cursing), Twain creates for himself the aesthetic license to contrive a scene far more ridiculous than he otherwise could have without losing all sense of subtlety.

What, finally, can we glean from this extended discussion of **(9)**, and the comparison between **(9)** and **(8)** (keeping in mind that, for current purposes, we're treating **(8)** as a situational irony)? We should notice that we have neither dispensed with the importance of pretense in describing some ironies, nor have we found it to be a necessary component of *all* ironies. In this sense, Currie's "irony is pretense" declaration is an overstatement, but one that inscribes a more limited truth—namely, that pretense is central to the mechanism by which many ironies are produced. Nor is Currie correct to suggest that the pretense mechanism is the key to understanding the continuity between verbal irony and its dramatic and situational siblings. That way of rendering the relationship leads, as we've seen, to unsatisfying accounts of literary and situational irony.

But put continuity concerns aside for the moment. Do pretense theories at least provide a complete accounting of *verbal* irony? I don't think they do. In the next section, I explain why.

§4 Irony and Sincerity

Pretense plays a central role in the production of many artificial ironies, but it isn't necessary for irony—even ironies that are intentionally produced. Two final examples: The first is drawn from a

review, by David Foster Wallace, of Joseph Frank's multi-volume biography of Dostoevsky. The review has a peculiar structure: most of the sentences are just the kind you'd expect in a review, but those sentences are occasionally interrupted, without warning or explanation, by an ongoing series of self-searching philosophical questions, set off from the main text by line breaks and before-and-after asterisks. For example, immediately after opining that the fourth volume of Frank's biography is so good that it "ensures Frank's status as the definitive literary biographer of one of the best fiction writers ever," Wallace writes,

> **Am I a good person? Deep down, do I really even want to be a good person, or do I only want to *seem* like a good person so that people (including myself) will approve of me? Is there a difference? How do I ever actually know whether I'm bullshitting myself, morally speaking?**

The review then proceeds as if nothing odd has happened: "In a way, Frank's books aren't literary biographies at all, at least not in the way that Ellmann's book on Joyce and Bate's on Keats are."[48] In the review's final phase, Wallace (who was, at the time of writing, in the process of finishing his great novel, *Infinite Jest*), reflects on how Frank's rendering of Dostoevsky is relevant to readers and writers in contemporary America:

(10) Upon his finishing Frank's books, though, I think that any serious American reader/writing will find himself driven to think hard about what exactly it is that makes many of the novelists of our own place and time look so thematically shallow and lightweight, so morally impoverished, in comparison to Gogol or Dostoevsky (or even to lesser lights like Lermontov and Turgenev). Frank's bio prompts us to ask ourselves why we seem to require of our art an ironic distance from deep convictions or desperate questions, so that contemporary writers have to either make jokes of them or else try to work them in under cover of some formal trick like intertextual quotation or incongruous juxtaposition, sticking the really urgent stuff inside asterisks as part of some multivalent defamiliarization-flourish or some such shit.[49]

Many readers will join me in registering a strong sense of irony at the passage's conclusion. Yet Wallace apparently means what he says – there doesn't seem to be any tricky figure of speech at work, here, threatening to confound our usual theories of meaning, or requiring some specialized interpretive mechanism. Nor is there any fictional character involved, into whose mouth a novelist or playwright has slipped some *double entendre*. Nonetheless, this is doubtlessly an artificial irony: Wallace *intends* the reader to recognize that the most vital point he hopes the essay will make is undermined by the form in which he's trying to make it. If the reader misses this irony, in fact, she will have missed something central to the essay's meaning: Wallace is *showing*, rather than merely *saying*, the harrowing predicament he finds himself in as a novelist in contemporary America.

Here's a punchier example, with a similar moral. As we've already noticed, Paul Grice posits a "Cooperative Principle" that governs conversation, comprised of a set of maxims and sub-maxims that conversationalists expect one another to observe. One such maxim (in Grice's taxonomy, the third maxim in the category of Manner) instructs speakers to be efficient in expressing their meaning. This is how Grice states the maxim's content:

(11)   "Be brief (avoid unnecessary prolixity)."[50]

When I first read the sentence, in a graduate seminar, I assumed that Grice had blundered into an almost cosmically inapt pleonasm. Prolixity is, after all, unnecessary by definition, and the entire parenthetical that exhorts us to avoid it seems to add little or nothing of significance to the imperative "Be brief." Fortunately, before I had a chance to publicize my clever observation, the professor teaching the seminar remarked that Grice was indulging in a bit of irony in (11). But

151

here, again, Grice certainly *means* what he says; nonetheless, many readers will join my old professor in recognizing Grice's ironic intent.

Still, you might think pretense theorists can account for both of these examples. Following Currie, we might say of **(11)**: yes, Grice *really does* mean that we ought to avoid unnecessary prolixity in conversation; he doesn't merely *pretend* to make that proscription. Nevertheless, **(11)** is ironic because Grice *pretends* to inhabit a limited perspective, from which he is unaware of the redundancy in his own expression. Similarly, we might say of **(10)**: of course Wallace doesn't *pretend* to disdain the practice of "sticking the really urgent stuff inside asterisks as part of some multivalent defamiliarization-flourish or some such shit"—he *does* disdain it, and that is what he sincerely intends to communicate. Nevertheless, his concluding statement is ironic because he *pretends* to occupy a limited perspective from which he is unaware of his own hypocrisy—as if he has forgotten that he himself has just been engaging in the practice that he now bemoans, or has failed to realize that his own review is an obvious example of the sort of thing he's now disdaining.

Maybe these interpretations appear plausible, but neither is right. We can certainly tell a story about **(11)** in which Grice is pretending in the way Currie might suggest, but no such story is necessary to recognize **(11)**'s irony. Suppose that Grice *isn't* pretending to be unaware of the fact that his way of stating the third maxim of Manner is itself (unnecessarily) prolix: our sense of **(11)**'s irony is undiminished. That's because pretense isn't essential to **(11)**'s irony, it's ancillary. The essence of **(11)**'s irony is that it expresses a rule for proper expression, but it is itself an

152

expression that violates that rule. The essence of its irony, in other words, is (twofold) dynamic reversal.

A forced reading of **(10)**'s irony is worse than extraneous: it actively disfigures Wallace's meaning. Wallace isn't merely *pretending* to find himself in a predicament in which the conventions of his culture and his chosen art form can no longer authentically express the themes that make art important: he *is* in this predicament. The effect of the irony in the passage is to invest Wallace's ensnarement with a kind of experiential vitality: as a reader I appreciate it more viscerally than I would have if Wallace had phrased the point as I paraphrased it above. The pretense interpretation, by positing insincerity in either Wallace's being in the predicament, or in his awareness of it, hollows out this sense of authenticity. The picture of irony as dynamic reversal does it justice.

A final possibility: if ironies like **(10)** and **(11)** can't be accounted for by echo theory, or pretense theory, or some other unified mechanistic account of irony production and comprehension, perhaps we shouldn't count them as verbal ironies at all. Maybe, for example, we could class them with situational ironies that just *happen* to manifest themselves in words. I'll address this general methodological strategy in a bit more detail below, but, for now, here's one reason I don't think it's viable. Had the situation depicted in **(9)** really happened, many of Pap's statements would fit this description (situational ironies manifesting themselves in words), but the suggestion can't fully characterize **(10)** or **(11)** because both are artificial ironies. Unlike Pap (or Oedipus), Wallace and Grice are *using* irony; they're not merely subject to it. Wallace and Grice are *being ironic*. Pap is not. But neither Wallace nor Grice is pretending.

153

I've highlighted some of the limitations of the pretense theory of irony, but I don't want to undersell its merits. Theories like Clark and Gerrig's and Currie's do a good job with what they regard as their primary task: providing an account of a figure of speech by which irony is often employed in language (and other forms of communication). And, though pretense theory doesn't provide the "center of gravity" for irony that Currie promises, it does draw our attention to important points of commonality between the way this figure of speech operates in verbal irony and the way that irony is employed in drama—just as its proponents suggest.

§5 The Virtues of Dynamic Reversal

I'll conclude by highlighting the virtues of understanding irony as dynamic reversal. But in the process, I'd like to consider more carefully a methodological concern I touched on above. In my choice of examples, haven't limited myself to the relatively traditional, simple instances of irony as a figure of speech that currently dominate the literature on verbal irony. Further, one of my chief objections to the pretense account was its failure to provide the convincing explanation of irony's continuity that its advocates often promise. But what if the ambition to account for irony's several kinds with a single theory is misguided, and the right approach to the continuity question isn't to answer it, but to discard it?

§5.1 What Should Theorists of Irony Study?

Deirdre Wilson has given the most compelling articulation of that sort of view. Faced with a slew of counterexamples to the echo theory, Wilson has insisted, not altogether unfairly, that the full discursive range of the concept of "irony" isn't a proper target of interest for pscyholinguists and

cognitive science. Responding to Raymond Gibbs' observation that irony's diversity "poses an important challenge for cognitive science theories of irony," Wilson asks, "but why assume that the goal of a cognitive science theory of irony should be to capture the very broad and vague extension of the ordinary language sense of the term?" Instead, she thinks, "The goal of a theory is to identify mechanisms and see what range of phenomena they explain." Having identified the mechanism central to "typical" cases of verbal irony, further empirical studies should be directed toward illuminating these mechanisms, rather than muddying the water with fringe cases. When experimentalists include as legitimate cases of irony examples that don't exploit the favored mechanism, "propos[ing] to merge these into a general category of 'irony,'" they "detract from, rather than enhance, our understanding of how irony works."[51]

Wilson is right to insist on the danger of presuming that the categories of classical rhetoric will map neatly onto the cognitive structure of the mind—a distinctive mechanism for each trope. She's also right to insist that the study of such cognitive architecture is a central concern of cognitive science and pscyholinguistics. Her concern that the variety of cases counted as "ironic" in the growing experimental literature is unprincipled in ways that lead to confusion is well founded. But she's wrong to advocate for an approach to irony studies—in cognitive science or anywhere else—that makes the identification of irony as a legitimate object of study with any particular cognitive mechanism a matter of methodological stipulation. I'll give two reasons why this is so: one external to Wilson's professed concerns, one internal. In each case, the deficiency is remedied by understanding irony as dynamic reversal.

§5.1.1 Begging Questions

First, the internal reason. Suppose that Wilson is right (as I believe she is) that the empirical study of irony has drifted too far toward an indiscriminate ecumenicalism about irony at the level of the utterance. (In a passage Wilson quotes disapprovingly, for example, Raymond Gibbs notes that, because "irony is a single category of figurative language, but a variety of types," a corpus study he conducted on the use of irony among college students "did not even distinguish irony from sarcasm, hyperboles, [and] understatement."[52]) Methodological restrictivism presents itself as a principled solution: identify "core cases" of irony with some natural kind (a mechanism), and distinguish between the variety of cases Gibbs includes according to whether their production and comprehension employs the privileged mechanism.

The problem, of course, is that a methodological stipulation like this risks begging the question with respect to some of the central issues empirical studies hope to resolve. For example, a central question in the empirical literature is whether irony can be produced and interpreted by a variety of cognitive mechanisms, as Gibbs and Colston contend it can, or whether a single mechanism is essential to all instances of pragmatic irony, as Sperber and Wilson claim. Instead of making this question empirically tractable by providing a clear, shared concept of irony that can be used to design and interpret experiments, methodological restrictivism obviates the question as an empirical concern altogether, settling it instead by methodological stipulation. The same general dynamic applies to the frequent disputes about whether particular classes of utterance should be counted as "ironic" for the purpose of psycholinguistic experiments whose aim is to discern which mechanisms are involved in irony production. Wilson, for example, has

contended that praise-by-blame and jocularity commonly counted as instances of irony in empirical studies (e.g. Filippova & Astington, 2010,[53] and Glenwright & Pexman, 2010[54])—cases like

(12) You're really bad at lifting weights! [to someone who has just lifted an impressively large weight]
(13) I'm not all that good in the sack anyways, so you're not missing out. [uttered with the transparent intention of suggesting the opposite][55]

—shouldn't be counted as genuine cases of irony. If speakers in such cases aren't "echoing a manifest doubt or suspicion that someone had previously thought or expressed," they should be regarded as "non-ironic playfulness, banter, or teasing"—even in the cases where interpreters successfully recover the speaker's intended meaning in spite of the absence of an echo. Because "there is no evidence from Gibbs' discussion that his examples were echoic…including them in experimental studies of irony sheds no light on the mechanisms for irony comprehension."[56] At its extreme, mechanistic restrictivism begs rather than answers questions about irony's mechanism.

By contrast, the approach underlying the dynamic reversal account of irony provides the needed discriminatory power without begging questions. Identifying irony with a formal property of interpretive structures is neutral with respect to the question of which cognitive mechanisms can be used to produce or interpret such structures, but it offers a principle for distinguishing irony from related phenomena. Wilson is right to object to a lump inclusion of understatement in the category of irony, for example: some understatements are used to ironic effect, while others aren't. The dynamic reversal account provides the basis for distinguishing these classes:

understatements that *reverse* the figure's usual function, amplifying the content that understatement is conventionally used to downplay, are the properly ironic instances. Similarly, irony provides the resources for distinguishing, along a continuum, irony from sarcasm: the more conventional cues or blatant signaling serve to effect the reversal, rather than a twofold interaction between medium and meaning, the more likely we are to call an utterance sarcastic rather than ironic. (Few dispute that there is at least great overlap between these categories, though, so we place them on a continuum rather than in a dichotomy.) As we noted above, sarcasm becomes more itself the more overt it is, but with irony, the opposite is true.

§5.1.2 Why Does Irony Interest Us?

Now, the external reason. Mechanism restrictivism does a poor job of accounting for such a wide range of the manifestations of irony that interest us aesthetically, ethically, and historically, that it can't provide a good basis for interdisciplinary collaboration. If cognitive scientists and psycholinguists adopt a methodology that treats discourse-length irony as either spurious or ancillary to their endeavor, then artists, literary theorists, cultural critics, and historians will understandably be less inclined to regard the work in those fields as relevant to their own endeavors. For those of us who think, first, that an understanding of irony's potential range of use is an important step in developing good theories of its proper or elegant uses (its ethics and aesthetics); second, that understanding the psychological mechanisms by which irony can be produced and understood sheds light on the question of how irony can be used; and, third, that the study of irony at the level of the utterance has much to reveal about the mechanisms by which

158

irony can be produced and understood, the disciplinary balkanization that restrictivism

encourages is regrettable.

By providing a thorough, satisfying answer to the continuity question, the dynamic

reversal account of irony encourages the opposite trend, offering theorists of irony a point of

intersection they can use to learn profitably from one another. In addition, by acknowledging the

full breadth of irony's manifestations, the dynamic reversal account broadens our appreciation of

irony's possible uses, so that we can understand not only irony's power to undercut (as Twain

used it so effectively to reveal the ridiculous bankruptcy of Pap's racism), but also to express a

certain kind of experience more deeply and authentically (as Wallace used it to bring his artistic

predicament home to the readers he hoped to engage), or covertly (as, several commentators have

argued, Dmitri Shostakovic's symphonies did during Stalin's rule in the Soviet Union).[57]

§5.2 Irony from the Ground Up

Even theories with ambitions to account for continuity, like many versions of pretense, begin with

the presumption that irony should be understood on the model of a mechanism, and try to extend

their accounts from there: situational and dramatic irony are understood in terms designed

primarily to accommodate their dominant, figurative cousin. This is why imaginary but oddly

inert audience members or speakers are such familiar figures in the irony literature: irony requires

a commentator and an interpreter, and when none are to be found, they need to be invented.

The dynamic reversal account of irony suggests that this traditional approach has the

proper order of explanation reversed. It suggests that we should share Jonathan Lear's conviction

that the *experience* of irony, rather than its intentional production, or even its interpretation, is explanatorily fundamental.

The best way to see what I mean by this is to compare the way the dynamic reversal account comprehends irony with the way the CI account, from chapter 1, comprehends meaning. At first glance, the two strategies look congruent: meaning's varieties are comprehended according to their shared role in the structure of interpretation; irony is similarly united across various media by appeal to certain features of interpretation. Irony, as I've said, characterizes a class of interpretive structures that share a distinctive formal feature.

But careful attention to the discussion above shows my claim to have been incomplete, because, unlike meaning, irony isn't confined to *interpretive* structures. Though meaning is united by the structural role it plays in interpretation, irony is united by dynamic reversal, a formal property that obtains analogously in interpretive *and* practical structures. Consider the way I integrated **(8)**, the Oedipus myth (treated as a natural rather than dramatic irony) into the dynamic reversal account. Oedipus's intentions, I claimed, occupied the role of the medium in twofold interpretation, and the consequences of the actions he took on the basis of those intentions occupied the role of meaning. But these are *role* analogies, not identifications. Oedipus's flight to Thebes, his murder of Laius, his marriage to Jocasta—these aren't mere acts of interpretation. And though Oedipus needs to interpret these actions from the proper perspective to *appreciate* their irony, the irony isn't itself *effected by* his interpretation; *its* structure may not be ironic in the way interpreting Hockney's painting, or Wallace's review, is. Instead, it's the practical structure—the complex of Oedipus's intentions and the consequences that issue from

160

the actions they inspire—that's a paradigm case of dynamic reversal. The congruence between CI integration of meaning and the dynamic reversal integration of irony extends only to artificial ironies, which are, fundamentally, matters of meaning and interpretation. Artificial ironies intentionally occasion (or, in weaker forms, mimic) ironic experience, or ironic predicaments. We understood the irony in **(9)** (Pap's rant) *first* in terms of the relationship between his intentions (in this case, mostly communicative, but unironic, intentions) and their consequences, and only secondarily from the broader perspective of the ironist who produced the scene. We understand artificial ironies in terms of natural ironies, not the other way around.

§5.3 Concluding Thoughts: The Ethics and Aesthetics of Irony

Restructuring our understanding of the relationship between the kinds of irony along the lines advocated in **§5.2** has enormous implications for the ethical and aesthetic questions I raised at the outset. Explanations of irony that treat understand it fundamentally on the model of figurative language invariably privilege its artificial modes, and become preoccupied with the question of how the artifice is accomplished. They provide real insight into some of irony's ethical and aesthetic modes: the particularly vivid ways irony can be used to undercut illusion, or to subvert the forces of repression by manufacturing a distance between meaning and overt responsibility that can only be bridged by a degree of shared knowledge and sympathy; the dangers of emptiness, alienation, self-deception—of becoming a person, to paraphrase Proust's description of Swann, who always takes care to sterilize his opinions by using a "special intonation,

mechanical and ironic, as though he had put the phrase or word between inverted commas, and was anxious to disclaim any personal responsibility for it."[58]

But approaches preoccupied by the mechanisms by which irony is intentionally produced too often fail to appreciate the shared, fundamental feature of human experience out of which irony arises. As Socrates was eager to remind us: it's possible to occupy a limited perspective while retaining our awareness its limitations. Because recognizing the limitations of a perspective is not, unfortunately, to transcend them, self-conscious occupation of limited perspective doesn't need to be *pretended*, as Currie suggests it must: if Socrates is right, it's the human condition. Nonetheless, despite our predicament, we develop beliefs, convictions, and values; and on their basis we put forward claims, develop intentions, and take action: we "pretend" in the capacious sense that David Hume uses the term. But when we pretend in this way, we become subject to an odd and particularly poignant sort of reversal: sometimes, like Oedipus, we not only fail to achieve our goals, but by our very efforts we undermine them. We become the authors of our own undoing. This fact, I have argued, is irony's center of gravity, and so there is a sense in which Currie was nearly right to declare "irony is pretense" – just not precisely the sense in which he intended. Sincere modes of ironic modes of interpretation and expression enable us to understand and share these experiences of reversal: to be aware of the gaps between appearance and reality, to be patient with the unfamiliar, and to find some source of beauty, forgiveness, or, at least, humor, in the times when we are the agents of our own undoing.

162

[1] A representative instance of each of these categories: Language: Dan Sperber and Deirdre Wilson, "Irony and the Use-Mention Distinction," in *Radical Pragmatics*, ed. Peter Cole (New York: Academic Press, 1981), 295–318; Ethics: Jonathan Lear, *A Case for Irony* (Harvard University Press, 2011); Ancient: Gregory Vlastos, *Socrates, Ironist and Moral Philosopher* (Cornell Univ Pr, 1991); Metaphysics/Politics: Richard Rorty, *Contingency, Irony, and Solidarity*, First Edition (Cambridge University Press, 1989).

[2] James Wood, "The Digressionist," *The New Republic*, August 9, 2004, http://www.newrepublic.com/article/the-digressionist.

[3] Wayne C. Booth, *A Rhetoric of Irony* (University of Chicago Press, 1975), 2.

[4] Lear, *A Case for Irony*, 180.

[5] Sperber and Wilson, "Irony and the Use-Mention Distinction," 298.

[6] Norman Knox, *The Word Irony and Its Context: 1500-1755* (Duke University Press, 1961), 4.

[7] Christy Wampole, "How to Live without Irony," *New York Times*, November 18, 2012, sec. Review.

[8] Maryam Ronagh and Lawrence Souder, "The Ethics of Ironic Science in Its Search for Spoof," *Science and Engineering Ethics*, December 16, 2014, 1–13.

[9] Raymond W. Gibbs, "Irony in Talk among Friends," in *Irony in Language and Thought: A Cognitive Science Reader*, ed. Herbert L. Colston and Raymond W. Gibbs, 1 edition (Routledge, 2007).

[10] Deirdre Wilson, "Irony, Hyperbole, Jokes and Banter," in *Papers Dedicated to Jacques Moeschler*, ed. Cristina Grisot et al., 2014, www.unige.ch/lettres/linguistique/moeschler/Festschrift/Festschrift.php.

[11] Gibbs, "Irony in Talk among Friends," 341.

[12] Paul Grice, "Logic and Conversation," in *Studies in the Way of Words* (Harvard University Press, 1991), 26.

[13] Ibid., 34.

[14] Sperber and Wilson, "Irony and the Use-Mention Distinction," 301. Sperber and Wilson's points, though cogent, are overstated. In the first place, though Grice himself does not explicitly explore this possibility, he might be able to accommodate at least some of the cases in which the implicated ironic meaning does not appear to be an inversion of the conventional meaning by interpreting the ironic meaning as a higher-order implicature generated by a lower-order implicature, rather than directly by the conventional meaning of the uttered sentence (as, for example, when metaphors are used ironically). (This is essentially the technique Sperber and Wilson themselves use to account for the intuitively communicative functions of irony—as higher order implicatures generated by attitudinal expressions.) In the second, the charge that the meaning-inverted interpretation of "The weather is lovely today!" renders it totally uninformative in the circumstances in which it is uttered and successfully understood implies that "This weather is really foul," would, in the same circumstances, be equally uninformative. Yet we aren't mystified by the fact that people often do, in such circumstances, feel moved to make this sort of remark. The upshot of this observation is that, to the extent that the uninformativeness charge tells against the ironic remark being regarded as an instance of used propositional meaning, it also tells against the unironic, literal version being so regarded. So, to the extent that the uninformativeness charge should move us to treat the ironic utterance as a breed of meaning mention rather than use, we should also conclude (on Sperber and Wilson's theory) that the literal version is a kind of mention, rather than a use. This is at least an odd consequence of the strong uninformativeness argument. The reason, I think, is that we can use language to simultaneously assert a proposition *and* express an attitude about its content. But we should note that the uninformativeness

argument is far from the only motivation Sperber and Wilson offer for their theory of irony—it isn't as if their whole theory hangs on its force.

[15] Sperber and Wilson's attitude toward the continuity question is, if not inconsistent, at least difficult to decipher. On one hand, as their comment about black bile and atrabilious temperament indicates, their interest in verbal irony arises from more general interests in the idea of "figurative meaning," and how it ought to be conceived of within a larger psycholinguistic theory of meaning. From this perspective, questions about the continuity and integrity of our existing concept of irony, in all of its uses, might seem ancillary. On the other hand, they often dedicate space in their discussions to address continuity, and to tout their own theory's relative proficiency in this respect as another of its virtues Dan Sperber, "Verbal Irony: Pretense or Echoic Mention?," *Journal of Experimental Psychology: General* 113, no. 1 (1984): 133, doi:10.1037/0096-3445.113.1.130.. In more recent work, Wilson has more explicitly endorsed (not without good reason) a more restrictivist approach, limiting irony to pragmatic phenomena, and analyzed in terms of allegedly modular interpretive capacities. For a more detailed, critical discussion of this approach, see Robert Willison, "In Defense of an Ecumenical Approach to Irony," in *Irony in Language Use and Communication*, ed. Angeliki Athanasiadou and Herbert L. Colston (John Benjamins Publishing Company, 2017).

[16] Sperber and Wilson have updated their terminology in various ways over the increasingly long history of this influential account. They now favor "interpretive use" to refer to much of what they previously called "mention," and hence refer to irony as a kind of "echoic allusion" rather than "echoic mention." Nonetheless, the position itself has remained largely unchanged in its details, and, under the great influence of their first paper, the irony literature has generally preserved "echoic mention" as the name of their theory. See Chapter 6 of Deirdre Wilson and Dan Sperber, *Meaning and Relevance* (Cambridge University Press, 2012) for a brief discussion of the terminology change

[17] Sperber and Wilson, "Irony and the Use-Mention Distinction," 306.

[18] Ibid., 308.

[19] Jane Austen and Tony Tanner, *Pride and Prejudice*, Penguin Classics edition (London ; New York: Penguin Books, 2002), 187. Quoted in Sperber and Wilson, "Irony and the Use-Mention Distinction," 307.

[20] Sperber and Wilson, "Irony and the Use-Mention Distinction," 307.

[21] Ibid., 310.

[22] Herbert Clark and Richard Gerrig, Gregory Currie, and Hideki Haamoto have all accused Sperber and Wilson of stretching echo beyond recognition. See Herbert H. Clark and Richard J. Gerrig, "On the Pretense Theory of Irony," *Journal of Experimental Psychology: General* 113, no. 1 (1984): 121–26; Gregory Currie, "Why Irony Is Pretense," in *The Architecture of the Imagination: New Essays on Pretense, Possibility, and Fiction*, ed. Shaun Nichols, 1st ed. (Oxford University Press, USA, 2006); Hideki Hamamoto, "Irony from a Cognitive Perspective," in *Relevance Theory: Applications and Implications*, ed. Robyn Carston, Nam-Sun Song, and Seiji Uchida (John Benjamins Pub Co, 1998). #Add to this note my own belief that the terminological complaint isn't terribly important, and more a symptom of the pattern of interpretive uncharity understandably spurred by the pretense-echo dialectic, if time.

[23] Sperber, "Verbal Irony," 130.

[24] In fact, I think this scene as it actually appears in *Pride and Prejudice* is best interpreted along the lines suggested by the second of these three possibilities. Instead of simply attempting to ironically implicate that Wickham "has not been the victim of misfortunes" (as Sperber and Wilson's reading suggests), I think Darcy is affecting an attitude of unsympathetic inscrutability, meant less to implicate to Elizabeth that she is being contradicted, and more to suggest that there is a broader perspective on the situation (Darcy's)

according to which Elizabeth's rhetorical question may be taken to mean something quite different than she intends. Darcy has at least two related reasons for employing this technique: the first is that it would clearly be ungentlemanly and beneath him to openly reveal Wickham's indiscretions (but nonetheless a part of him must want Elizabeth to find out about them); and second, while people (and Austen characters, especially) deeply desire the attention of their love interests, they do not like to appear to others or to themselves as openly campaigning for it. Such an appearance compromises their sense of pride in the former case, and their confidence in the authenticity of the attention shown in the latter. This brief foray into literary interpretation is less digressive than it might seem, because it illustrates the fact that irony's potential uses are richer and more complex than the sort of straightforward ridicule that echo-and-dissociate model emphasizes.

[25] Interestingly (and unfortunately), despite the welter of empirical research conducted on irony in the last decades, no one (to my knowledge) has tested the hypothesis made in the body text—that distinctively ironic effects are stronger near the edge of transparency than they are in the distance. This seems to me a question worth empirical investigation, particularly because (as I'll argue above) it's a rare point of genuine asymmetry between the echo and pretense hypotheses about irony's mechanism. At any rate, despite having not yet been the subject of an empirical study, the idea that transparency lacking subtlety vitiates irony is a commonly expressed intuition. Grice, for instance, observed that "while one wants [ironic] pretense to be recognized as such, to announce it as pretense would spoil the effect." And Elisabeth Camp plausibly cites this intuition as one of a suite of features often taken to distinguish verbal irony from sarcasm. Paul Grice, "Further Notes on Logic and Coversation," in *Studies in the Way of Words* (Harvard University Press, 1991), 55; Elisabeth Camp, "Sarcasm, Pretense, and The Semantics/Pragmatics Distinction," *Noûs* 46, no. 4 (2012): 587–634.

[26] Richard Wollheim, "Seeing-as, Seeing-in, and Pictorial Representation," in *Art and Its Objects* (Cambridge University Press, 1980), 213.

[27] Sperber, "Verbal Irony," 131.

[28] Wollheim, "Seeing-as, Seeing-in, and Pictorial Representation," 224.

[29] Though I turn to them in more detail later, this paragraph also introduces the two featured of irony in the account's title ("dynamic reversal"). The *dynamism* condition is meant to generalize the stipulation, in Wollheim's account, that the two visual experiences can't be left so that one merely "floats above the other": instead, they must interact. The *reversal* condition specifies the nature of this interaction: instead of establishing a rapport, it must be disharmonious, or undermining.

[30] Alexander Sturgis, *Telling Time* (National Gallery Company Ltd, 2000), 43.

[31] H. W. Fowler, *A Dictionary of Modern English Usage: The Classic First Edition*, ed. David Crystal, Reissue (Oxford University Press, USA, 2009), 295.

[32] Clark and Gerrig, "On the Pretense Theory of Irony," 122.

[33] Ibid., 124.

[34] Currie, "Why Irony Is Pretense," 116.

[35] Ibid., 118.

[36] Ibid., 113.

[37] Ibid., 128.

[38] Ibid.

[39] Ibid., 129.

[40] Ibid.

[41] This is one of the more disturbing confirmations of David Foster Wallace's observation that "of course, you end up becoming yourself."

[42] Morissette's song is routinely pilloried for misrepresenting irony—more frequently than is fair, I think, given that it's a pop song, not a usage guide. A characteristic example, from a *The New York Times* review of one of Morissette's more recent albums: "Nine years after she wrote the unironic 'Ironic,' she has discovered actual irony." Jon Pareles, "MUSIC; The Solipsisters Sing Out Once Again - New York Times," *The New York Times*, May 16, 2004, http://www.nytimes.com/2004/05/16/arts/music-the-solipsisters-sing-out-once-again.html?pagewanted=all&src=pm.

[43] The moralism is pretty perverse by today's standards. The ancient Greek gods were largely sympathetic with Exodus 20:5 when it came to the sins of the fathers—excepting, of course, those cases in which the sinful fathers were their own. They therefore regarded Oedipus's fate as just retribution for his father's iniquities: Laius brought a curse on himself and his descendants by abducting and raping a prince whom he was supposed to be teaching to race chariots. On the other hand, taking road rage to the point of murder is probably more blameworthy by today's standards than it was in Oedipus's time, so maybe it all evens out and he deserved what he got, one way or another.

[44] Mark Twain, *Adventures of Huckleberry Finn*, First Edition (Dover Publications, 1994), 21.

[45] I hope the reader will forgive me, in this passage and any relevant others, if I don't qualify all of my language to indicate that I'm side-stepping the variety of the well-worn metaphysical questions about the status of the "authors" of fictions, e.g. whether they are actual, implied, or dead. Suffice to say that in the specific case of "Mark Twain," I'm fairly confident that the answer is: all three.

[46] I intend "artificial" in its non-perjorative sense: artificial ironies are called "artificial" because they are intentionally made, not because they are any less genuinely ironic than situational ironies.

[47] Currie draws out this connection by claiming that all the sentences in fiction are pretend assertions; and though I agree with him in noting the clear connection between pretense and fiction, I hesitate to follow him in the specifics of that formulation. Curry, "Why Irony Is Pretense," 120. Clark and Gerrig note the same clear connection, citing it as an advantage over echo theory. Clark and Gerrig, "On the Pretense Theory of Irony," 123-124.

[48] David Foster Wallace, "Joseph Frank's Dostoevsky," in *Consider the Lobster: And Other Essays*, First Edition (Little, Brown and Company, 2005), 257.

[49] Ibid., 271.

[50] Grice, "Logic and Conversation," 27.

[51] Wilson, "Irony, Hyperbole, Jokes and Banter," 3.

[52] Gibbs, "Irony in Talk among Friends," 356.

[53] Eva Filippova and Janet Wilde Astington, "Children's Understanding of Social-Cognitive and Social-Communicative Aspects of Discourse Irony," *Child Development* 81, no. 3 (June 2010): 913–28, doi:10.1111/j.1467-8624.2010.01442.x.

[54] Melanie Glenwright and Penny M. Pexman, "Development of Children's Ability to Distinguish Sarcasm and Verbal Irony*," *Journal of Child Language* 37, no. 02 (March 2010): 429–451, doi:10.1017/S0305000909009520.

[55] Gibbs, "Irony in Talk among Friends."

[56] Wilson, "Irony, Hyperbole, Jokes and Banter," 12–13.

[57] Jennifer Gerstel, "Irony, Deception, and Political Culture in the Works of Dmitri Shostakovich," *Mosaic: An Interdisciplinary Critical Journal* 32, no. 4 (1999): 35–51, doi:10.2307/44029848. As Shostakovich described the conditions in Leningrad at the debut of his Fifth Symphony, "Even before the war, in

Leningrad there probably wasn't a single family who hadn't lost someone, a father, a brother, or if not a relative, then a close friend, Everyone had someone to cry over, but you had to cry silently, under your blanket, so that no one would see. Everyone feared everyone else, and the sorrow oppressed and suffocated us." M. T. Anderson, *Symphony for the City of the Dead: Dmitri Shostakovich and the Siege of Leningrad*, 1st Ed edition (Somerville, Massachusetts: Candlewick, 2015), 134.

[58] Marcel Proust, *Swann's Way*, ed. Susanna Lee, Centenary ed. edition (New York: W. W. Norton & Company, 2013), 91. "What other life did he set apart," the narrator wonders, "for saying in all seriousness what he thought about things, for formulating judgments which he would not put between inverted commas?"