

Not a Matter of Interpretation

STEVEN KNAPP*
WALTER BENN MICHAELS**

Justice Scalia famously begins his dissent in *Johnson v. Santa Clara County* by describing the text of Title VII of the 1964 Civil Rights Act as written “[w]ith a clarity which, had it not proven so unavailing, one might well recommend as a model of statutory draftsmanship.”¹ It is the “unavailing,” of course, that gives the remark both its wit and its pathos. Scalia’s point is that, despite the clarity with which they wrote, the authors of the statute were unable to keep it from being misinterpreted. Although the text says it shall be “unlawful” to “refuse to hire . . . any individual . . . because of such individual’s race, color, religion, sex, or national origin,”² the majority decision makes it lawful for the county of Santa Clara to have refused to hire Paul Johnson because he was a man. Therefore, Scalia insists, the majority “completes the process of converting” Title VII “from a guarantee that race or sex will *not* be the basis for employment determinations, to a guarantee that it often *will*.”³ What the text says you cannot do, its interpreters now say you can do.

The question of whether Title VII forbids affirmative action is obviously an important legal question, but our interest here is in an interpretive problem rather than a legal one. Put in Scalia’s terms, the interpretive problem is the following: How can a text written with such

* Steven Knapp, The Johns Hopkins University.

** Walter Benn Michaels, University of Illinois at Chicago.

1. *Johnson v. Transp. Agency, Santa Clara County, Cal.*, 480 U.S. 616, 657 (1987) (Scalia, J., dissenting).

2. *Id.* (citing Title VII of the Civil Rights Act of 1964, 42 U.S.C. § 2000e–2(a) (1964)).

3. *Id.* at 658.

clarity nonetheless be misinterpreted? One of the few rules of statutory interpretation that everyone accepts, Scalia says in his influential contribution to *A Matter of Interpretation*, is the rule that “when the text of a statute is clear, that is the end of the matter.”⁴ But, Scalia thinks, the text here is a very model of clarity, and yet that has not been the end of the matter. Why not?

One possibility, of course, is that there really is no dispute over what the statute means, and Scalia himself may be pointing in this direction when he describes the way in which the same majority’s earlier decision in *Weber*⁵ thoroughly “rewrote . . . the statute it purported to construe.”⁶ If “purported” here means something like “pretended” (a meaning the dictionary sanctions, and dictionary definitions matter to those convinced by Scalia’s theory of interpretation), then Scalia may be suggesting that the majority is not so much misinterpreting Title VII as ignoring or replacing it. And if that is true, then of course there is no interpretive disagreement, and the authors’ efforts were unavailing only in the sense that they could not get the Court to obey their text not in the sense that they could not get the Court to understand it. The actions of the majority in ignoring the meaning of the statute may be ethically and/or politically problematic, but they are theoretically uninteresting. We do not need a theory of interpretation to explain what is going on when people make up new laws instead of interpreting the old ones. So if there is to be an interpretive issue here, we have to assume that the disagreement between Scalia and the majority in *Johnson* is not a disagreement about whether affirmative action is a good thing (although they may well disagree about that) but is instead a disagreement about the entirely different question of whether affirmative action is permitted by Title VII.

Perhaps, then, the meaning of the statute is not quite as clear as Scalia says it is. Indeed, it does seem that if the dispute here really is an interpretive one—if it really is a dispute over what the text means and not a disagreement about the right thing to do—the statute cannot possibly be as clear as Scalia says it is. A look at the statute suggests, however, that Scalia is at least in some sense right about its clarity. It is not in any obvious way ambiguous or vague. Again, if you looked up the words in the dictionary, you would not find surprising or confusing secondary definitions. So how can two different readers—in good faith (and for there to be a genuinely interpretive dispute we must, as we have just noted, assume good faith)—disagree? What are they disagreeing about?

4. Antonin Scalia, *Common-Law Courts in a Civil-Law System: The Role of United States Federal Courts in Interpreting the Constitution and Laws*, in *A MATTER OF INTERPRETATION: FEDERAL COURTS AND THE LAW* 16 (Amy Gutmann ed., 1997).

5. *United Steelworkers v. Weber*, 443 U.S. 193 (1979).

6. *Johnson*, 480 U.S. at 670 (Scalia, J., dissenting).

Scalia's answer emerges if we begin by asking what he calls "the basic question" of statutory interpretation: "What are we looking for when we construe a statute?"⁷ In light of this question, the problem with the majority justices' position is not that they cannot understand the plain meaning of the statute but that they are not looking for its plain meaning. They are rejecting the idea that the plain meaning of the text counts as the text's meaning and are looking instead for something else: what the legislature intended. The disagreement here, in other words, is a theoretical one—it is a dispute about what the proper object of interpretation should be, about whether it should be what Scalia calls "the objective indication of the words" (what the authors said), or whether it should be "the intent of the legislature" (what the authors meant).⁸

Scalia's main argument for what he calls "textualism" and against "intentionalism" is that intentionalism "is simply incompatible with democratic government."⁹ And, although we will return shortly to the question of whether that is true, we need to begin by asking a more fundamental question: Do textualism and intentionalism, as Scalia understands them, really give us two different objects of interpretation? The reason for asking this question is obvious. It has, in fact, been provided by Scalia himself in his description of the clarity with which Title VII was written, and in his implication that not all texts are so clearly written. What is the difference between authors who express themselves clearly and authors who do not? The seemingly obvious answer is that authors who write clearly are authors who make it easy for their readers to understand what they mean, while authors who do not

7. Scalia, *supra* note 4, at 16.

8. For our previous arguments in defense of intentionalism, see Steven Knapp & Walter Benn Michaels, *Against Theory*, 8 CRITICAL INQUIRY 723 (1982); Steven Knapp & Walter Benn Michaels, *A Reply to Our Critics*, 9 CRITICAL INQUIRY 790 (1983); Steven Knapp & Walter Benn Michaels, *A Reply to Richard Rorty: What Is Pragmatism?*, 11 CRITICAL INQUIRY 466 (1985) (reprinted along with preceding two items in *AGAINST THEORY: LITERARY STUDIES AND THE NEW PRAGMATISM* (W.J.T. Mitchell ed., 1985)); Steven Knapp & Walter Benn Michaels, *Against Theory 2: Hermeneutics and Deconstruction*, 14 CRITICAL INQUIRY 49 (1987); Steven Knapp & Walter Benn Michaels, *Intention, Identity, and the Constitution: A Response to David Hoy*, in *LEGAL HERMENEUTICS: HISTORY, THEORY, AND PRACTICE* 187 (Gregory Leyh ed., 1992); Steven Knapp & Walter Benn Michaels, *Reply to George Wilson*, 19 CRITICAL INQUIRY 186 (1992); Steven Knapp & Walter Benn Michaels, *Reply to John Searle*, 25 NEW LITERARY HIST. 669 (1994). See also Steven Knapp, *Practice, Purpose, and Interpretive Controversy*, in *PRAGMATISM IN LAW AND SOCIETY* 323 (Michael Brint & William Weaver eds., 1991).

9. Scalia, *supra* note 4, at 17.

write so clearly make it hard or even impossible for their readers to understand them. Therefore, when we read the texts of unclear authors, we are more likely to misunderstand them and to come up with incorrect interpretations. But what are the incorrect interpretations incorrect interpretations of?

Take Scalia's example of the meaning of the word "bay." "If you tell me, 'I took the boat out on the bay,'" he says, "I understand 'bay' to mean one thing; if you tell me, 'I put the saddle on the bay,' I understand it to mean something else."¹⁰ Thus, someone who understood "I took the boat out on the bay" to mean something like "I loaded the boat onto the back of the reddish-brown horse and took it for a ride" would have misunderstood the sentence. But what was misunderstood would not be what the word "bay" meant. The whole point of the example is that if you look the word up in the dictionary, you will see that it means a "reddish-brown" horse, and that it also means "a body of water partially enclosed by land" (not to mention "any of certain . . . trees or shrubs with aromatic foliage" and a surprising number of other things).¹¹

What you want to know, however, is not just what meanings the word "bay" can have according to the rules of English, but what meaning it has when used on this particular occasion by this particular speaker. And, of course, looking it up will not help you with that. If, in other words, you have misunderstood the sentence "I took the boat out on the bay," your problem is not that you do not know the many meanings of the word "bay." On the contrary, it is precisely because you do know at least two of those meanings that you do not know what meaning the word "bay" has in this sentence. This is to say, you have misunderstood not what the word means, but what the author meant. Indeed, the very idea of a clear text (a text written in such a way that it is not likely to be misunderstood), or rather the very idea of the difference between a clear and an unclear text, relies on the presumption that the object of interpretation (what the interpreter is trying to figure out) is the author's intention.

We can see this by simply imagining how we might respond to the sentence about the bay if we were not interested in what its speaker intended. If we felt we did not fully understand the meaning of the word "bay," we might look it up in the dictionary, and thereby be reminded of other possible meanings (for example, "a part of a building marked off by vertical elements"¹²). But, of course, the dictionary would not tell us

10. *Id.* at 26.

11. THE AMERICAN HERITAGE DICTIONARY OF THE ENGLISH LANGUAGE 154 (Joseph P. Picket et al. eds., 4th ed. 2000) [hereinafter AMERICAN HERITAGE].

12. *Id.*

which of the meanings listed was the right one. Indeed, the very existence of the list would make the question of which of the listed meanings was the right one incomprehensible; if the word “bay” means all those things, then it can not be right to choose just one of them. So we not only would not get any help in figuring out which of the possible meanings was the right one, we would not have any reason to look for the right one because we would not have any reason to think that one was more right than the others. The point here is not epistemological—it would be hard to figure out which one was right. The problem instead is ontological—they would all be right. They would all be right because if we thought that the meaning of the sentence was a function of what the words meant, rather than a function of what the author meant by the words, then every text would mean everything its words could mean in the language in which the text was written. None of the dictionary meanings would be merely possible; all would be actual.

Insofar, then, as the opposition between a clear text and an unclear text relies on the difference between a text’s being understood and its being misunderstood, that opposition requires a reference to the author’s intention. The only way a sentence like “I took the boat out on the bay” can be interpreted to mean “I went sailing” instead of “I used the horse to carry the dinghy,” is by reference to what its speaker means. And, again, the point here is not that if we do not look for the speaker’s intention, we will have no way of figuring out what the right meaning is; it is that if we are not looking for the speaker’s intention, we have no reason to think that any meaning should count as the right one.

A clear writer, then, is a writer who is good at making herself understood, a writer who chooses or arranges words in such a way that her reader will understand what she means by them. But to say this is just to say that a clear writer is a writer who is good not at making you understand what the words mean (the dictionary does that) but at making you understand what she intends. The sentence “I am heading out on the bay” clearly means one thing when spoken by someone sticking her foot in the stirrup and (just as clearly) means something else when spoken by someone stepping into the dinghy, and its meaning may not be clear at all when it is spoken by someone with one foot in the stirrup and the other foot on the boat. In any case, the difference between when it is clear and when it is not clear is entirely a difference between when it communicates what the author intends to communicate and when it does not. As soon as you start thinking of the meaning of a text as the sort of

thing that authors can try to express clearly, and of clarity as something they can either achieve or fail to achieve, then you have already committed yourself to the idea that what the text means is what the author meant by it. The Justice Scalia who believes that texts can be clearly or unclearly written is not a textualist but an intentionalist, and the choice he presents us with—the choice between what the authors said and what the authors meant—is a false choice.

This is why Scalia's explanation of how we figure out whether the "bay" is a horse or a body of water is mistaken. According to Scalia, what we do here is invoke the "principle that a word is given meaning by those around it."¹³ The mention of a boat makes the bay a body of water; the mention of a saddle makes it a horse. It is not the intent of the author that matters, according to Scalia, it is the "context."¹⁴ But why does the context matter? If you have just asked me how I am going to get the boat out of here, and, as I strap the dinghy onto the horse, I say, "I am taking the boat out on the bay," the context no doubt makes it clear what "bay" means. And even though one of the words around "bay" is "boat," it does not mean "body of water." The context is relevant because the context is evidence of what I mean. The word is not "given meaning by [the words] around it,"¹⁵ but the words around it may well provide evidence of what meaning the speaker is giving to it, as may the speaker's actions, the look on his face, the tone of his voice, etc.

With Title VII, or with almost any written text, we do not know much about the looks on peoples' faces, of course. And the words around the disputed ones will no doubt play a crucial role in helping us decide what we think the disputed words mean. But that is what the records of congressional debate, and all the other texts to which intentionalists have recourse, also do. Indeed, that is the only reason that the one source textualists really approve of—the dictionary—is at all relevant. A standard literary critical example of this point is the reference to his "vegetable love" by the speaker in Andrew Marvell's "To his coy Mistress."¹⁶ What does the word "vegetable" mean? Here is what the *American Heritage Dictionary*¹⁷ says about the word's history in its entry on "vegetable":

Andrew Marvell's "To his coy Mistress" contains many striking phrases and images, but perhaps most puzzling to modern readers is one in this promise from the speaker to his beloved: "*Had we but world enough, and time . . . /My*

13. Scalia, *supra* note 4, at 26.

14. Antonin Scalia, *Response, in A MATTER OF INTERPRETATION: FEDERAL COURTS AND THE LAW* 147 (Amy Gutmann ed., 1997).

15. *See id.* at 26.

16. ANDREW MARVELL, *To his coy Mistress, in ANDREW MARVELL* 24 (Frank Kermode & Keith Walker eds., 1990).

17. AMERICAN HERITAGE DICTIONARY, *supra* note 11.

vegetable love should grow/Vaster than empires and more slow.” One critic has playfully praised Marvell for his ability to make one “think of pumpkins and eternity in one breath,” but *vegetable* in this case is only indirectly related to edible plants. Here the word is used figuratively in the sense “having the property of life and growth, as does a plant,” a use based on an ancient religious and philosophical notion of the tripartite soul. As interpreted by the Scholastics, the *vegetative* soul was common to plants, animals, and humans; the *sensitive* soul was common to animals and humans; and the *rational* soul was found only in humans. “Vegetable love” is thus a love that grows, takes nourishment, and reproduces, although slowly. Marvell’s 17th-century use illustrates the original sense of *vegetable*, first recorded in the 15th century. In 1582 we find recorded for the first time the adjective use of *vegetable* familiar to us, “having to do with plants.” In a work of the same date appears the first instance of *vegetable* as a noun, meaning “a plant.” It is not until the 18th century that we find the noun and adjective used more restrictively to refer specifically to certain kinds of plants that are eaten.¹⁸

Set aside the dictionary’s own intentionalism (“Here the word is used figuratively”¹⁹—what are you making a decision about when you decide whether a word is being used figuratively or literally?). The relevant issue here is the description of the word “vegetable” as having an “original sense”—still operative, the dictionary tells us, in the seventeenth century—and a more “restrictive” and more modern one, not yet operative in Marvell’s time. So when we try to figure out what “vegetable” means in Marvell’s poem, do we want the meaning that a seventeenth-century dictionary would have given us—“growing”? Or the meaning that a twenty-first-century dictionary would give us—“edible plant-like”? Suppose we choose the seventeenth-century dictionary. How can we justify our choice to someone who prefers the twenty-first-century one? The obvious answer—the text was written in the seventeenth century, why should it matter what the word came to mean later—will just as obviously not work if we remain true to our textualism. Intentionalists will prefer the seventeenth-century dictionary because intentionalists will regard the definition it gives as evidence of what Marvell intended. But textualists do not care what Marvell intended.

Perhaps, as textualists, we could say that we are not interested in knowing what Marvell intended, we are just interested in knowing what language the text is in and finding out what “vegetable” means in that language. So what language is it in? The problem here, of course, is that the question of what language the text is in is a question about what language it was written in, and the answer to that question involves

18. *Id.* at 1906.

19. *Id.*

asking about Marvell again. Was he using seventeenth-century English and attaching the signified “growing” to the signifier “vegetable,” or was he using twenty-first-century English and attaching the signified “edible plant-like” to it? Even if, as textualists, we are interested in what the words mean in the language, and not in what Marvell meant by them, it turns out that the question of what language the words are in is really only a disguised version of the question of what Marvell meant by them. Thus, once again, the answer to Scalia’s question—what are we looking for when we construe a text²⁰—turns out to be (even for those who think they are looking for what the words mean instead of what the author means) the author’s intention.

Put more generally, we could say that the opposition between intentionalism, as a historical inquiry into what some author or authors did, and textualism, as a linguistic inquiry into the ways in which some word is defined, is a false opposition. We cannot even ask the linguistic question of how the word is defined until we have some account of what language the word is in, and our sense of what language the word is in is already a function of our understanding of an historical event: the production of the text. To see this, all we have to do is imagine two languages that look exactly alike—the only thing that distinguishes them is their semantics; the words in L1 (call it English) mean something different from the words in L2 (call it Schmenglish). The fundamental question we confront with any text written in either of these languages is the question of which language it is in, and that question can only be answered by an account of which set of semantic rules the author of the text was following.

And once we realize that the question of what we are looking for can only be answered by an account of which set of semantic rules the author was following, we realize that the English/Schmenglish hypothetical actually describes a fairly standard situation. Any time anybody uses a word to mean something even slightly different from the meaning it has in English, he or she is speaking some kind of Schmenglish, even if this Schmenglish differs from English not with respect to its entire semantics but only with respect to the one word in question. From that point of view it makes sense to say that the meaning of every text is governed by the semantic rules of the language in which it is written, but only because it also makes sense to say that what those rules are—what language it is and, therefore, what rules matter—is entirely up to the author. From this standpoint, judges inquiring into the records of legislative debates are just trying to figure out which language the text they are being asked to read is written in.

20. Scalia, *supra* note 4, at 16.

But this way of putting the point is, of course, precisely what makes textualists unhappy and the political motives of textualism relevant—the sense that only textualism corresponds to the needs of a democracy and makes possible the rule of law. How can we be expected to follow laws written in languages we do not understand? How can we be expected to understand texts that do not follow rules we know? It is important to recognize that this is a perfectly reasonable, but—from the standpoint of the theory of interpretation—completely empty question. If I speak in English to a group of people who only speak French, it is surely the case that I will have a very hard time being understood, and if I am saying something important to them and to me, it may well be disastrous that I did not use my less fluent but still serviceable French. But the fact that not a single person understood me or could plausibly have been expected to understand me in no way affects the meaning of my utterance. And, of course, the same thing is true if I speak not in English but in Schmenglish. The question of meaning and the question of communication are two different questions, as is obvious any time anyone says something that is understood by some people, but not by everyone. If you understand my utterance, I have succeeded in communicating with you; if your friend does not understand it, I have failed to communicate with him. But the meaning of my utterance, whatever it is, is the same and would be the same if both of you understood it, or if neither of you did. And if I am speaking Schmenglish, your insistence on applying the rules of English will be of no greater use to you than it would be if I were speaking French.

This does not mean that you cannot insist on treating my utterance as if it were in English. It just means that if you do, you are insisting on treating the actual historical event that took place—the use of some set of semantic rules to mean something—as if it were a different historical event—the use of some other set of semantic rules to mean something else. Treating texts as if they were produced through the use of a certain set of rules, regardless of whether they actually were, textualists are not choosing between what the author intended and what the text says. They are choosing between what the author actually intended and a fictional alternative, for example, what they think the author should have intended. Properly understood, then, textualism is a theory about what the text should have meant, and not about how it should be interpreted. When there is a discrepancy between the rules the authors actually followed and the rules textualists believe they ought to have followed,

textualists want to invoke the rules the authors ought to have followed. From the standpoint of the theory of interpretation, they are in the position of someone who believes that everyone ought to speak English and who, therefore, treats every utterance as if it were in English, even if the actual author of the utterance may sometimes be speaking French.

From this standpoint, it cannot exactly be true that the problem with textualism is that, like intentionalism, it “treats statutes as static texts and assumes that the meaning of a statute is fixed from the date of enactment.”²¹ The date of enactment may be relevant if we are interested in figuring out the actual language the author used, but it is by no means necessarily relevant to the question of what language we think the author *ought* to have used. That question is a question about us, about what we think. So textualism is not quite as “dreary” as the proponents of a more “dynamic” interpretation make it out to be. But it does not follow that what one of the most prominent of these proponents, William Eskridge, calls the “interpreter’s perspective”²² is relevant to legal (or any other kind of) interpretation. For one thing, as we have already seen, textualism achieves its (unwanted) dynamism only through a kind of arbitrary stipulation (treat the text as English whether or not it is English)—a stipulation that, if they acknowledged it, textualists would presumably give up. At which point, they would indeed become truly static and dreary because, as we have seen, they would become intentionalists. Furthermore, Eskridge’s own account of the relevance of perspective to interpretation—“[t]he interpreter’s perspective makes a difference in statutory interpretation: two different people acting in good faith often interpret the same text in different ways”²³—makes a point exactly opposite to the one he hopes to make.

Eskridge illustrates the point he hopes to make by citing his own experience in reading Charlotte Brontë’s *Jane Eyre*, first when young, and later as an adult, when, “inspired by scholarship concerning female friendships in the nineteenth century and my own identity as a gay man,” he came to “consider Jane’s and Helen’s relationship in romantic terms”²⁴—a thought, he says, that would have been “inconceivable” to him when younger.²⁵ The idea here is that the change in his own “view” of the text is also a change in the text’s “meaning,”²⁶ and that the difference between the young reader of Brontë and the older one

21. WILLIAM N. ESKRIDGE, JR., *DYNAMIC STATUTORY INTERPRETATION* 9 (1994).

22. *Id.* at 49.

23. *Id.*

24. *Id.* at 59–60.

25. *Id.* at 59.

26. *Id.*

exemplifies the situation in which the readers of both literary and legal texts “disagree with one another” about what the text means.²⁷

But how can this be true? How do people who read Bronte without reference to their experience as gay men (because they are not gay or they are not men) count as *disagreeing* with Eskridge? What are they disagreeing with him about? You and I are not disagreeing if *Jane Eyre* makes you think of the intensity of certain kinds of same-sex friendship, and makes me think of how constricted women’s lives were in the nineteenth century; we are just being made to think about different things. The fact that *Jane Eyre* makes you think of the former (subversion of gender roles) only counts against its making me think of the latter (the rigidity of gender roles) if you also believe that *Jane Eyre* ought to make me think of what it makes you think of, and that if it does not, I am not only interpreting the text but misinterpreting it. If, in other words, your experience makes you think of one thing and mine makes me think of other things, we differ, but we do not yet disagree. We only disagree when each of us thinks the other is not just different, but mistaken, and for the other to be mistaken, we must not only have different responses to the text; we must also have different beliefs about what the text means. And the minute we have different beliefs about what the text means, we have become intentionalists, because the author’s intention is the only thing there is for us to disagree about.

Why? We have already seen that we cannot disagree—at least for long—about the text’s plain meaning. That would just be a disagreement about what its words meant in English (or at least in the version of English Charlotte Bronte was writing in), and it would either be eliminated by consulting the relevant dictionary, or, once we analyzed what it was that made the relevant dictionary relevant, revealed as a disagreement about what she intended. And we have also seen that we cannot disagree about our responses to the text because if you, as a gay man, think about one thing when you read it, and I, as a straight man, think about something else, we are not disagreeing, we are just thinking about different things. So we do not disagree about what the words mean, and we cannot disagree about what they make us think of, but we certainly can disagree about what Charlotte Bronte meant by them. That is why the question of whether women really did have intense and eroticized (if not yet what we would call lesbian) friendships in the nineteenth century matters. If

27. *Id.* at 61.

they did, it gives us a reason to believe that *Jane Eyre* might plausibly be about such a friendship, which is to say a reason for believing that Bronte might have intended to explore the conditions that created such friendships. If they didn't, it doesn't.

Just as the question of whether women did have such friendships is a question about a matter of historical fact, so is the question of what Bronte intended, and thus of what *Jane Eyre* means. This is to say that its meaning is indeed "fixed from the date of enactment." But the fact that it is fixed does not, of course, mean that our accounts of it are fixed. On the contrary, we have every reason to believe that people will continue to come up with new interpretations of old texts, and that our arguments about the meanings of at least some of them will never be resolved. It should be clear, however, that the point of the intentionalism defended here is not to resolve such arguments, but to explain how they are possible, why they can count as arguments, and what they are arguments about. Our point is that the possibility of interpretive disagreement—disagreement about what a text means as opposed, for example, to disagreement about what we should do—is entirely dependent upon our treating the text as an expression of its author's intent.

In his contribution to this volume, Jeffrey Goldsworthy appears to grant that interpreters, insofar as they are interpreting the text and not doing something else with it, necessarily treat a text as an expression of its author's intent. According to Goldsworthy, however, one can accept this thesis without succumbing to what he calls "strong intentionalism," the thesis that a text means (if it means anything) what its author intended it to mean. Instead, he proposes a position of "moderate intentionalism," which adds to intentionalism the proviso that

[f]or a text to mean what its author intended it to mean, it is necessary that its intended audience be capable of ascertaining its intended meaning. This requires that the audience have access to persuasive evidence of that meaning, provided partly by the conventional meaning of the text and partly by other clues of the author's intention, such as the context in which it was written.²⁸

Certainly it is true that, for an author to *communicate* what he intends to communicate, he is better off following the conventions he intends to follow, and providing the evidence he intends to provide. But Goldsworthy goes further than that, arguing not only that the audience's understanding is determined by the evidence available to it but also that the text's *meaning* is determined by that evidence. This is why, according to Goldsworthy, "[i]f the audience does not have access to such evidence, then not only will the author's attempt to communicate his

28. Jeffrey Goldsworthy, *Moderate versus Strong Intentionalism: Knapp and Michaels Revisited*, 42 SAN DIEGO L. REV. 669, 670 (2005).

intention fail, but the text will also not mean what he intended it to mean.”²⁹ What, then, will it mean? Goldsworthy’s answer is the following: The text will mean whatever the “readily available” evidence happens to *suggest* that it means. “Moderate intentionalism is the thesis that the meaning of an utterance is the meaning which evidence readily available to its intended audience suggests that it was intended to mean.”³⁰ Or as he also writes, “[m]oderate intentionalists . . . hold that what someone *appears* to have said, in the light of all the evidence readily available to his intended audience, is what he *did* say.”³¹

It does not matter, then, if what an author appears to have intended—what the readily available clues suggest that she intended—differs from what she intended in fact, which of course will happen whenever the clues she intended to give are different from the clues her audience picks up. When it comes to meaning, according to Goldsworthy, appearance is reality. But now suppose that some members of her audience pick up one set of clues, while others pick up another set, so that the meaning, as it appears to some, is different from the meaning, as it appears to others. Which appearance counts?

While acknowledging that it is an “arguably counterintuitive consequence” of his thesis,³² Goldsworthy embraces the only possible answer—they both do:

For example, if I intend a remark made during a conversation with Bob and Joe to be a joke, Bob might understand this but Joe might not. Joe might take it to be an insult, not through any fault of his, but because I failed to appreciate that he does not know me as well as Bob, and therefore needs more clues of my jocular intention. According to the moderate intentionalist, my remark has two meanings: it means one thing in the case of Bob, and another in the case of Joe, because they do not have ready access to the same evidence of my intention.³³

In other words, what Joe, out of ignorance, takes the remark to mean is one of the things it actually does mean, because, on moderate intentionalism, a text or utterance means whatever it appears to mean in light of whatever evidence happens to be available to its audience, or indeed to any subset of that audience.³⁴ But of course there is no limit to the

29. *Id.* at 670.

30. *Id.* at 680.

31. *Id.* at 677.

32. *Id.* at 680.

33. *Id.*

34. It is not clear whether Goldsworthy believes that if he tells the joke and Joe does not get it, Joe is correct in thinking that Goldsworthy has insulted him. It is not

number of ways in which the available evidence can vary across members of the audience. The available evidence—the way things appear—will suggest to some (those who, like Joe, do not know any better) that Marvell meant by “vegetable” what “vegetable” means in twenty-first-century English, to others what it means in Schmenglish, to still others what it means in any of a potentially infinite variety of future successor-languages. And if all those meanings count as the right meanings, we are once again left without any way of understanding what people are up to when they think they are disagreeing about what a text means.

In short, moderate intentionalism has all the faults of textualism; indeed, precisely the same faults. It misses the main point of intentionalism: we can only make sense of what people are doing when they disagree with each other about the meaning of a text if we suppose that they regard the text as meaning what its author actually (and not just apparently) intended. From that point of view, moderate intentionalism turns out not to be a version of intentionalism at all.

If we want to make sense of interpretation, then, there is no getting away, even moderately, from the author’s intention. Language counts; linguistic rules matter. But we cannot even decide what language a text is written in—and therefore what rules matter, without deciding what language the author intended to write it in—and therefore what rules the author was following. And once we decide what rules the author was following, we cannot choose among the meanings those rules make possible without deciding which meaning the author actually intended.

Our examples so far have all involved texts with single authors. But there are many texts, including most legal texts, that are the products of multiple authors, sometimes working in close collaboration, but often separated by time, space, or institutional role. So if the object of interpretation is what the author intended, what happens when a text has more than one author? Its professed ability to answer that question is no doubt an abiding source of textualism’s appeal. According to the

clear, in other words, whether he thinks that just because Joe reasonably believes that he has been insulted, he has, in fact, been insulted. It is pretty clear, however, that a commitment to treating reasonable beliefs as true beliefs will have interesting consequences not merely for the practice of interpretation, but for life in general (because it will now be possible for contradictory beliefs to be true). But perhaps we can address what Goldsworthy is truly concerned about by noting that we might well want to hold people responsible for saying things that can be reasonably misinterpreted without also having to claim that the misinterpretation must be true. And it is precisely (and only) with respect to the question of responsibility that the criterion of the reasonable becomes relevant—if we think Joe’s misinterpretation is reasonable, then we may hold Goldsworthy responsible; if we think it is not, we probably will not. But if we think Joe is right—if we think what Goldsworthy said really was an insult—then it does not matter whether Joe’s interpretation is reasonable or not.

textualist, we do not have to worry about multiple authors because the text is the text, no matter who wrote it, and to interpret it is just to determine and apply its linguistic rules. But that answer, as we have seen, is no answer at all, because without referring to authorial intention, we cannot even tell what set of rules is the right one, let alone how to apply them.

What does happen, then, when we have reason to think we are dealing with more than one author? Of course, nothing prevents two or more authors from sharing the same intention, in which case there is clearly no problem in saying that the text was written in the language its multiple authors intended to write it in, and that what the text means is what its multiple authors intended it to mean in that language. We can sharpen the question by asking: What do we do when we have reason to think that the text has multiple authors whose intentions did not agree, who were writing in different languages, or who were writing in the same language but used different dictionary meanings of the same words?

This is an important question, but, for better or worse, intentionalism offers no more help in answering it than textualism does, although intentionalism perhaps helps to clarify what is going on when circumstances compel us to answer this question in the case of a particular text. Theoretically, there is only one option. In cases where we have reason to think that a text's authors were writing in different languages, or intending different meanings within a single language—either way, giving the signifiers different signifieds—the text has turned out to instantiate more than one object of interpretation, and we can only interpret it by choosing which of those objects we wish to interpret.

Suppose, for instance, that Marvell's "To his coy Mistress" had been written collaboratively by two different poets, Marvell and, say, Schmarvell. Imagine that they wrote at that moment in the history of English when "vegetable" was shifting from its older meaning of "growing" to its now-current meaning of "edible plant-like." Marvell (we somehow discover) intended the older meaning, while Schmarvell intended the newer one. It would be tempting under those circumstances to say that the word "vegetable," as it appears in Marvell/Schmarvell's poem, has *both* meanings. But that solution turns out to be just as mistaken as the textualist's claim that the word has whatever meaning the rules of the language impose on it. For in the example, as we have imagined it, neither author intended the word "vegetable" to have both meanings. And if it was not their intention that gave the word both

meanings, what did? Of course Marvell himself, in the actual case, could have intended both meanings (making “vegetable” a pun), and Marvell/Schmarvell, in our hypothetical case, could together have intended both meanings (again, producing the pun). In either of those cases, however, we would be dealing with a single authorial intention to mean two different things, not with two different authorial intentions. We would have a single object of interpretation with two different meanings, not two different objects of interpretation. But where the intentions differ, so do the objects of interpretation, and the interpreter cannot even begin without deciding which of those objects to interpret.

Again, consider a case in which the difference between intended meanings would be even sharper than the difference between the successive English meanings of “vegetable.” Two historians coauthor a book on the French Revolution. As they trade drafts back and forth, at some point one of them inserts the following string of signifiers: “In 1789, the streets of Paris were awash in canards.” They intentionally publish the book with these signifiers included. But they do not realize at the time that, by the signifier “canard,” one of them intends to signify (in English, as glossed by American Heritage³⁵) “unfounded or false, deliberately misleading stor[ies],” while the other intends to signify (in French, as glossed by Larousse) ducks (because she wants to convey that the streets were unhealthily cluttered with livestock). Would it make sense to argue that, although neither author intended this, the text is correctly interpreted as asserting that Paris was filled with baseless rumors and ducks? Or should we simply say that the two authors inadvertently attached two different signifieds to the same signifier, thereby inadvertently producing two different signs, with two unrelated meanings?³⁶

The point here is that while there is no object of interpretation without authorial intention, a single set of signifiers—and therefore, in one sense, a single text—can instantiate multiple objects of interpretation.³⁷

35. AMERICAN HERITAGE, *supra* note 11, at 270.

36. For a somewhat different use of the term “canard,” see Antonin Scalia, *Assorted Canards of Contemporary Legal Analysis*, 40 CASE W. RES. L. REV. 581 (1989–1990).

37. A single text in one sense because the term “text,” in debates of this kind, is sometimes used to mean a string of signifiers (just the marks without their meanings), and at other times—a string of signs (the marks and their meanings combined). This terminological difference is irrelevant to the present point, which is that multiple authors can give multiple meanings to a single string of signifiers. For the classic literary treatment of this phenomenon see JORGE LUIS BORGES, *Pierre Menard, Author of the Quixote*, in COLLECTED FICTIONS 88 (Andrew Hurley trans., Viking Penguin 1998) (1989). The story is presented as a postmortem review of Menard, a twentieth century French author of the *Symboliste* school whose main literary ambition was “to produce a number of pages which coincided—word for word and line for line—with those of

When applied to poetry, this point does not seem especially consequential. Not only do poets collaborate rather rarely, but it is also hard to see what harm would result from discovering two different poems where there seemed to be only one. When applied to the law, on the other hand, this point can be very consequential. Legal texts are often the products of multiple stages of drafting and ratification, not to mention the revision that some say occurs in judicial review. But nothing in the logic of interpretation itself can tell us which of those stages should count as the one that confers on the text the meaning we are trying to interpret when we try to obey the law or to adjudicate it.

We have argued that it does not make sense to wonder whether to interpret a text according to authorial intention or linguistic rules; there is never an option, because without authorial intention, we do not even know what rules to apply. But it can indeed make sense to wonder which author or authors' intention is the one that matters for a particular legal, institutional, or political purpose. Unfortunately, knowing that the first question is not a real one provides no help in answering the second question, which is. So we have answered Scalia's "basic question," but only by replacing it with another one. The basic question was the following: "What are we looking for when we construe a statute?"³⁸ The answer is (not only with statutes but with all texts) that we are looking for the author's intention. It is only if we are looking for the author's (or authors') intention that the problem of interpretive disagreement—the problem that interpretive theory is designed to solve—can arise. In the course of answering the basic question, however, we have also shown that some of the real problems have nothing to do with interpretive disagreement after all. If we are trying to decide whether "vegetable" means "growing" or "edible plant-like," we may not be trying to choose

Miguel de Cervantes." *Id.* at 91. He had no interest in composing a different version of the *Quixote* or in merely copying it, and abandoned as too easy the idea of somehow becoming Cervantes; he preferred "continuing to be Pierre Menard and coming to the *Quixote through the experiences of Pierre Menard.*" *Id.* In other words, his ambition was literally to produce from scratch the same set of signifiers but with a different meaning. Borges' fictional reviewer illustrates Menard's achievement by quoting identical strings of signifiers side by side and offering comments like the following: "The Cervantes text and the Menard text are verbally identical, but the second is almost infinitely richer. (More *ambiguous*, his detractors will say—but ambiguity is richness.)" *Id.* at 94. Whatever one thinks of his aesthetic judgments, Borges' reviewer has exactly the right theoretical take on what happens when a single string of signifiers is produced with two different authorial intentions.

38. Scalia, *supra* note 4, at 16.

between two different meanings of the same poem—we may be trying to decide between two different poems, Marvell’s and Schmarvell’s. And if we are trying to decide whether “canards” means “false stories” or “ducks,” we may not be trying to choose between two different meanings of the same sign, but between two different signs.

In the context of legal controversy, there are, no doubt, situations in which it is hard to tell what the authors of a law intended their signs to mean. In those situations, the problem we are trying to solve is an interpretive one, and solving that problem involves finding the best evidence we can of what the authors intended. But the harder problem may be to decide which among the various layers of authorship is the one that counts. And if that is the problem we are trying to solve—not about what the signs mean, but about whose signs are the ones that matter—then the solution, whatever it is, will not be a matter of interpretation.