

*University of Southern California Law
School*

Legal Studies Working Paper Series

Year 2012

Paper 86

Farewell to Conceptual Analysis (in
Jurisprudence)

Andrei Marmor*

*mf676@cornell.edu

This working paper is hosted by The Berkeley Electronic Press (bepress) and may not be commercially reproduced without the permission of the copyright holder.

<http://law.bepress.com/usclwps-lss/art86>

Copyright ©2012 by the author.

Farewell to Conceptual Analysis (in Jurisprudence)

Andrei Marmor

Abstract

I have two main purposes in this essay: First, to show that conceptual analysis is not nearly as central to legal philosophy as typically assumed. The main methodological thrust of analytical jurisprudence, and in particular of legal positivism, is reductionism, not conceptual analysis. Consequently, the main objections to legal positivism are best seen as arguing against the possibility of reduction. Second, I aim to show that the interpretivist challenges to analytical jurisprudence bark up the wrong tree in this respect, and actually fail to engage with the methodological stance they aim to replace. Along the way I offer a partial defense of reductionism and the limited essentialism that comes with it.

Farewell to Conceptual Analysis (in Jurisprudence)

Andrei Marmor

Analytical legal philosophy is not an exercise in conceptual analysis. This is the main thesis I hope to establish in this article. Historically, conceptual analysis may have played some role in H. L. A. Hart's jurisprudence, but it was rather tangential to the main project of his theory. Hart's main objective in *The Concept of Law* was not essentially different from that of Austin, namely, to provide a reductionist theory of law. The main purpose of Hart's theory was to offer an explanation of law in terms of something more foundational in nature, that is, in terms of social facts, which, in turn, can be explained by reference to people's actual conduct, beliefs and attitudes. Hart's objection to Austin's reductionism was not to the idea of reduction but to the particular building blocks that Austin used in articulating his theory. And I think that this is generally the case: My claim is going to be that the main methodological thrust of legal positivism is reductionism, not conceptual analysis. And the main objections to legal positivism are best seen as a denial of the possibility of such a reduction.

1. Conceptual Analysis

The title of Hart's seminal book, *The Concept of Law*, gave the impression that the *concept* of law stands for something of great importance, and that it is somehow different from an ordinary understanding of what the word "law" means, yet awaiting philosophical elucidation. Of course the word "law," at least in English, has many uses, such as in "laws of nature," "the second law of thermodynamics" and the like, that have nothing, or almost nothing, to do with law in its juridical sense.¹ But if we focus our attention on the use of "law" in the juridical context, it is very doubtful that there is something that the *concept* of law stands for, distinct from what people ordinarily mean by the term.

¹ The etiology of "law" is rather complex, of course. Distinctions and very different uses we take for granted today may not have emerged for centuries.

Let us begin with the obvious question: What is a concept? And what is it that we try to do when we purport to explain or analyze the concept of X? Gilbert Ryle gave us a pretty good sense of what he meant by “concept” and what conceptual analysis consists in:

...concepts are not things, as words are, but rather the functionings of words, as keeping wicket is the functioning of the wicket-keeper. Very much as the functioning of the wicket-keeper interlocks with the functioning of the bowler, the batsman and the rest, so the functioning of a word interlocks with the functioning of the other members of the team for which that word is playing. One word may have two or more functions; but one of its functions cannot change places with another.²

Clearly enough, according to Ryle, concepts are not abstract objects; They are not things of any kind. A concept designates the myriad ways in which a word is used by competent speakers of the relevant language in a given language game – that is, ways in which it plays specific roles in making moves within an interlocking set of other concepts and arguments. Some concepts are more regimented than others, depending on the particular discipline or setting in which they are used.³ But even with regimented or scientific concepts, Ryle thought, it would be hopeless to look for a set of necessary and sufficient conditions that would provide the set of functionings a concept has in a language game. Rather we are looking at piecemeal examination of families of conceptual connections, ways in which the functioning of a word is dependent on another. Furthermore, Ryle, like other philosophers of the ordinary language analysis school at the time, believed that these conceptual connections are epistemically transparent. A careful examination of the ways in which we use language at particular settings is obvious to any competent user of the relevant language. Because we know, and necessarily so, the meaning of the words we use as competent members of the relevant linguistic community, we will always be struck with the undeniable correctness of any genuine conceptual connection whenever it is presented to us. How much and how far we can generalize from these connections, and how deep the solution they may provide for various philosophical questions, is not entirely clear. Ryle, like Wittgenstein, clearly thought that most philosophical problems are conceptual problems.

Still, the relation between the meaning of a word and the concept it designates is somewhat unclear. Words can mean different things, depending

² Gilbert Ryle, *Dilemmas*, 32.

³ *Dilemmas*, at 33-34.

on the contexts or settings in which they are used. Ryle thought that one and the same word can label different concepts. The concept designated by the word “law” for example, would be quite different when used in a juridical context from the same word used, say, in the context of a theory in physics. Quite right, but this does not answer our question: Once we identify the concept – that is, the general functioning(s) of a word in a given type of setting, say, “law” in the juridical sense – is the concept different from what the word *means* in that general setting? To use the familiar Fregeian terminology: How is *concept* different, if it is, from *sense*? I am not trying to suggest that a concept is, after all, some abstract object or entity (as Frege thought about sense⁴). The question is more practical, as it were: When we try to elucidate or analyze a concept, is there anything else to it than figuring out what the word, in its relevant settings, means in the language in question? It is difficult to see how it would be different. Take the word “chair,” for example. We use this word in English to designate a certain type of artifact, those manufactured for people to sit on, with a certain typical shape or form, etc. In short, the word “chair” in English has a fairly clear meaning or sense, designating a certain type of artifact. Now we can call this a concept, if we like; we can say that all reasonably competent users of English have a concept of chairs. But it is difficult to see how the *concept* of chairs is different from whatever this word means in its standard use in English expressions. If I asked you to tell me what “chairs” are, would you feel compelled to ask me, “Do you mean the concept of chairs, or just the meaning of the word in English?” Giving me an answer to the first would also give me an answer to the second.

Now you might think that there is a difference when less mundane words are at stake, words such as democracy, courage or knowledge. As we know, people tend to have different conceptions of what democracy is, and they may have somewhat different conceptions of what constitutes courage or what knowledge really is. And thus it has become common to distinguish the concept from its possible conceptions. We all understand the concept of democracy, it is assumed, while we may disagree about its favored conception. I have no objection to putting things this way. But let us focus on the concept. People have a concept of democracy insofar as they understand what the word means, know what it stands for and, by and large, use it correctly in their expressions. And thus, by elucidating what the word democracy stands for, you have given its meaning; at the same time, you have explicated what the concept is. Suppose, to illustrate, that someone who grew up in, say, North Korea, has

⁴ Frege’s view about the idea of sense was, curiously, a form of Platonic realism. On the significance of this stance, see Soames, *Philosophy of Language*, 12-16.

absolutely no idea what democracy is; let us assume that he has never heard the word. Now, having escaped from North Korea, he hears the word for the first time and asks you what it means. Let us assume that you give this person a correct answer. You explain everything that would enable him to use the word correctly – that is, you give him the correct meaning of the word in English. Is there anything left for him to know in order to grasp the concept of democracy? Is there any point at which you can say, “You know perfectly well what ‘democracy’ means, but you have not yet quite grasped the concept”?⁵ I fail to see what this residual piece of information might be.

One might object to this by pointing out how the mastery of numerous concept-words in a natural language requires not simply a form of knowing-that, but also some form of knowing-how; mastery of a concept is often a matter of habituation, like a skill we acquire, which takes time, practice, etc. This is absolutely true, and those of us who had to learn a second language know exactly how frustrating and difficult this process of habituation often is.⁶ But this would not show that there is a gap between mastering the meaning of a word in a natural language and acquiring its concept. You do not really know what a word means without an adequate grasp of what it stands for and how to use it in different contexts. If it takes habituation and practice to acquire this kind of knowledge, that is what it takes to grasp the meaning of the word and be able to use it correctly. Once you have acquired that competence, you have acquired the concept, and vice versa. No residue is left here, in either direction.

Perhaps in one sense this conclusion is too strong. Philosophers sometimes argue for conceptual claims that are explicitly acknowledged as *revisionist*; such concepts are not meant to reflect an agreement in judgments about the concept’s application to all its standard cases. Therefore, counterexamples to the conceptual claim, based on prevailing linguistic intuitions or common usage, would not necessarily refute a revisionist concept. I am not denying here the possibility of such claims; it is certainly possible to define a concept that is revisionist in some sense. But then the concept would need to be supported by theoretical arguments about its usefulness or theoretical purpose, not by anything traditionally called *conceptual analysis*. An analysis of concepts is an analysis of the ways in which words function in our actual language games, which must be based on observation of linguistic practices and prevailing linguistic intuitions. Conceptual analysis cannot be

⁵ Or does it work the other way around? Is there a point at which we can say, “Well, you may have the concept of democracy, but you do not quite know what the word means”? That I find even more mysterious.

⁶ See my *Social Conventions*, pp. 68-69.

revisionist. Concepts can be revisionist in various ways if they are properly defined and serve some useful theoretical purpose, but that is quite a different matter.

Thus, if we assume that reasonably competent speakers of a natural language know the meaning of the word law (in its juridical sense), then they understand the concept of law.⁷ And, crucially, there is nothing more to the concept than that which is designated by the word. But now you might be tempted to conclude that this line of thought vindicates Dworkin's famous "semantic sting" argument. Dworkin claimed that conceptual theories of law, like Hart's, are, essentially, about the meaning of the word "law." He argued that this quest for the analysis of meaning is misguided because it flies in the face of essential aspects of legal practice, most importantly, the fact that the main participants in the practice have theoretical arguments about its concept – that is, about what law really is.⁸ One standard reply to Dworkin has been that Hart was seeking to elucidate the concept of law, not the ordinary meaning of the word "law." People (myself included) have pointed out that Hart explicitly rejected the idea that his theory is about the meaning of the word "law," or that such an attempt to define the meaning of law would be theoretically fruitful.⁹ If I am right, however, that there is no significant difference between understanding the meaning of a word and understanding its concept, then this line of response is taking the wrong tack. In other words, if you think (I do not) that Hart's theory of law is basically an attempt to elucidate the concept of law, you would have to agree with Dworkin that Hart's theory is an attempt to articulate what the word "law" means; so it is a "semantic theory" after all.

I think that Dworkin is right, up to a point. He was right to assume that conceptual analysis is, essentially, a linguistic inquiry. At least it has been so conceived by the ordinary language analysis school of Wittgenstein, Ryle and Austin, to which Hart (half-heartedly, I think) belonged and which he was certainly influenced by. As we saw in Ryle's quotation, the idea of a concept stands for the "functionings of words" in their settings, and ways in which these functionings are intertwined with the functionings of related concepts with which they interact in a given setting. The idea that conceptual analysis is

⁷ Knowing what a word means, and consequently, understanding its concept, admit of degrees; in many cases it makes perfect sense to say that A has a better understanding of the meaning of X than B; or that A's understanding of the concept of X is more complete than that of B. Nothing that I say in the text should be taken to imply otherwise.

⁸ Dworkin, *Law's Empire*, chapter 1.

⁹ *The Concept of Law*, 204; see also my *Interpretation and Legal Theory*, 6-8, and Raz, *Between Authority and Interpretation*, 58-66.

transparent and reveals something that can be recognized upon reflection as obvious is only made plausible by the assumption that it is our own linguistic practices that we seek to elucidate here. Conceptual connections are transparent because they are constituted by our language, and language is public and knowable to every competent user. Conceptual analysis is, essentially, about language. I do not mean to suggest that philosophical inquiry into the conceptual connections of language use is not interesting or worth pursuing – far from it. The ways we use language typically reveal a great deal about the ways we think. Thought, however – even if it is collective and public, as revealed by language use – is never a guarantee of truth. The fact that we collectively think about something in a certain way does not mean that things are that way. Systematic errors are not unheard of, even if they are deeply entrenched in our language use.

None of this is meant to vindicate Dworkin’s “semantic sting” argument. One main weakness of that argument, as pointed out by several critics,¹⁰ consists in Dworkin’s assumption that criterial semantics, which he attributes to Hart (not unreasonably, given Wittgenstein’s influence), is incapable of explaining the kind of “theoretical disagreements” in law that he focuses on. Dworkin’s main assumption, that knowledge of the meaning of a word is necessarily tantamount to agreement in judgments about the word’s application to all (non-borderline) cases, is questionable indeed. But this is not the line of response I want to explore here. In fact, as I indicated above, I share Dworkin’s view that conceptual analysis is not distinguishable from an analysis of the meaning of words in a given setting. The question is whether Hart’s theory of law is a form of conceptual analysis at all and, more generally, whether analytical jurisprudence is best seen as a form of conceptual analysis. I want to give a negative answer to both of these questions.

2. The Possibility of Reduction

Many courses in analytical jurisprudence begin with Austin’s command theory of law, and for good pedagogical reasons. Austin’s theory seems very intuitive, almost compelling at first sight, and thus, by revealing its weakness (usually on the basis of Hart’s critique), we can motivate a discussion of the hard questions about the nature of law, opening the way to more sophisticated philosophical accounts presented to us by philosophers of the 20th century. But there is a pervasive danger in this pedagogical approach: By systematically demolishing

¹⁰ See, for example, Raz, *Between Authority and Interpretation* at 62-76.

Austin's construction, we may be throwing out the baby with the bathwater. And this baby is the methodological aspiration that motivated Austin: to explain the law in terms of something else, more foundational in nature. In short, we recognize that Austin's theory of law purported to provide a reduction of law to facts of a social kind, explicable in sociological terms. And this reduction strikes us as crude and ultimately unsuccessful. Many think that Hart must have given up on this reductionist project, since he demonstrated many of its failures. But this assumption is seriously misguided. Hart's own theory of law is as reductionist as Austin's. Hart clearly shared Austin's view that a theoretical explanation of the nature of law should explain what the law is in terms of social facts, facts that can be explained by more foundational truths about how people behave, the kind of beliefs they have about their conduct, and the kind of attitudes and dispositions that tend to accompany those shared beliefs. In other words, the hallmark of Hart's theory is the idea that social rules are at the foundations of law, and that social rules, in turn, can be explained reductively in terms of people's actual conduct, beliefs and attitudes.

All of this requires careful demonstration, of course, but since I have done this elsewhere,¹¹ I will not try to reiterate the whole argument here – with one exception that will prove useful for the argument in the sequel. I venture to guess that Hart's reductionist project has been blurred by his own cryptic remarks on the importance of the *internal point of view*. The standard lesson that people draw from his remarks is the exact opposite of what Hart actually says in those pages.¹² According to the common misunderstanding, Hart's suggestion that an adequate account of the normativity of law must explicate the ways in which participants in the practice regard legal requirements as reasons for their action, is taken to indicate that Hart shared Kelsen's anti-reductionist critique of Austin. Hart must have meant, people assume, that there is something misguided about an explanation of law's normativity in purely sociological terms. But in fact, it is the other way around. Hart's discussion of the internal point of view is actually a critique of Kelsen, not so much of Austin. There was not much of a need, after Kelsen, to remind us that an account of law as a normative system must include the recognition that the law is taken to constitute reasons for action for those who regard it as binding.

¹¹ The first three chapters of my *Philosophy of Law* are devoted to articulating the reductionist project that was in dispute between Austin, Kelsen and Hart. Nevertheless, a reviewer (see Patterson's review in NDPR 07/17/11) complained that my book presupposes the soundness of conceptual analysis and invited me to say more in defense of this method; as I try to explain in this paper, this is an invitation I have to decline.

¹² *The Concept of Law*, p 87.

To regard a norm as binding is to see it as providing reasons for action. The point Hart makes about the internal point of view is that it can be explained reductively, by observing people's actual practices, beliefs and attitudes. There is no need, he claimed, to postulate a presupposition of a basic norm, as Kelsen argued; what we need is only to recognize that, where there is a functioning legal system in place, most key participants *regard* legal norms as binding – which is to say, as something that gives them reasons for action. And, crucially, we can do all this by reporting on the kind of beliefs and attitudes that the insiders share. We do not have to endorse their point of view, or talk about it as if we do, only to recognize that a social rule exists iff the relevant population shares certain beliefs and attitudes:

For the observer may, without accepting the rules himself, assert that the group accepts the rules, and thus may from the outside refer to the way in which they are concerned with them from the internal point of view.¹³

And perhaps this is even clearer when Hart explains the difference between the rules of recognition and Kelsen's basic norm:

First, a person who seriously asserts the validity of some given rule of law...himself makes use of a rule of recognition which he accepts as appropriate for identifying the law. Secondly...If the truth of his presupposition were doubted, it could be established by reference to actual practice: to the way in which courts identify what is to count as law, and to the general acceptance or acquiescence in these identifications.¹⁴

In short, Hart's disagreement with Kelsen is precisely about the possibility of reduction. They both shared the view that legality can only be explained on the basis of some normative framework that grants certain actions and events in the world the legal significance that they have. While Kelsen thought that this normative framework must, ultimately, be postulated or presupposed, Hart strove to show that there is no need for such a presupposition; the norms that confer legal significance on acts and events are social norms that can be "established by reference to actual practice," as he put it.¹⁵ And by practice Hart means the actual patterns of conduct people exhibit, the kind of beliefs they share about it and the attitudes that accompany those beliefs.

¹³ *The Concept of Law*, at 87

¹⁴ *Ibid.*, at 105.

¹⁵ In *Philosophy of Law*, chapter 1, I argue that Kelsen's anti-reductionist argument fails on its own terms; Kelsen's argument about the basic norm cannot withstand the pressure for a reductionist account, one that Hart provided a few years later.

Let us get back on track now. If, as I claim, the main methodological thrust of Hart's theory of law, like that of Austin's, consists in a reductionist project, where does that leave conceptual analysis? Is there a sense in which reduction of some type of discourse or class of statements to another type or class is a form of conceptual analysis? That depends on the kind of reduction in play. In some contexts, philosophers strive to offer a *semantic* reduction of one type of discourse to another. A semantic reduction would have to satisfy the condition that the entire vocabulary of a given type of discourse or theory, say, D1, can be fully expressed in terms of the axioms and vocabulary of a different type or class of statements, say D2. If this condition can be met, we have a semantic reduction of D1 to D2. I am not sure that there is a very good example of a successful semantic reduction in any domain,¹⁶ but I am willing to assume that, if such a reduction is possible, it might constitute a form of analysis that is, essentially, about language. We could then say that the concepts ordinarily used in one type of discourse actually *mean* something else, fully explicable in the terms of some other type of discourse or semantics.¹⁷

The type of reduction discussed in legal philosophy, however, is not of this semantic kind.¹⁸ Rather, it is a metaphysical or constitutive form of reduction. The idea of a metaphysical reduction is to show that a distinct type of phenomenon is actually constituted by, and fully reducible to, some other, more foundational type of phenomenon. In our case, the idea is to show that law is constituted by social practices that can be fully explained by the way people actually behave, the kind of beliefs they share about their behavior, and the attitudes and dispositions that they exhibit in the relevant contexts. In short, we try to reduce one type of facts to some other, more foundational type of facts.

Constitutive or metaphysical reduction is not about concepts. It is about the nature of things – that is, about the actual properties of objects or phenomena and their constitutive elements. Consider, for example, the debates about the reduction of biology to chemistry and physics. The question here is whether we can fully explain the phenomena manifest in the biological sphere

¹⁶ Some versions of expressivism in metaethics come very close to attempting a semantic reduction of ethical discourse to expressions of desires, preferences and the like. How successful this project is, needless to say, remains highly controversial.

¹⁷ I am not suggesting that this would be a form of conceptual analysis in the traditional sense, but somewhere in the neighborhood.

¹⁸ It is possible to construe Austin's legal theory as attempting a semantic reduction as well. If this is the case, then Hart certainly did not share the objective; nothing in Hart's work suggests that he strove to provide a semantic reduction or that he thought such a project to be feasible.

in terms of chemical reactions and physical properties. This is not a philosophical debate about the meaning of concepts deployed in the biological sciences. The success of a reductionist explanation is not determined by how successful the theory is in explaining our use of language. The criteria of success are determined by the question of how complete an explanation of biological phenomena in terms of chemical and physical theories can be. The explanation we seek is of the observable phenomena, not of the use of concepts in a given setting.

One might object to the example: Arguably, most concepts employed in biology are of the *natural kind* – that is, they are concepts purporting to refer to some set of objects or phenomena, whatever their real nature turns out to be. So there is a sense in which the use of such concepts already assumes, as it were, that our grasp of their reference might be partial and incomplete. I certainly do not assume that law is a natural kind of concept. But it still remains the case that an explanation of law is an explanation of observable phenomena, even if the phenomena are of a social kind, constituted by the collective behavior, shared beliefs and social interactions of human beings. Hart's reductive theory of law is not an attempt to explain how people use the concept of law; it is an attempt to explain what law is in terms of observable social behavior of human beings. More specifically, it is an attempt to explain what constitutes legal practices and institutions, what makes it the case that people regard some such practices as legal while others they do not. The fact that the reference of "law" is a social construction does not entail that a theory of this reference is a theory about concepts.¹⁹

None of this means, of course, that the social nature of the object of a reductive theory does not impose some constraints on what would count as a successful explanation. In particular, we must examine the question of how people's self-understandings of the social practices they engage in are part of what needs to be explained, and what kind of constraints this imposes on the nature of the relevant explanations. I will get to this indirectly, however, first considering the challenge of interpretivism.

¹⁹ By way of illustration, consider a Marxist explanation of religion. Marxism clearly assumes that religion is a social construct, if anything is. But surely a Marxist explanation of religion is not a conceptual analysis; it purports to explain the relevant social constructions reductively in terms of the material forces and means by which class struggles generate false consciousness, etc.

3. The Allure of Interpretivism

Presumably, nobody denies Raz's suggestion that "In large measure what we study when we study the nature of law is the nature of our own self-understandings. ... it is part of the self-consciousness of our society to see certain institutions as legal. And that consciousness is part of what we study when we inquire into the nature of law."²⁰ Yet some philosophers argue that it is precisely this truism that should lead us to realize that an inquiry into the nature of law is essentially interpretative in nature. And, they argue, interpretation of this collective self-consciousness is partly, but essentially, a matter of making moral sense of the concepts and ideas we associate with the practice. Interpretation, in other words, is partly a matter of evaluative judgments, a matter of figuring out the best moral justifications we have for the practices in question.

The plausibility of interpretivism in legal philosophy comes from two main sources. One is the assumption, sometimes explicitly stated, that the alternative to it, or rather, its object of critique, is conceptual analysis.²¹ The second consideration lending support to interpretivism comes from analogies to other areas in which it seems rather compelling, such as self-understandings in the realm of the arts, or various social practices, such as making promises or conventions of civility. In these areas it seems clearer, perhaps, that a philosophical explanation of what the practices are is very closely tied to the kind of values we find in them, and the ways in which we articulate what makes them valuable and worth paying attention to. There is, however, a subtle difference between two lines of thought here. According to one line of thought, the need for an interpretative theory comes from the limits of conceptual analysis – that is, from the dead end we reach when we focus our attention on the nonnormative facts allegedly constituting legal practice. According to a different line of thought, the debate between traditional legal philosophy and interpretivism is about the order of explanation. Instead of beginning with our shared conceptual understandings of what law is, based on the kind of claims law allegedly makes on us, and then perhaps asking a moral question about the legitimacy of those claims, interpretivism asks us to reverse the order; it poses the moral question first, asking what makes certain kinds of institutional

²⁰ *Between Authority and Interpretation*, at 31. Needless to say, philosophers may disagree about the question of how robust "our" self-understandings are, and how much of a pre-theoretical consensus or agreement in judgments we can assume here. I will dwell on this in greater detail below.

²¹ See Dworkin, *Law's Empire*, Stavropoulos "Obligations, Interpretivism, and the Legal Point of View", and Greenberg "The Standard Picture and its Discontents".

demands legitimate, and then it purports to extrapolate some constraints on what those institutional or political institutions have to be, and what kind of claims they can make, and how, in order to satisfy the favored conception of moral legitimacy.²²

One of the main problems with interpretivism is that it takes conceptual analysis to be its main rival, its main focus of criticism. As I argue, however, conceptual analysis is really not so central to the mainstream philosophical enterprise of analytical legal philosophy, or to the “orthodox” view (or “standard picture”), as interpreters typically label the main traditions of 20th century legal philosophy, particularly the legal positivist tradition. The orthodox view is mostly about the possibility of reduction.²³ And the type of reduction that is sought by the orthodox view is not about concepts; it is about observable phenomena, including, of course, social phenomena exhibited by the collective actions and shared beliefs of a population. So the question we need to examine here is whether interpretivism is the kind of view that can engage with, and form a genuine critique of, the reductionism that forms the main object of traditional legal positivism. My answer is going to be that interpretivism engages with this methodological debate if, and only to the extent that, it denies the possibility of reduction.

Let me start with the question about the order of explanation.²⁴ There is, of course, nothing to prevent anyone from starting with any kind of philosophical question they like. If you want to start with a moral question about the legitimacy of state coercion or the legitimacy of political authority or whatever, that is fine. It has been done for centuries and is mostly referred to as political philosophy. The question is, of course, whether anything about the nature of *legal* philosophy forces this order of explanation on us. Interpretivism is committed to the thesis that, indeed, something forces on us the primacy of the moral question. The order of explanation is forced on us, the argument

²² I am not suggesting that different philosophers adhere to one or the other of these two versions of interpretivism; on the contrary, Dworkin, Stavropoulos and Greenberg (if I understand him correctly) advance both of these versions concomitantly, shifting from one to the other as if they are one and the same, or at least very closely linked. At least Stavropoulos makes it quite clear that he does not assume that the demands in question are necessarily institutional or somehow pre-theoretically “legal” in any sense.

²³ Scott Shapiro’s more recent “planning theory of law” is also quite clearly a reductionist project, aiming to reduce law to activities of planning, and certain types of collective intentions (based on Bratman’s theory that is, itself, reductionist in nature). See S. Shapiro, *Legality*.

²⁴ This argument is most clearly stated in Stavropoulos “Obligations, Interpretivism, and the Legal Point of View.”

seems to be, because the opposite order takes too much for granted.²⁵ If you start with an observation about the nonnormative facts allegedly constituting the practice you strive to explain, you necessarily privilege certain prevailing concepts and self-understandings as defining the very subject matter of what you strive to explain. In other words, even if the method of legal philosophy is not necessarily conceptual analysis, you must start with some concepts that define the subject matter, concepts or conceptions that determine what is in need of explanation.²⁶ But, interpretivism tells us, nothing about the prevailing concepts of a practice or an institution really justifies this privilege of determining what is the true nature of the subject matter that calls for theorizing. Concepts, prevailing and entrenched as they may be, can get things totally wrong, even if they are not contested at the time. Just as astronomy would be wrongheaded to take for granted that stars are holes in the sky, as was once widely believed, and take it from there, so legal philosophy would be wrongheaded to assume that law is, say, what political authorities decree, and just take it from there.

The analogy is misleading, of course; we can start with a theory that stars are holes in the sky and see how far it takes us. As it turns out, it did not take us far enough, so we had to revise the assumption about what stars are in light of the relevant evidence. Philosophy, however, does not have the luxury of the scientific methods allowing hypotheses to be tested empirically.²⁷ So the danger is that if we take concepts for granted, we might be stuck with trying to explain how stars are holes in the sky.

If we do not want to make this argument into an indictment of philosophy as such – which would be self-defeating, of course – we must limit the concern to the place where it belongs. In other words, I doubt that interpretivism purports to challenge the very possibility of any kind of nonnormative philosophy. But then, if nonnormative philosophy is possible, what makes philosophy of law different? The answer might be that law is different from other philosophical subjects because it is, in itself, a normative practice – that is, the kind of practice that makes moral demands on us. But at

²⁵ *Ibid.*

²⁶ Raz comes very close to admitting as much when he claims that what we strive to explain in legal philosophy is what the concept of law, as we now have it in contemporary western societies, stands for, even if it is also true that, in different societies and at different times, people entertained rather different concepts of law (or none). See *Between Authority and Interpretation*, 40.

²⁷ Brian Leiter, and others who share an empiricist view of philosophy, would disagree. This is not the place, however, to discuss the controversial idea of “naturalizing” jurisprudence and other forms of philosophical empiricism.

this point we are already beginning to lose sight of the argument. The indictment of orthodoxy is based on the claim that it is bound to privilege concepts and self-understandings of law in a way that predetermines the issues; we cannot assume, this argument goes, that our concepts, and the ways in which legal practice is understood by those whose practice it, correctly identify certain central features of law as the appropriate objects of philosophical inquiry. At best, such self-understandings form a possible hypothesis that competes with others we can come up with – but a hypotheses of what? If we do not make any assumptions about the object of our inquiry, hardly any philosophical explanation can get off the ground. The answer we are given is that the alternative to privileging prevailing concepts is to take certain paradigmatic examples as tentatively defining the subject matter and potentially revising those paradigms as we go along.²⁸

The main problem here is that the philosophical method that interpretivism recommends is vulnerable to the same problem that it accuses orthodoxy of. If your starting point is a moral question about the law, and then you assume that the important aspects of law are those that figure in, or follow from, your moral answer, you may have also assumed too much. Consider, for example, the moral question that Dworkin (and, it seems, Stavropoulos) takes to be central – namely, the question about the justification of the use of collective force or state coercion. Quite clearly, the question itself presupposes that coercion is an essential aspect of law; if it is not, why focus on coercion? It may turn out that many of law's functions in society have very little to do with the use of force or coercion. Are those noncoercive functions not essential or not characteristic of law? Interpretative, partly moral, questions about the law are just as much *about the law* as any nonnormative questions. They must start somewhere, and not just anywhere; they must start with some understanding about what law is and which aspects of it are more central to it than others. (I will return to this point later.)

Now, true enough, Dworkin sometimes gives the impression that he actually cares little about the question of whether his interpretative-normative theory captures some pre-theoretical essential aspects of law or not. But then we are back to the problem that not just about any normative question in political philosophy has something interesting to tell us about the nature of law. In other words, unless it is assumed that coercive force (or whatever else one deems to be the focus of the relevant moral question) is an essential or central aspect *of law*, the theory of legitimacy offered would not necessarily have any

²⁸ See Dworkin, *Law's Empire*, 91-92 & Stavropoulos, "Obligations, Interpretivism, and the Legal Point of View" at ...

bearing on the nature of law. Political philosophy is not legal philosophy by default.

At this point, the second line of thought that leads to interpretivism may be brought in. What is special about the law, interpretivists seem to claim, is its very normative nature, the fact that it makes certain normative, including moral, demands on us, and that we regard at least some of these demands as morally warranted or binding. By itself, however, this observation, sound as it is, does not lead to the conclusion that any philosophical theory about a normative practice must begin with a moral question, or that it must include moral arguments as part of the theory. The argument has to be more subtle than this. Perhaps something like the following analogy might do the work: Consider, for example, the realm of art. We regard certain human creations as works of art, and countless others as having no artistic merit whatsoever. So we can ask: What makes anything a work of art? Quite plausibly, we cannot give an answer to this question without having some views about where value lies in art, or what makes art have the kind of values we associate with it. Interpretivism in the realm of art, or philosophy of art, if you like, seems rather compelling. Why would it not be equally compelling in philosophy of law?

There are several reasons to doubt the analogy. First, interpretation is central to art because works of art are created as objects of interpretation; they are created (partly) as an invitation to appreciate their aesthetic and artistic features, to appreciate the kinds of achievement they manifest, and the like. Law, however, is not created to become an object of interpretation – law is created to guide human conduct; it has practical purposes. Making a law is not an invitation to the public to offer interpretations of it or to appreciate its legal qualities. Second, there is a clear sense in which art is an essentially contested concept. That is, the word “art” stands for a certain form of human achievement, an aspiration for excellence that is essentially contested, inviting different views about what kind of achievement it is and what its standards of excellence are. Each work of art contributes to this cultural debate, as it were, making an implicit statement about its conception of art, what the creator values in it, etc. And that is why the question of whether a given artifact is a work of art or not, and the question of what we value about art, are very closely linked.

Legality, however, is not an essentially contested concept. We do not regard legality as a form of human achievement, inviting, as it were, different conceptions of what makes it an achievement, what it is an achievement of and

the standards of excellence we associate with it.²⁹ The making of *good* law – morally, politically, economically or otherwise – is of course a form of achievement. We may have different conceptions of what would be a good law in this or that domain, but the relevant issue that is essentially contested here is not the legality of the lawmaking but the moral or other evaluative dimension of it – that is, the contested element here is the “good,” not the “law.”³⁰

None of this shows interpretivism to be misguided as a philosophical method. It only shows that we have not yet found a way in which interpretivism actually engages with, not to say refutes, the reductionist project that forms part of the “orthodox” view. It seems, so far, just an alternative philosophical interest in law, or something in that neighborhood, guided by moral concerns about the potential impact of certain types of demands on our reasons for action. In order to engage with the philosophical method of the orthodox view, interpretivism has to argue that reduction of law and legality to some social facts is not attainable. In other words, interpretivism is a challenge to legal positivism iff it amounts to the denial of the possibility of reduction. I think that it may amount to that, but I will explain this a bit later. First, I want to consider the challenge posed by interpretivism to the kind of essentialism that has informed contemporary legal positivism.

4. Essentialism in Perspective

You cannot offer a reduction of a certain type or class of phenomena to some other type or class of facts unless you have a fairly robust sense of what constitutes the main features of the reduced class. Furthermore, the success of a reductionist theory depends on how complete the explanation it offers is, and we cannot have any sense of that without knowing what needs to be explained. If we want to offer a reductionist theory of law, we must have a pretty good idea of what is essential to law and what is only incidental and contingent. Otherwise, we would lack any criteria of success for the reduction offered. And this is where, interpretivists claim, we get into trouble. Assumptions about

²⁹ The *rule of law* can be seen as a form of achievement that clearly admits of degrees of success or failure. It may well be an achievement of a given society and its government that it adheres to the rule of law, and societies or regimes achieve these goals to a greater or lesser degree. But this is an entirely different matter. I have elaborated on this distinction in my *Philosophy of Law*, pp 132-134.

³⁰ This is not meant to be an argument against a teleological conception of law, such as John Finnis’, maintaining that law is essentially concerned with the common good.

essential aspects of law are profoundly suspect, they claim, because they privilege certain conceptual assumptions or self-understandings of law that might be as erroneous or misguided as the idea that stars are holes in the sky. In other words, the idea here is that any assumption about the essential aspects of law is bound to be contestable, morally or otherwise, and hence cannot form a benchmark against which we can test the plausibility of a reductionist theory.

The debate typically focuses on one prominent example, offered by Joseph Raz, claiming that it is an essential aspect of law that it purports, and necessarily so, to claim legitimate authority over its subjects.³¹ Law's authoritative nature, Raz famously argued, is an essential aspect of law, and thus any theory about the nature of law must give an explanation of how law can be authoritative and how it can make the kind of claims that potentially meet the conditions of constituting legitimate authoritative directives. Interpretivism challenges this assumption not by denying that law is authoritative, but by challenging the essentialism that Raz's position assumes here. Simply put, the question is: How can we tell that anything is essential or necessary to law? If it is a conceptual claim – as most commentators understand Raz to be assuming – we are back to the difficulties of conceptual analysis and the unwarranted assumption that prevailing concepts are somehow privileged in defining the subject matter of the inquiry. And if it is not a conceptual claim, what kind of claim is it?

This is a fair question. Before I try to answer it, however, let me reiterate a point I made earlier – that I do not see how interpretivism avoids the same problem. Suppose you do not make any assumptions about essential aspects of law and instead ask yourself what justifies the use of collective force, or any other similar question that puts a moral concern about law at the forefront. We would still need to know that your moral concern is one about the law, and, more importantly, that the conclusions you reach are such that they tell us something of interest about the nature of law. I just do not see how any of this can be done without making certain assumptions about aspects of law that we deem more essential to it, or more characteristic of it, than others. What makes Dworkin's question about the legitimacy of the use of collective force more relevant to law, than, say, a question about the moral legitimacy of progressive taxes? The answer has to be this: We can have law and a functioning legal system in place without a progressive tax system, while, presumably, we cannot have law without some coercive element at work. Which is to say that we have

³¹ This is not the only essential feature of law that Raz argued for; law also essentially claims normative superiority to any other normative system, and it claims the authority to regulate any aspect of human conduct. See his *Practical Reason and Norms*, 149-154.

already assumed that coercion is somehow more essential to law, more central to what law is, than a progressive system of taxation. And I hope it is clear that the example generalizes to all cases.

Dworkin might respond that the choice of the moral question is itself morally determined; we need to focus on law's coercive aspect because it is the morally pressing issue here (and, of course, it may not be the only moral question in the neighborhood; others might be just as relevant). That may be true; it is possible that there are very good moral reasons to be concerned with law's coercive aspect. But again, none of this would guarantee that the answer to the moral question you come up with tells us something about the nature of law. Unless you assume that coercion is central or essential to *what the law is*, a moral theory about the legitimacy of coercion would be just that – a moral theory about the legitimacy of coercion.

But I do not want to dodge the question: On what grounds can we claim that F is an essential feature of X? To begin with, it would be a mistake to equate essential properties of X with *a priori necessary* conditions. That “F is an essential property of X” is not necessarily an *a priori* truth. And it is not always the case that if F is an essential property of X, then F is a necessary condition for something to be an X. There are many essential properties of objects or categories that we learn from experience. We learn from experience, for example, that the ability to use language is an essential aspect of homo sapiens. Linguistic competence is partly, but quite essentially, what defines humans as a distinct biological species. This is hardly an *a priori* truth, though. In fact, it rests on an empirical generalization that could turn out to be inaccurate.³² Second, to continue with the example, when we say that linguistic competence is an essential feature of humans, we are not committed to the idea that a person with serious brain damage that causes her to lose the ability to use language renders that person nonhuman. Essential features are not necessary conditions. To take another example, it makes sense to suggest that “manufactured for humans to sit on” is an essential feature of chairs; it is what chairs are. Yet a chair designed for an art exhibit, not made for people to sit on, might still be a chair. Essential features of an X might be necessary conditions for central cases to be an X, but what counts as central cases admits of some vagueness, with borderline cases along the way to marginal and deviant cases.

³² We may come to learn, for example, that dolphins' linguistic competence is much closer to humans' than we think – perhaps even superior to ours! And then, of course, we may need to revise our views about how distinct a characteristic linguistic competence is, and how it defines our categorization of species. But it would not undermine the claim that linguistic competence is essential to humans; it may also be essential to dolphins.

Be this as it may, what constitutes the idea of an essential property of something is bound to vary between different domains or realms of inquiry. Essential features in number theory are probably different from essential features in biology. Let us focus on essential properties of social practices. (Now you know what is coming.) We can say that it is an essential aspect of chess that it is a competitive game. Perhaps we need to qualify this statement: It is an essential aspect of chess, *as we know it*, that it is a competitive game. For all we know, there might be Martians who “play” chess not as a competitive game but as a religious ritual; perhaps the winner becomes the high priest.³³ So it is not really a game on Mars. Whether Mars-chess is, in some sense, chess or not, is a difficult question to answer, but it is not what is needed here. What is needed is to realize that the characterization of chess as a competitive game is not a statement about concepts or, generally, about language. It is a claim about the nature of a certain type of human activity. And whatever we know about chess is not something we learn by mastering a language, but by observation, learning and participation in this and similar activities. If you do not know, for example, that chess is a *competitive* game, that players must aim to win the game, and that what counts as winning is determined by the rules, you will not be able to play chess, nor will you understand what is going on when observing others playing it. True enough, you will also not know what the word “chess” means, thus we can say that you do not quite have the concept of chess. But learning that chess is a competitive game is not something that comes about by learning to speak English; it comes about by learning about chess, what it is and how to play it (or, at least, how to understand what others are doing and aiming for when playing it).

Admittedly, I have only responded to half of the challenge; I suggested that essential properties of social practices are properties or features of observable and learnable activities, not of language or concepts. But I have not said anything about what makes a feature essential and how we can come to know it. I think we can safely say this: A feature of a social practice is essential to it if without it the practice would either not have existed at all or would have been radically different from what it is. At least, this is how we normally go about trying to determine whether a suggested feature is essential or not. If you suggest that some feature F is essential to law, we should ask whether a practice that is similar in all other respects, but lacks F, is still close enough to what we recognize as law in our society. This would easily show, for example, that progressive taxation is not essential to law; we would clearly recognize as a legal system one that has a flat tax rate, and even one that imposes no taxes at all.

³³ The example is from Schwyzer, “Rules and Practices.”

Some cases are harder to determine, of course. The question of whether coercion is essential to law is not so easy to answer. Thinking about a practice or institution that would be similar to law in all other respects, but for the fact that there is no coercion involved whatsoever, might stretch our imagination beyond our ability to come up with a determinate answer. It may be so, although I am not suggesting that this is the case. If it is, we may not have a determinate answer to the question of whether coercion is an essential aspect of law. Since essential features are neither *a priori* nor necessary, the possibility of indeterminacy about some cases is not philosophically problematic. I see no reason to deny that what is essential to a social practice is vague, allowing for some borderline cases.

Now let us assume that Raz's thesis that law is essentially an authoritative institution is on par with the suggestion that chess is a competitive game. First, we can conclude that this is not a conceptual claim, at least not in the sense that it is a proposition about concepts. It is a proposition about an essential feature of a certain type of human activity, a social practice. Second, we can have a sense of how to go about determining whether Raz's suggestion is true or not: Imagine a practice or institution that is somehow similar to law in all other respects but for the fact that it makes no claim to be authoritative. How close would it be to what we regard as legal practice in our world? I hope the answer is evident, but that is not what I need to substantiate here. Finally, it does not follow from the essential authoritative nature of law that each and every law must make a claim to be authoritative; some deviant cases are possible.³⁴

Authority was just an example; other essential features of law have been identified by Hart (that modern legal systems combine primary and secondary rules), by Kelsen (that laws must come in systems, with some hierarchical structure) and, if I may, by myself (that law has conventional foundations). And these features, and others, are not mutually exclusive, of course. Needless to say, the list of features that philosophers identify as essential to law is not randomly selected. The features serve some theoretical purposes, highlighting an aspect of law that calls for theoretical explanation and has some significance for the kind of theory offered. But there is nothing question-begging about that. Theoretical questions and interests always come up against the background of previous theories and the difficulties they give rise to; our interest in particular

³⁴ One example I have in mind are cases where legislatures enact laws that have no prescriptive action-guiding content. In the U.S., for example, states often legislate that X is the "state bird" or something like that, which is a symbolic act without any authoritative directive included.

questions and the ways in which we phrase them is path-dependent, shaped by traditions constituting the relevant discipline. This is quite universal and applies to moral theories as well.

5. The Limits of Reductionism

I mentioned earlier that interpretivism can critically engage with what it calls the orthodox view iff it denies the possibility of reduction. In his earlier criticism of Hart's jurisprudence, this is precisely the path that Dworkin followed. Though not phrased explicitly as an anti-reductionist critique, Dworkin's argument about legal principles aimed to show that Hart's rules of recognition cannot provide a complete answer to what constitutes legal validity, because legal validity is partly deduced by moral reasoning. This is quite clearly an argument that purports to show that legality cannot be reduced to social facts.³⁵ Unfortunately, the interpretivist turn in Dworkin's theory obscured this anti-reductionist argument, and led many to suspect, including Hart himself, that interpretivism and legal positivism talk past each other. Legal positivism, as Hart clearly saw it, aims to provide the essential features of law quite universally – that is, wherever law exists – whereas the impression that Dworkin left about interpretivism was that it articulates a theory of adjudication for Anglo-American law.³⁶ I think that Hart was not quite right about this. There is a clear sense in which Dworkin's interpretative theory challenges the reductionism inherent in Hart's legal positivism. It aims to show that legality cannot be reduced to facts about people's conduct, beliefs and attitudes because legality is partly, but necessarily, constituted by truths about moral issues. If truths about morality necessarily figure in the determination of what counts as a legal requirement, then it is clearly not the case that legality admits of reduction to social facts. And, of course, Dworkin does not need to show that this is the case in every legal system we are familiar with; in order to refute the possibility of reduction it is sufficient to show that reduction is not possible at least in some cases. Since reductionism is a universal claim, it can be refuted by counterexamples.

This is not the place to evaluate Dworkin's anti-reductionist argument; I have done that elsewhere.³⁷ I want to focus, instead, on a question we have left open so far, about the kind of constraints imposed on any theory about the

³⁵ See Dworkin, "The Model of Rules I"; I explain this in greater detail in my *Philosophy of Law*, chapter 4.

³⁶ See Hart, postscript to *The Concept of Law* (2nd ed.)

³⁷ *Philosophy of Law*, chapter 4.

nature of law stemming from the fact that part of what we aim to explain is people's self-understandings of what law is. Dworkin maintains that a theory about the nature of law must provide the kind of explanation that participants in the practice can recognize, at least upon reflection, as something that rationalizes for them the ways in which they understand the practice and its normative requirements. And he clearly thinks that this constraint is general, one that applies to philosophical explanations, viz., *constructive interpretations*, of social practices in general. Call this the *rationalization constraint*. Now, of course, if you hold the view that any philosophical explanation of a social practice is necessarily a form of constructive interpretation, the rationalization constraint makes a lot of sense. Constructive interpretation, as Dworkin understands this idea, is the kind of explanation that is aimed at those whose practice you interpret; it aims to make sense for the participants what they should regard as valuable about the practice and how they should rationalize their own understanding of it.

A reductionist account of a social practice, however, is not committed to the rationalization constraint. In some clear cases, reductionism in social explanations clearly violates this constraint. A Marxist explanation of religion, for example, cannot offer the kind of explanation that religious people could come to recognize as their own, so to speak, at least not as long as they still deem themselves religious. In other words, a reductionist explanation of a social practice is not expected to provide a kind of "aha!" moment for the subjects, providing them with an insight that would rationalize the practice for them.

This does not mean, however, that the social nature of the object of explanation of a reductive theory does not impose some constraints on the kind of explanations we can offer. It is difficult, if not impossible, to explain anything in the social realm without invoking people's beliefs (and attitudes shaped by those beliefs, etc.). Yet most (not to say all) of our beliefs are parasitic on natural language; there are countless beliefs that people cannot have unless they speak a natural language. Semantics, and perhaps more generally, linguistic competence, play a crucial role in belief formation, the individuation of the content of beliefs, and so on and so forth. Therefore, social explanations that rely on the articulation of people's beliefs, attitudes and the like, are inevitably constrained by the constitutive role that language plays in making those beliefs and attitudes possible and shapes their content.

Let me demonstrate with an example from Hart's account of social rules. For a social rule to exist, Hart maintained, certain things must be true: It must be the case that there is some regularity of behavior manifest in the overt

conduct of a certain population, and this conduct must be accompanied by certain types of beliefs and attitudes that most individuals in that population share about it. For this kind of reduction to be successful, it has to be shown that the explanation is complete, and that the elements offered are sufficient to distinguish social rules from other, similar type of phenomena.

Critics have long argued, however, that Hart's explanation is not complete.³⁸ In particular, it cannot distinguish between a case of following a social rule and one of complying with a generally or widely accepted reason. People often exhibit a regularity of behavior because they believe that there are some reasons that apply to the circumstances at hand, and thus they comply with the reason and exhibit the same kind of attitudes that Hart associates with the "acceptance" of a rule. But, nevertheless, it is not a rule they follow. The textbook example given here concerns strategies followed in games: Players are often expected to know, as they do, that there are certain things not done in a game, for example, because it would be a bad strategy. So, by and large, they do not do them, and would criticize others who do, etc. However, such widely recognized strategies do not form part of the rules of the game. If you fail to comply, you might be criticized for being foolish or amateurish, but not for violating a rule.³⁹

What Hart may have missed here (perhaps, I am not claiming that he did) is the fact that following a rule requires the attribution of a reflexive type of belief to those whose rule it is. The difference between complying with a widely believed strategy and complying with a rule consists in the fact that participants regard the existence *of the rule* as a reason for their action – which is to say that we cannot account for the relevant type of belief people have without invoking the very concept that we try to explain.

Is this as serious a worry as it seems? Let's see where we stand. To account for the existence of a social rule, we must attribute certain beliefs to those whose rule it is. One of those beliefs must be a belief in the existence of a rule, which means that participants must understand the concept of a rule to have the relevant kind of belief. As I argued earlier, however, to understand the concept of a rule is not different from knowing what the word (in the relevant context) means and how to use it correctly. Now, this would be a serious problem for Hart's account if we thought that his theory about the nature of social rules was meant to be a conceptual definition. Then we would have to admit that he was trying to define what the concept of a rule is in a way that

³⁸ Myself included, see my *Positive Law & Objective Value* at 2-4, and reference there.

³⁹ The game example is taken from Warnock, *The Object of Morality*, 43-6.

presupposes that people understand the concept. But I do not quite see it that way. What Hart was trying to do is to provide an account of *social* rules, not a conceptual definition of what rules are. This account purports to identify the kind of beliefs and attitudes people must have when following a social rule. One of those beliefs may well require that participants know, at least roughly, what the word “rule” means in their natural language. In other words, I do not think that it was part of Hart’s project to provide a conceptual analysis of “rules,” or “social rules” for that matter. The project was to show what facts constitute the existence of *a social* rule in a given population, and to show that we do not need to have more in our toolbox for this purpose than the elements he identified, namely, regularities of behavior coupled with certain types of beliefs and attitudes. The fact that beliefs and attitudes depend on natural language may well show that the kind of reductive explanation Hart offered cannot go all the way down; at some point it hits the bedrock of semantic competence shaped by the natural language we speak.

It is crucial to note, however, that the semantic competence we must assume here does not commit us to agree with the views, no doubt often partial and incomplete, that people have about the nature of the things the relevant words stand for. People who have reasonable knowledge of English use the word “rule” correctly on the majority of occasions, though few would be able to define with any precision what the word means or stands for; even philosophers find that task difficult. Furthermore, we often use words correctly even if we know (or believe, if you prefer) that the word stands for something that does not exist, or that it stands for a confused idea. (I find the idea of God profoundly confused and incoherent, but I am confident that I know what the word means and I can certainly use it correctly. You may have other examples in mind.)

In other words, the constraints imposed by the need to rely on semantics do not entail that we have to take all that much for granted. It only means that we must be aware of the fact that the semantics of a natural language play a very important role in belief formation. Language, in other words, imposes certain constraints on social explanations, but it is not the main object of them. It certainly does not mean that philosophical problems are conceptual problems, or that conceptual analysis forms the relevant methodological foundations. Every reductive explanation presupposes that the class reduced to

is, in some relevant sense, more foundational than the reduced class; but it does not have to assume that it is as foundational as it gets.⁴⁰

⁴⁰ I am greatly indebted to Nicos Stavropoulos and Wil Waluchow for helpful comments on earlier drafts of the paper.

REFERENCES:

Dworkin, R. M., *Law's Empire* (Harvard, 1986)

---- "The Model of Rules I" in *Taking Rights Seriously* (Duckworth, 1977), chapter 2.

Greenberg M., "The Standard Picture and its Discontents", in Green & Leiter (eds.), *Oxford Studies in Philosophy of Law*. Oxford University Press, 2011.

Hart, H.L.A., *The Concept of Law* (Oxford, 1961) and *Postscript* in 2nd ed, (Oxford 1994)

Marmor, A. *Philosophy of Law* (Princeton, 2011)

--- *Social Conventions: from language to law* (Princeton, 2009)

--- *Interpretation and Legal Theory* (Revised 2nd ed.), (Hart Publishing, 2005)

--- *Positive Law & Objective Value* (Oxford, 2001)

Raz, J., *Between Authority and Interpretation* (Oxford 2009)

---- *Practical Reason and Norms* (2nd ed, Princeton 2009)

Ryle, G., *Dilemmas* (Cambridge, 1954)

Schwyzler, H., "Rules and Practices", 78 *Philosophical Review* (1969), 451.

Shapiro, S. *Legality* (Harvard, 2011)

Soames, S. *Philosophy of Language* (Princeton, 2010)

Stavropoulos N., "Obligations, Interpretivism, and the Legal Point of View", forthcoming in (Marmor, ed.) *Routledge Companion to Philosophy of Law*.

Warnock, G., *The Object of Morality*, (London, 1971), 43-6.