

March 2007

March 2007

**Interpreting Legal Texts:  
What is, and What is not, Special about the Law**

by

**Scott Soames**

**School of Philosophy**

**University of Southern California**

**To be Presented at**

**An International Conference on**

***Law, Language, and Interpretation***

**University of Akureyri**

**Akureyri, Iceland**

**April 1-2, 2007**

How is the content of positive law related to its authoritative sources – including written constitutions, statutes and administrative rules, understood in light of the beliefs, intentions, and presuppositions of those who produced them, and those to whom they are addressed? Progress can, I think, be made in answering this question by seeing it as an instance of the more general question of what determines the contents of ordinary linguistic texts. I will, therefore, look at recent advances in semantics and pragmatics, and extract lessons for legal interpretation by tracking implications for different kinds of “hard cases” in the law.

I will call a legal case *genuinely hard* iff its (legally correct) outcome is not determined by all nonlegal and nonmoral facts plus the *linguistically-based content* of the relevant legal texts -- including everything asserted and conveyed therein. These cases divide into three types: (i) those in which the texts say too little to produce any result, (ii) those in which the texts are inconsistent, and thus generate contradictory results, and (iii) those in which the texts yield a single result, which is, paradoxically, legally incorrect. In a genuinely hard case, a correct outcome can be reached only by an innovative judicial decision – which, effectively, creates new law. In contrast, a case is *semantically hard* iff the meanings of the relevant legal texts, plus all nonlegal and nonmoral facts, fail to determine its (legally correct) outcome. The distinction between semantically, and genuinely, hard cases turns on the distinction between the semantic contents of legal texts and their complete, linguistically-based contents.

Appreciating this distinction involves grasping several important lessons. First, literal meaning is more austere, abstract, and less transparent than it is often taken to be. Thus, even if fidelity to the text is the touchstone of interpretation, a text’s meaning is often not readily identifiable. Second, meaning is not what faithful interpreters should be looking for anyway – since even when it is identified, it may fail to determine the text’s content. That content, which

encompasses everything conveyed or asserted by the text, often includes information that goes well beyond the semantic contents of the sentences involved. Typically, an agent produces a sentence in a context with a communicative goal and topic, a record of what has been supposed or established up to then, and assumptions about the beliefs and intentions of participants. This pragmatic information interacts with the semantic content of the sentence to add content to the discourse. In recent years, we have learned that the pragmatic determinants of this content are not minor add-ons to semantic content. Semantic content is often merely a vehicle for getting to pragmatically enriched content, and sometimes the semantic content of a sentence is not itself asserted, or even included in what the speaker is committed to. The semantic-cum-pragmatic information-generating process governing the routine interpretation of linguistic texts and performances may start with literal meaning, but it doesn't end there.

The import of this for judicial interpretation is that the slide from semantically hard cases to genuinely hard ones should be resisted. It's true that the meanings of legal texts, plus the facts of a case, often fail to determine its outcome. But this *shouldn't* be taken to show that *the content of the law* embodied in those texts doesn't determine the outcome, and *mustn't* be used to invite judicial legislation. Just as *what I say*, and commit myself to, by uttering a sentence, is often a function of more than its semantic content, so "what the law says," and is committed to, is often a function of more than the semantic contents of relevant legal texts. Just as you have no standing to reinterpret my remark to conform to your moral and political views, *simply because the meaning of my sentence doesn't fully determine the content of my remark*, so judges applying the law have no standing to reinterpret it, *simply because the linguistic meanings of the relevant legal texts don't fully determine the content of the law*. There are other principles at work filling the gap between sentence meanings and the contents of texts, legal or otherwise.

Finally, even when a case *is* genuinely hard -- because the full linguistic content of authoritative texts fails to determine a correct outcome -- there are further principles to consult that routinely guide the interpretation of incomplete, inconsistent, or otherwise defective linguistic materials. Applicable in both legal and nonlegal contexts, these principles constrain the resolution of hard cases, and limit the moral and political discretion of judges in making new law. In what follows I will illustrate these points.

### **The Austerity and Non-Transparency of Meaning**

One of the chief results of recent years is the non-descriptionality and indefinability of many terms once thought to be definable -- ordinary proper names and natural kind terms, among them. Although these expressions may carry descriptive information, that information doesn't define them. For example, virtually everyone who uses the words 'Bill Clinton' and 'Washington D.C.' knows that the former names a past president of the United States, and the latter names the capital of the country. Similarly for 'water' and 'gold'. One would be hard pressed to find speakers who didn't know that water is a clear, drinkable liquid that falls from the sky, or that gold is a valuable yellow metal. However, this information doesn't define the terms. A childhood friend of Clinton, marooned on a desert island for forty years, would, on returning, still understand the name, use it to refer to the same man we do, and successfully communicate with us in so doing -- despite not knowing what most people regard as the most important things about Clinton. Although users of a name standardly associate *some* information with it, this information varies from speaker to speaker, and virtually none of it is required for mastering the name. Thus, the only semantic content left for it is its referent -- which explains why it is neither an apriori nor a necessary truth that Bill Clinton is (if he exists) a former president of the United States, as it would be if the name were definable in terms of common knowledge about Clinton.

These points carry over to natural kind terms, which are really just names of kinds, rather than individuals. The “definition” of ‘water’ is not, *the clear, drinkable, liquid that falls from the sky, etc.*, but *That stuff*, said pointing at water. We can, of course, be pretty sure that a man in the desert asking for water, wants a drink, and that a farmer who checks for rain knows that it will water his crop. However, truths like *Water is a clear, drinkable liquid that falls from the sky* are neither necessary nor apriori, as they would be if ‘water’ were descriptively definable. Nor must such truths be known in order to understand the word. An unfortunate, whose only knowledge of water was of the cloudy stuff draining from a nearby laundry, could use and understand the word, based on his limited acquaintance with water, just as our modern day Robinson Crusoe could use, and understand, the name ‘Bill Clinton’ based on his acquaintance with Clinton.

To be acquainted with a natural kind is to be acquainted with its instances. What makes a kind *natural* is that observable properties of its instances share a common explanation. In introducing a general term for such a kind, we stipulate that it applies to certain encountered particulars, and that other particulars will count as instances iff they share whatever features causally explain the properties observed in the original sample. For ‘water’, ‘gold’, and ‘green’, these features turn out to be *having the molecular structure  $H_2O$* , *having atomic number 79*, and *having a certain spectral reflectance property*. These causal-explanatory features of the original samples determine the extension of the terms, even though they were initially unknown, and needn’t be known now in order to understand the terms. Thus the reference of such a term isn’t determined by satisfaction of descriptive conditions encoded in its meaning. It has no such meaning. Rather, its meaning is the kind it names, and its reference is determined by agreed upon instances, plus a special similarity relation, holding between instances of the kind, the specific content of which needn’t be known, even by competent speakers.

A modified version of this picture extends beyond natural kind terms to other indefinable terms, including Wittgenstein's example of the term 'game'. Although there is no one thing that all and only games share, there are paradigmatic examples of games, plus a similarity relation such that anything that bears it to something that clearly is a game itself counts as a game. In this case, however, the similarity relation is not a causal-explanatory one, but something more closely tied to speakers' interests and priorities.

The relevance of these points to the law is illustrated by two legal cases. The first is *PGA vs. Martin*, in which the question is whether the Americans with Disabilities Act, requires the Professional Golfers Association to waive its rule requiring all tournament competitors to walk the course. The Supreme Court ruled (i) that the act does apply to the PGA, and (ii) that the PGA is so required because the rule change needed to accommodate the disabled doesn't affect golf's essential, or fundamental, nature. Justice Scalia dissents, correctly, in my opinion, on (i), but incorrectly on (ii), which is the only part of the case I will discuss. His position seems to be that the question of what is essential to golf – beyond the rules codified by its authoritative rule-making body – is nonsensical. This, I believe, is an error.

The crucial points, as I see them, are these: First, *the word* 'golf', like 'game', can't be defined by specifying necessary and sufficient conditions for its application. Rather, its meaning is the activity it designates, which is a kind, or type (just as the Word Series is a kind or type), instances of which are sequences of events of a certain sort. Second, although *the activity golf* is "defined" or "constituted" by its rules, the sense in which this is so isn't that of defining a word. You can know what the word 'golf' means, as well as what the activity golf is, without knowing all the rules. Also, the meaning and the activity can both survive rule changes. The word 'baseball' didn't change meaning, or reference, when the American League adopted the designated

hitter rule – and the Red Sox didn't cease playing the game and start playing something else, when that happened. Similarly, ordinary golfers don't cease playing golf when the seasons change, and they switch from so-called summer to winter rules.

In what sense, then, is golf constituted by its rules? In the sense that each round of golf is made up of actions guided by the rules. Just as Admiral Nelson's ship *Victory* was constituted by the planks of wood, pieces of iron, sheets of canvass and strands of rope that made it up, so a round of golf is constituted by the rule-guided actions that make it up. The *Victory* was constituted by its parts, even though an inventory of its parts doesn't give the meaning of the name, and having precisely those parts isn't essential to the ship (since it could have existed with slightly different ones). *The activity golf* is similarly constituted by its rules, even though they don't give the meaning of the word, and having precisely those rules isn't essential to golf (since the rules can vary a bit, without the activity being lost)

Thus, there was nothing absurd about the Court's finding that one of the rules – namely that tournament players must walk the course – isn't essential to golf. Scalia's mistake was, I think, due to one or both of two errors. The metaphysical sense in which the rules of golf *constitute* the activity (without being essential to it) must not be confused with the linguistic sense in which certain conditions *define* a word (and hence are essential to its meaning). If one conflates these two, one may wrongly think that since the rules constitute golf, they are essential to what we mean by 'golf' and so cannot be changed without fundamentally altering our conception of the game. The second error is to run together the *correct* Wittgensteinian observation -- that when questions of proper application arise for words like 'golf,' there is no higher authority than its use by competent speakers -- with the *incorrect* idea that the PGA's authority over its tournaments makes

it's rules committee the arbiter of competence in determining the application of the word. These errors illustrate how confusion about meaning can adversely affect legal reasoning.

Another legal example illustrating the austerity and nontransparency of meaning is *Nix v. Hedden*. The case turns on whether tomatoes count as vegetables rather than fruits, and so are subject to a tariff on vegetables that excludes fruits. The issue is whether the meaning of the word “fruit” excludes all vegetables, while including the edible “ripened seed-bearing ovary of the plant”, as my dictionary puts it – thereby precluding tomatoes, peas, beans, eggplants, cucumbers, squash, and peppers from being vegetables. The court found otherwise, ruling that ‘vegetable’ and ‘fruit’ have no specialized meaning in commerce, and so must be understood in their ordinary meaning, according to which tomatoes, along with the others, *are* vegetables.

*Botanically speaking*, tomatoes are fruit of a vine, just as are cucumbers, squashes, beans and peas. But *in the common language of the people* ... all these are vegetables, which are grown in kitchen gardens, and which, whether eaten cooked or raw, are, like potatoes, carrots, parsnips turnips, beets, cauliflower, cabbage, celery and lettuce, usually served at dinner, with, or after the soup, fish or meats which constitute the principal part of the repast, and not, like fruits generally, as desert. [*my italics*]

In my view, the court's decision was correct, and its reasoning defensible. It correctly notes that the term ‘fruit’ has a botanical meaning encompassing any edible “fruit of the vine,” containing the seed-producing part of the plant. On this meaning, its a natural kind term the extension of which includes tomatoes, beans and peas. Less obviously, but still plausibly, the court suggests that ‘fruit’ also has an ordinary meaning in which it contrasts with ‘vegetable’, both of which being “defined” by paradigmatic instances – including, one imagines, oranges, lemons, plums, and melons, in one case, and carrots, cucumbers, lettuce, and tomatoes, in the other. The way to think of these ordinary meanings is, of course, not in terms of verbal definitions, but along the lines suggested by Wittgenstein's example of games. The end result is a plausible analysis of



the statutory language, from which the court's decision follows, given the reasonable assumption that the words are to be taken in their ordinary meaning.

Suppose however -- for the sake of argument -- that this view of language is incorrect, because 'fruit', like other natural kind terms, has only a single "scientific" meaning -- the "fruit of the vine" one. Suppose further that vegetables are, by definition, edible parts of plants that aren't fruits. On this view, the common classification of beans, peas, peppers, and tomatoes as vegetables is not evidence of linguistic ambiguity, but a mistake based on ignorance (on a par with ignorance of the molecular structure of water). It would, therefore, follow that a literal interpretation of the statutory language excludes tomatoes from the tax. *Would it also follow that the court's decision was legally incorrect -- or, if not incorrect, correct only by virtue of a justifiable substitution the court's superior wisdom for that of the legislature?* I don't think so. However, to explain why, I must turn to my second linguistic lesson for legal interpretation.

### **The Gap between the Content of a Text and the Meaning of its Sentences**

The main point is that the semantic content of a sentence doesn't always determine what is asserted and conveyed by literal uses of it. Sometimes more than the semantic content is asserted or conveyed, and sometimes the semantic content isn't asserted at all. Since *the content of the law* includes everything asserted and conveyed in adopting the relevant legal texts, meaning is sometimes merely a guide to interpretation, to be supplemented by other things. Consider, for example, definite descriptions, [the F]. The basic *semantic* fact about these expressions is that a sentence [The F is G] expresses a truth iff there is a single thing satisfying F (in the relevant circumstance), and that thing also satisfies G. The basic *pragmatic* fact is that these descriptions are often used referentially, to focus attention on a particular individual, about whom the speaker makes a statement.

These points are illustrated by (1a), the literal meaning of which is given by (1b).<sup>1</sup>

1a. The man in the corner drinking champagne is a famous philosopher.

b. [the x: x is a man & x is in the corner & x is drinking champagne] x is a famous philosopher

When the description in (1a) is used referentially, the speaker focuses attention on a particular man m. Since it is clear to all that m is the intended referent, this information combines with the meaning of (1a) to produce (1c), which is the primary proposition asserted by the utterance.

1c. [the x: x is a man & x is in the corner & x is drinking champagne & x = m] x is a famous philosopher (where the content of 'm' is the man it designates)

Since (1d) is an obvious consequence of (1c), it, too, counts as asserted.

1d. m is a famous philosopher

What about the semantic content (1b) of the sentence uttered? Although it isn't a necessary consequence of the primary assertion (1c), it is an obvious consequence of (1c), plus the proposition (Pre), identifying m as the man in the corner drinking champagne.

Pre. [the x: x is a man & x is in the corner & x is drinking champagne] x = m

In many contexts, this will be presupposed by speaker-hearers. When it is, the semantic content of (1a) will also count as asserted. When it isn't, the semantic content won't be asserted, but rather will serve merely as raw material from which the real assertion is constructed. Thus, when there are two men in the corner drinking champagne, but it is obvious that the speaker is talking only about (the new guy) m, (Pre) won't be presupposed, and (1b), won't be asserted -- though (1c) and (1d) will. In this case the speaker says something true, and nothing false, even though the semantic content of her sentence is false. She is, of course, open to the charge of error in cases of misdescription – in which m is not a man in the corner drinking champagne. In these cases (1c) will be false. However, this culpable error may be mitigated, if (1d) is true. Since in many cases that assertion will be more important to the conversation than whether m is in fact

drinking champagne, or really in the corner, the fact that the speaker has, strictly speaking, said something false will matter less than her having stated an important truth.

With this in mind, let's return to *Nix vs. Hedden*, and the linguistic controversy about the terms 'fruit' and 'vegetable'. When we left off, we were considering the hypothesis that 'fruit' has only the single, botanical meaning, "fruit of the vine," and that vegetables are, by definition, edible parts of plants that aren't fruits. On this hypothesis, *the literal meaning* of the statutory language excludes tomatoes from the tax. Nevertheless, I maintained, even if this were so, it still wouldn't follow that the court had either decided the case incorrectly, or correctly substituted its own amended statute for the one passed by the legislature. It should now be clear why. The content of the statute is what the lawmakers asserted and committed themselves to, in adopting the statutory language. If, as seems plausible, they were using the terms 'fruit' and 'vegetable' referentially – on analogy with referential uses of descriptions – then they thereby asserted, among other things, that tomatoes, cucumbers, and peas are subject to the tax, even if they misdescribed them. If, in addition, this assertion was their central legislative point, the court's decision was both justified and deferential.

This is just one example of how the meaning of a sentence may differ from what is asserted by uttering it. There are many ways to create a gap between meaning and assertion. Here are some obvious ones:

- (a) Driving an out-of-towner to a destination in Los Angeles, I say, "This is going to take some time", or "It's some distance from here". In so doing, I assert that this is going to take a *substantial* amount of time, or that it is a *considerable* distance from here -- *not* that this is going to take an amount of time greater than zero, or that the distance is greater than zero, which are the semantic contents of the sentences uttered.<sup>2</sup>
- (b) A doctor examining a gunshot wound says to the patient, "You aren't going to die," thereby asserting that the patient isn't going to die *yet*, or *from this*, not that death will never come.<sup>3</sup>

- (c) Calling my friend from the airport I say “I have arrived”, thereby telling him that I have arrived *in Boston*, not that I have arrived somewhere. When he asks “Have you eaten?”, I take him to be asking whether I have eaten *dinner this evening*, not whether I have ever engaged in the activity of eating.<sup>4</sup>
- (d) When I say *I have two children*, in response to *How many children do you have?*, what I assert is that I have *exactly* two. When I say *Yes, I have two beers*, in response to *Do you have two bears in the fridge?*, what I assert is that I have *at least* two. When I say *Students are required to take 3 classes and allowed to take 5* in response to, *How many classes are students allowed to take?*, what I assert is that they are required to take *at least* 3, and allowed to take *from 3 up to* 5. When I say, *I live 400 miles from San Francisco* in response to *How far do you live from San Francisco*, what I assert is that I live *approximately* 400 miles from San Francisco. ‘Exactly’, ‘at least’, ‘up to’, ‘approximately’ – these are additions to what is asserted, generated by special features of the context of utterance, over and above semantic content.<sup>5</sup>

In all these cases, what is communicated isn’t the semantic content of the sentence uttered, but something richer, to which meaning and obvious background assumptions have both contributed.<sup>6</sup>

Another linguistic construction – involving sentences that report the use of something -- brings us to a well-known legal case about which Stephen Neale has illuminatingly written.<sup>7</sup> When one uses something, one always uses it in some way, to do some thing. But in reporting that use, how it has been used, and even what it has been used for, isn’t always explicitly mentioned. For example, what is explicitly said using the sentences in (2), can be implicitly said using (3) – provided that the context makes clear what’s being taken for granted.

- 2a. I used a hammer as a tool to pound in the nails.  
 b. I used a hammer as a weapon to fight off the burglar.  
 c. I used a hammer as a doorstop to keep to the door from closing.
3. I used a hammer.

In each case, the phrase *used a hammer* carries its ordinary meaning. In (2), this meaning is supplemented by the ordinary meaning of the phrase following it. In (3), the extra information comes from context. Nevertheless, *what the speaker says* may be the same in both cases. Of course, it's also possible to report that one used a hammer, while giving no information about how it was used -- in which case one might either utter (4), or utter (3) in a context in which it is clear that (4) is all it is being used to say.

4. I used a hammer to do something.

With this in mind we turn to *Smith v. the United States*. The question is whether petitioner's attempt to trade a gun for drugs constituted *a use of a firearm* during [what was acknowledged to be a] *drug trafficking crime*. The Supreme Court ruled that it did. Here is the statutory language:

Whoever, during and in relation to any crime of violence or drug trafficking ... uses or carries a firearm, shall, in addition to the punishment provided for such crime of violence or drug trafficking crime, be sentenced to imprisonment for five years ...

The majority opinion reads, in part: [*Italics* in all quotes are mine]

Petitioner argues that exchanging a firearm for drugs does not constitute "use" of the firearm *within the meaning of the statute*. ... In essence, petitioner argues that he cannot be said to have "used" a firearm unless he used it as a weapon, since that is how firearms most often are used. ... We confine our discussion to what the parties view as the dispositive issue in this case: whether trading a firearm for drugs can constitute "use" of the firearm *within the meaning of* [the statute].

When a *word* is not *defined* by statute, we normally construe it in accord with *its ordinary or natural meaning*. ... Surely petitioner's treatment of his [gun] can be described as "use" of the firearm *within the every day meaning of that term*. Petitioner "used" his [gun] in an attempt to obtain drugs by offering to trade it for cocaine.

The Court here tacitly assumes that *the content of the statute* – what it *says* about additional penalties – is given by *its meaning*, which is identified as *the ordinary or natural meaning of its words*.

Petitioner makes the same assumption, while claiming that the ordinary meaning of ‘use a gun’ is to use it for the purpose for which it is normally intended – *as a weapon*. As the majority puts it:

In petitioner’s view, [the text] should require proof not only that the defendant used the firearm but also that he *used it as a weapon*.... Petitioner and the dissent ... contend that the average person on the street would not think immediately of a guns-for-drugs trade as an example of “using a firearm.” Rather, that phrase normally evokes an image of the most familiar use to which a firearm is put – use as a weapon. Petitioner and the dissent therefore argue that the statute excludes uses where the weapon is not ... employed for its destructive capacity.... The dissent announces its own, restrictive *definition of “use”*. “To use an instrumentality,” the dissent argues, “*ordinarily means* to use it for its intended purpose” ...

Having framed the issue as that of determining the *ordinary meaning* of the phrase “uses a firearm,” the majority holds that those who use a firearm as a weapon don’t exhaust those to which the phrase truly applies, which also include those who merely trade it for something. Thus, the Court concludes, the ordinary meaning is controlling, and the petitioner is subject to the additional penalty.

Against this, Justice Scalia argues in dissent:

*In the search for statutory meaning, we give nontechnical words and phrases their ordinary meaning...* To use an instrumentality ordinarily means to use it for its intended purpose. *When someone asks “Do you use a cane?” he is not inquiring whether you have your grandfather’s silver-handled walking-stick on display in the hall; he wants to know whether you walk with a cane.* Similarly, to speak of “using a firearm” is to speak of using it for its distinctive purpose, i.e., as a weapon. To be sure, “one can use a firearm in a number of ways,”... including as an article of exchange...but that is not the *ordinary meaning* of ‘using’ the one or the other.

The Court asserts that the “significant flaw” in this argument is that “to say that the ordinary meaning of ‘uses a firearm’ includes using a firearm as a weapon” is quite different from saying the ordinary meaning “also excludes any other use.” *The two are*

*indeed different – but it is precisely the latter that I assert to be true. The ordinary meaning of “uses a firearm” does not include using it as an article of commerce. I think it perfectly obvious, for example, that the objective falsity requirement for a perjury conviction would not be satisfied if a witness answered “no” to a prosecutor’s inquiry whether he had ever “used a firearm,” even though he had once sold his grandfather’s Enfield rifle to a collector.*

The striking thing about this case is how a shared conflation of *the meaning of the statutory language* with *the content of the resulting statute* leads to an incoherent debate in which the dissent and majority talk past each other. The majority is right to maintain (i) that the phrase ‘uses a firearm’ occurs in the text with its ordinary, literal meaning, and (ii) that this meaning doesn’t preclude use other than as a weapon. However, the majority is wrong to think that this shows that the content of the statute – i.e. the law that the statutory language was used to enact – covers uses of firearms other than as weapons. By contrast, Scalia is wrong in claiming that the ordinary meaning of the phrase excludes uses of firearms for sale or trade. The ordinary meaning is *silent* about the manner of use. Thus, when the phrase occurs in a sentence, the resulting assertion must be completed – either by the content provided by an explicit phrase (in the manner of (2)), or by pragmatically supplied content from the context of utterance (in the manner of (3)). Since, the latter option was employed by the Congress, the job of the Court was to infer *what Congress asserted* from the incomplete semantic content provided by the statutory language.

For this, Scalia’s examples are telling. What would one be *querying* if one asked, “Do you use a cane?” Unless something special about the context indicated otherwise, one would be asking whether you use a cane *for walking*. What would one be *asserting*, on the witness stand, if one said “No,” to the question, “Have you ever used a firearm?” Not that one had never sold a firearm, but – unless something special in the context signaled something different -- that one had never fired, or used a firearm as a weapon. Similarly, unless there is something in the Congressional record not

explicitly mentioned by the majority, what Congress should be taken to have asserted is that an additional penalty will be added to crimes in which a firearm is used *as a weapon*.

The central confusion in the Supreme Court’s reasoning – between the meaning of the statutory language and the content of law enacted using that language – also marred the decision of the Circuit Court of the District of Columbia, when the case was heard there – as is evident from the following passage, quoted with approval by the Supreme Court majority.

“It may well be that Congress, when it drafted the language of [the statute] had in mind a more obvious use of guns in connection with a drug crime, but the language [of the statute] is not so limited; nor can be imagine any reason why Congress would not have wished its language to cover this situation. Whether guns are used as a medium of exchange for drugs sold illegally or as a means to protect the transaction or dealers, their introduction into the scene of drug transactions dramatically heightens the danger to our society.”

The Circuit Court seems to be saying (i) that Congress may well have understood itself to be prohibiting only a narrow range of uses of a firearm, but (ii) that *even if that was its understanding*, that isn’t what it achieved, because *the language* it used doesn’t explicitly rule out a broader class of cases, and (iii) that the Circuit Court itself thinks there is good reason for wanting the statute to apply more broadly. However, (ii) is confused. If the Congress did understand itself to be asserting that an additional penalty is attached to a narrow range of cases, then the fact that the language used to do this could’ve been used (in another context) to assert that the penalty attached to a broader range of cases is irrelevant. The content of the statute is what Congress *did* say, using that language, not what it *could*, and in the opinion of the Circuit Court *should*, have said.

### **Constraints on the Resolution of Genuinely Hard Cases**

So much for my first two lessons – the austerity of meaning and the gap between the content of a text and the meaning of its sentences. Given these, one can distinguish *genuinely hard cases* in



the law from those that are merely *semantically hard*. In the latter, legally correct outcomes aren't determined by the *meanings* of the authoritative texts, plus the relevant facts. In the former, they aren't determined by the facts plus the *full linguistic contents* of the texts, encompassing everything asserted or stipulated in adopting them. Roughly put, a genuinely hard case is one in which the totality of facts plus pre-existing law doesn't determine a correct outcome.

One subtype of this sort occurs when an outcome is determined, but it is legally incorrect. At first, this may sound incoherent. If facts plus the controlling legal texts really determine a single outcome, what could it mean to say that it is incorrect? The outcome could, of course, be morally or politically suboptimal, but that's beside the point. Courts doesn't have a general mandate to correct moral and political imperfections, and an imperfect law is still the law. No, the subtype is one in which a single outcome *is* determined by the full linguistic contents of the authoritative legal texts, but that outcome is *legally* incorrect – and not because it conflicts with other, equally authoritative texts. For if that were so, then the law as a whole wouldn't determine a single outcome -- which, by hypothesis, it does.

A toy example is provided by Lon Fuller's fable about sleeping in the train station

Let us suppose that in leafing through the statutes, we come upon the following enactment: "It shall be a misdemeanor, punishable by a fine of five dollars, to sleep in any railway station." *We have no trouble in perceiving the general nature of the target toward which the statute is aimed.* Indeed, we are likely at once to call to mind the picture of a disheveled tramp, spread out in an ungainly fashion on one of the benches of the station, keeping weary passengers on their feet and filling their ears with raucous and alcoholic snores. .... Now let us see how this example bears on the ideal of fidelity to law. Suppose I am a judge, and that two men are brought before me for violating this statute. The first is a passenger who was waiting at 3 A. M. for a delayed train. When he was arrested he was sitting upright in an orderly fashion, but was heard by the arresting officer to be gently snoring. The second is a man who had brought a blanket and pillow

to the station and had obviously settled himself down for the night. He was arrested, however, before he had a chance to go to sleep.<sup>8</sup> (my italics)

Here, the correct outcomes seem to be that the second defendant is guilty, while the first isn't – even though the innocent party did what the law explicitly prohibited, while the guilty party didn't. How can that be?

One possible answer is that the law isn't quite what it seems to be. Even though the sentence used to enact it mentions sleeping and only sleeping, one might maintain that the law itself – namely, *what was said or stipulated* in enacting it – prohibits one from *using, or attempting to use, a railway station as a place to sleep*. Since the first defendant, though dozing off, wasn't doing that, while the second defendant was, the second would properly be judged guilty, while the first wouldn't. The virtue of this account is that it gets what, intuitively, seem to be the right results, while capturing what the lawmakers would have taken pains to say, had they thought about these cases, and wished to remove doubt.

Nevertheless, the analysis is a stretch. Fuller's example is not a real-life case, but a partially specified fantasy that could be completed in various ways. Although one can imagine completions of the story in which the lawmakers really did understand themselves to be asserting the contextually enriched content that gives the desired results, one can also imagine completions in which they didn't give the matter much thought, and were content with the first formulation that sprang to mind. In these latter versions of Fuller's fable, it is hard to justify the idea that what the lawmakers asserted was sufficiently nuanced to dictate the desired results in his examples. There is, after all, a distinction between *what one actually says* in a given context, and *what one would say*, if one considered things more carefully. Are there grounds for thinking that what lawmakers actually assert *in all Fuller-type cases* must be nuanced enough to determine the correct outcome in every problematic future application? That doesn't seem very

likely. But then, we can't maintain both that the content of the law is always what was said or stipulated in enacting it, and that it determines the correct outcomes in all relevant examples.<sup>9</sup>

In my opinion, the most realistic analysis maintains the basic identification of a law with what was said in enacting it -- while adding the proviso that judges (and other officials) have the authority to make minor adjustments in its content to avoid transparently undesirable results in cases not previously contemplated. Although explicitness in writing the laws is desirable, everyone knows that it is impossible to anticipate every eventuality, and that, sooner or later, one reaches a point of diminishing returns. Recognizing this, we implicitly give judges a limited prerogative to make changes in the law they interpret -- a prerogative not found in all genres of interpretation. This is an important respect in which legal interpretation is special, and fundamentally different from some other forms of linguistic interpretation.

Adding this special interpretive prerogative to the general linguistic model for determining the content of arbitrary texts isn't difficult. Given a context in which lawmakers adopt statutory language, we apply the model to extract *what is said* in the usual way. Next, we prefix an operator -- "roughly" or "approximately" -- to that output, as the contribution to *what is said* made by the genre of lawmaking itself. The effect of this contribution is to make what is said -- which we continue to identify with the content of the law enacted -- somewhat indefinite, open-ended, and subject to authorized precisifications when required by previously unanticipated cases. Though Fuller himself seemed to think otherwise, I believe these authorized revisions must be severely restricted. His examples *don't* show that judges have the authority to rewrite laws the literal application of which would violate their own views of what *should, or shouldn't*, be legal. What is established is much more limited. *When the literal application of a law to a case violates the clear intention guiding the legislators in adopting it, judges are authorized to*

*make minimal modifications to bring it into line with what the legislators were trying to accomplish.*

The second type of genuinely hard case is one in which facts plus the full, linguistic contents of the controlling legal texts fail to yield any determinate outcome, because that content is vague. Understanding such cases requires understanding how vague language works. In illustrating this, I will draw on a theory of vague predicates according to which they are both *context-sensitive* and *partially defined*.<sup>10</sup> Although competing treatments are available, the main legal lessons I will draw can, I think, be translated into other theoretical frameworks.

Vague predicates are, as I said, partially defined. To say that P is partially defined is to say that it is governed by linguistic rules that provide *sufficient* conditions for P to apply to an item, and *sufficient* conditions for P not to apply, but no conditions that are both *individually sufficient and jointly necessary* for it to apply, or not to apply. Because the conditions are mutually exclusive, but not jointly exhaustive, there will be items not covered by the rules for which there are no possible grounds for accepting either the claim that P applies to them, or the claim that it doesn't. P is *undefined* for these items.

In addition to being partially defined, vague predicates are also context-sensitive. Given such a predicate P, one begins with a pair of sets. One of them, the *default determinate-extension*, is the set of things to which the rules of the language of the community, plus non-linguistic facts, determine that P applies. The other, the *default determinate-antiextension*, is the set of things to which the rules of the language plus the facts determine that P doesn't apply. P is undefined for o just in case o is in neither of these sets. Since the sets don't exhaust all cases, speakers have the discretion of adjusting the extension and antiextension of P so as to include initially undefined cases to suit their conversational purposes. Often they do this by explicitly predicating P of an

object *o*, or by denying such a predication. When one does this, and one's conversational partners go along, the extension (or antiextension) of *P* is expanded to include *o*, plus all objects that bear a certain relation of similarity to *o*. Consider, for example, the contextual expansion of the extension of the predicate 'blue' that occurs when I say of an object for which it is initially undefined "That's blue," and my conversational partners agree. Complications aside, the result is the adoption of a new contextual standard that counts *o*, all objects perceptually indistinguishable in color from *o*, and all objects discriminately bluer than *o* as *blue*. Note how my discretion in applying vague terms to borderline cases is constrained. If I dictate that *o* is to count as *blue*, I can't deny that objects bluer than, or indistinguishable from, *o* are too. Content-changing precisifications are possible, but only if they are principled changes.

This lesson applies to legal texts. Suppose that a legal text stipulates that things in category *C* are to be treated in a certain way, while things not in *C* are to be treated differently. Suppose further that *C* is a vague predicate, and so undefined for a range of items -- including the item *x* on which the case turns. Since *x* is neither in the default-determinate extension, nor the default-determinate anti-extension, of *C*, the rules of the linguistic community, plus relevant facts, don't determine that *C* applies to *x*, or that it doesn't. Still, application or non-application of *C* to *x* may be determined by discretionary standards adopted by authors of the text. Thus, the first step in adjudicating the case is to examine the legislative history to discover whether the lawmakers implicitly encoded such a standard into the law. If they did, then the content of the law determines the outcome of the case, and the job of interpretation is simply to discern that content from the linguistic meaning of the text, plus the context in which it was produced. If they didn't adopt such a standard, then the job of interpretation is to formulate one -- thereby creating new law by precisification. This generation of new content by filling gaps in the

original text is another respect in which legal interpretation is importantly different from some other kinds of linguistic interpretation. This special gap-filling role is justified by the fact that the court has no alternative, if it is to decide the case at all, and by the fact that legislators are well aware of this aspect of adjudication, and often deliberately employ vague language with future judicial precisification in mind. In so doing, they implicitly cede limited legislative authority to courts. What considerations *should* be employed in exercising this authority is controversial. Fullerian appeal to broad social purposes, Dworkinian balancing of the morally best outcome with deference to legal tradition and precedent, and Scalian passion for minimizing judicial intrusion into the democratic process are all philosophically defensible. Since no definitive resolution of these debates seems likely, it is hard to see how any of these positions can rightfully be denied a place at the legal table.<sup>11</sup>

Legal lessons about vagueness are illustrated by two well-known cases: *Curran v. Mount Diablo Boy Scouts*, and *Randall v. Orange County Council, Boy Scouts of America*. In *Curran*, the Boy Scouts were held to be a *business organization* and thereby subject to a 1959 California anti-discrimination law -- which was understood to require treating heterosexuals and homosexuals alike. In *Randall*, the Scouts were held not to be a *religious organization*, and therefore to violate the same anti-discrimination law by requiring belief in God. The cases turned on the vague terms ‘business establishment’ and ‘religious organization’. If the latter applies, or neither do, then both cases were wrongly decided. Only if the former applies, but the latter doesn’t were the initial outcomes legally correct.<sup>12</sup>

These outcomes can’t be justified by claiming that the rules of the language *determined* them. Although the Boy Scouts do have assets, sell a few things, and employ people, commerce is only distantly related to their main activity, and the organization certainly is *not* a clear case of

a business establishment. Either it is definitely not a business, or it is a borderline case -- and hence in the undefined range for the expression. As for being a religious organization, it has some aspects one would expect -- such as requiring belief in God and adherence to a moral code -- while lacking others -- such as prayers, worship services, and sectarian doctrines. Thus, a plausible case can be made that it is a genuine borderline case, and so in the undefined range of the predicate 'religious organization'.

Suppose, for the sake of argument, that this is true of both predicates. On this supposition, it is a matter of discretion whether or not to count them as applying to the Scouts. Did the legislators who enacted California's antidiscrimination statute exercise this discretion by implicitly adopting standards counting the Scouts as a business, but not a religious, organization? There is nothing to indicate that they did, which isn't surprising -- since, as the dissent in Randall points out, had it been understood in 1959 that passage of the act would have required the Scouts to condone homosexuality and atheism, "a fire storm of public protest would have descended on the Legislature." No, the legislators didn't themselves adopt standards dictating the outcomes later reached. Thus, the only justification for those outcomes can come from judicial resolution of the supposed vagueness.

The crucial question is whether the judges articulated *principled* criteria for narrowing the range of the terms 'business organization' and 'religious organization' so as to exclude the Scouts from the extension of the latter, while including it in the extension of the former. It is arguable that they didn't. As the dissent in Randall maintained, and the California Supreme Court later agreed, the limited range of commercial and financial activities cited provide an insufficient criterion to render the Boy Scouts a business organization, since -- if it were applied consistently -- the criterion would lead to the unwanted result that nearly all organizations of any appreciable size are

businesses – including most churches, charities, civic, and other private organizations. Thus, the initial rulings failed an elementary test. When the vagueness of terms is judicially narrowed, the result must not be a naked assertion that the term applies to a certain borderline case, or doesn't, but a principled explanation that narrows indeterminacy in a comprehensible way, and generalizes acceptably to similar cases.

The final type of genuinely hard case is one I have time only just to mention. These are cases that fall under two or more legal provisions which, when combined with the facts at hand, determine inconsistent outcomes. If neither provision is clearly subordinate to the other, no unique outcome is determined, and the job of the court is to resolve the conflict by revising existing law in some justifiable way. The problem is one of subtraction. As with other kinds of inconsistent language use – including inconsistent definitions and paradoxes of various sorts – what one wants is a purified, consistent content gotten by eliminating, or weakening, parts of the original content. Of course, not just any purification will do. The aim is to find the consistent content that is maximally faithful to the original – or, if more than one such content is maximally faithful, the intersection of all maximally faithful purifications. This, of course, is easier said than done. How one comes up with the relevant weakenings, and how, once one has them, one compares their relative fidelity to the original texts are problems for which no algorithmic solution seems possible. Still, the guiding principle is clear. In striving for resolution, one tries to preserve as much of the pre-existing law as possible. Although there is room for judicial discretion in this, it isn't essentially expansive in nature.

### **Conclusion**

This result fits my overall theme of the genuine, but limited, legislative role of legal interpretation. The meanings of legal texts *are* too austere to determine the content of the law. As a result interpretation is often needed to reach proper outcomes. However, a judge's view of the



legitimate social, moral, and political purposes of the law is *not* the crucial ingredient added by interpretation. The first and most important task of interpretation is to discern what a text says. As we have seen, its meaning is only one of the determinants of this. What speakers assert or commit themselves to is, quite standardly, a function of not only of the meanings of their words, but also of other information present in contexts of utterance. Linguistic contexts in which laws are passed are no exception. There are, to be sure, special features of such contexts that can affect how the gap between meaning and content is filled. The legislative process is governed by purposes that transcend, and sometimes conflict with, the conversational ideal of the efficient and cooperative exchange of information. Consequently, the way in which Gricean maxims, which are based on this ideal, contribute to filling the gap between the meaning and content of legal texts may, in some cases, differ from their contribution to filling similar gaps in ordinary conversations. In addition, the issue of who constitutes the audience for language use in the legislative context is up for grabs in a way it isn't in ordinary contexts. Other lawmakers are certainly part of the audience – but so, perhaps, are the law-enforcement officials who will implement the law, the lawyers who will explain it, the judges who will interpret it, and the population at which it is aimed. This complexity can, of course, affect the precise way in which the gap between the meaning of legal language and its content is filled, but it doesn't change the overall picture. The first, and most important, step in legal interpretation is not moral, social, or political, but linguistic. Presented with a legal text, one must first identify what the relevant actors said, stipulated, or otherwise committed themselves to in adopting it. For this, one needs to know how semantics and pragmatics interact to generate content.

A similar lesson applies to adjudicating genuinely hard cases, in which the full, linguistically-determined content of the law is inadequate. To resolve these cases, judges need to understand the semantics and pragmatics of vague language, as well as the ways in which consistent

contents can be extracted from inconsistent ones. After all this is taken into account, there may be a residue of remaining judicial discretion for which Fullerian or Dworkinian views of the broad social, political, and moral matters are relevant, but it isn't large.<sup>13</sup>

### Notes

---

<sup>1</sup> See Keith Donnellan, "Reference and Definite Descriptions," *Philosophical Review*, 75, 1966, 281-304 for the original discussion. For follow up, see my "Donnellan's Referential/Attributive Distinction," *Philosophical Studies*, 73, 1994, 149-168, and my "Why Incomplete Descriptions do not Defeat Russell's Theory of Descriptions," *Teorema*, XXIV, 2005, 7-30.

<sup>2</sup> Examples like these are given in Robyn Carston, "Implicature, Explicature, and Truth-Theoretic Semantics," in Ruth Kempson, ed., *Mental Representation: The Interface Between Language and Reality*, (Cambridge: Cambridge University Press), 1988.

<sup>3</sup> The example comes from Kent Bach, "Conversational Implicature," *Mind and Language*, 9, 1994, 124-162.

<sup>4</sup> Again, see Bach, "Conversational Implicature."

<sup>5</sup> For systematic discussion, see my "Drawing the Line between Meaning and Implicature – and Relating both to Assertion," forthcoming in *Nous*.

<sup>6</sup> For a more comprehensive overview of this conception of the relationship between meaning, semantic content, and assertion see my, "The Gap Between Meaning and Assertion: Why what we literally say often differs from what our words literally mean," in Hackl, M. and R. Thornton, eds. *Asserting, Meaning, and Implying* (Oxford: Oxford University Press), forthcoming.

<sup>7</sup> In "On Location," in n M. O'Rourke and C. Washington, eds., *Situating Semantics: Essays in Honor of John Perry*, (Cambridge: MIT Press), forthcoming.

<sup>8</sup> Page 664 of L. Fuller, "Positivism and Fidelity to Law: A reply to Professor Hart," *Harvard Law Review*, 71, 1957, 630-672.

---

<sup>9</sup> Similar remarks apply to the problematic case Fuller poses for H. L. A. Hart's example of a local ordinance excluding vehicles from the park. Hart's point was that because 'vehicle' is vague, its "core meaning" will provide clear cases to which it applies and clear cases to which it doesn't apply, while leaving a range of borderline cases "in the penumbra" for which there is no definite answer to the question of whether or not the term applies. As a result, Hart thought, there is a certain open-endedness and room for discretion in adjudicating the penumbral cases. Fuller's example is intended to extend this open-endedness and discretion even to clear cases. He writes, "What would Professor Hart say if some local patriots wanted to mount on a pedestal in the park a truck used in World War II, while other citizens, regarding the proposed memorial as an eyesore, support their stand by the "no vehicle" rule? Does this truck, in perfect working order, fall within the core or the penumbra?" (p. 663) Fuller's implicit contention is that the outcome of the case is unclear, and not fully determined by the local ordinance, despite the fact that the truck is not in the penumbra of 'vehicle', but rather is clear example of something to which the term definitely applies. Perhaps. However, the incompletely specified example is not entirely clear. Although there may be some ways of filling out the case for which Fuller's contention is correct, one can easily imagine other ways of filling it out for which the contention fails because the content of the ordinance is not identical with the semantic content of the sentence expressing it – e.g., "No vehicle is allowed in the park" -- but rather with a modest contextual enrichment of that content – *No use of a vehicle (as a vehicle) is allowed in the park*. In versions of the example in which this is what the rule-makers understand themselves to be prohibiting in passing the ordinance, the content of the ordinance fully determines that it does not apply to the vehicle-cum-statue, and further judicial discretion is not warranted.

<sup>10</sup> The account is spelled out in chapter 7 of my *Understanding Truth*, (New York: Oxford University Press), 1999, and my "Higher-Order Vagueness for Partially Defined Predicates," in J. C. Beall, ed., *Liars and Heaps*, (Oxford: Clarendon Press), 2003, 128-150.

<sup>11</sup> There is a third type of vagueness case that is similar to the second, but with a more limited range of indeterminacy. In these cases, it is clear that rule-makers adopted some discretionary standard, but indeterminate which standard, within a certain range, they adopted. It is thus determinate that the content of the law falls somewhere in a range of possible contents, but indeterminate where. In such cases, the job of interpretation is to discern the range of indeterminacy, and to create new, precisified, legal content, by making a principled selection from within that range. Complexity aside, there is nothing essentially new here. The nature of the interpretive task

---

in such cases is a combination of the tasks involved in the other two types of vagueness cases. However, there is something new worth noting. Once it is granted that what speakers say, or otherwise commit themselves to, may be to some extent indeterminate, it is evident that not all vagueness arises from vague language. Sometimes the move from the linguistic meaning of a legal text to that which the rule-makers were saying, or committing themselves to, in adopting it is itself an independent source of vagueness. In such cases, it is determinate that the content of the law falls somewhere within a range of possible contents, but indeterminate where.

<sup>12</sup> The cases were later overturned by the California Supreme Court.

<sup>13</sup> Thanks to my colleague Andrei Marmor for many helpful comments and discussions, as well as for teaching me most of the (all too limited) legal philosophy I know.