

The Moral Impact Theory, the Dependence View, and Natural Law

Mark Greenberg, UCLA¹

The central idea of the Moral Impact Theory is that legal obligations are those genuine obligations that obtain in virtue of the actions of legal institutions. (If we formulate this idea in terms of moral obligations, the paradigm of genuine obligations, we get the view that legal obligations are those moral obligations that obtain in virtue of the actions of legal institutions.) What philosophical question does the Moral Impact Theory seek to answer? Facts about the content of the law – for example, the fact that, under California law, contracts for the sale of land are not valid unless committed to writing – are not among the most basic facts of the universe. Rather, such *legal content facts* – *legal facts*, for short – obtain in virtue of other, more basic facts. A core question in philosophy of law is how to explain, at the most fundamental level, the obtaining of legal facts in terms of more basic facts. The Moral Impact Theory is an answer to this question.

The Moral Impact Theory thus offers an account of the more basic facts and how these more basic facts together make it the case that the legal facts obtain. Consider the issue of whether a statute's contribution to the content of the law is, say, the public meaning of the statute, the content of the legislature's communicative intention in enacting the statute, the content of the legislature's legal intention,² the principles that would have best justified the enactment of the statute, or some other factor or combination of factors. What impact, if any, the enactment of a statute has on our genuine obligations depends on reasons that derive from relevant values such as democracy and fairness. The benefits of, for example, coordination, avoidance of harm, stability, and the peaceful settlement of disputes yield many other moral reasons that bear on a statute's impact on obligations. On the Moral Impact Theory, very roughly, a statute's impact *on genuine obligations* constitutes a statute's contribution *to the law*. Therefore, a statute's contribution to the content of the law is determined, at the fundamental level, by all relevant values. More generally, the Moral Impact Theory maintains that, at the fundamental level, what the determinants of the content of the law are — and how they contribute to the content of the law — is determined by all relevant values.

The Moral Impact Theory fits naturally into a particular background understanding of what law is for. On this understanding, law is supposed to improve the normative or moral situation by changing our obligations, powers, rights, and so on. In order to fulfill this aim or function, legal institutions seek to change our genuine obligations (powers, rights, and so on) in certain characteristic ways. Just as individuals do things, such as making promises and giving consent, that affect their genuine obligations, so legal institutions take actions, such as enacting statutes and regulations and deciding cases, that have a more general impact on genuine obligations. According to the Moral Impact Theory, roughly speaking, the genuine obligations that come about in these ways are legal obligations.

The Moral Impact Theory is a member of a family of theories, which I call the *Dependence View*. (I use this term because the central idea is that legal obligations constitutively

¹ I am very grateful to Andrea Ashworth and Nicos Stavropoulos for valuable comments.

² For a brief explanation of what legal intentions are, see below p. 20 & nn. 48-49.

depend on moral obligations.) Members of the family vary along several dimensions; for example, they take different positions with respect to which moral or normative obligations constitute legal obligations.

In recent years, there has been a surge of interest in the Moral Impact Theory and its cousins in the Dependence View family – *Dependence Theories*, for short.³ Elsewhere I have argued for the Moral Impact Theory and the Dependence View more generally, developed the theory, and responded to common objections.⁴ Rather than duplicating those discussions here, this chapter takes on several other tasks. Section 1, after initial clarifications and discussion of dimensions along which Dependence Theories vary, provides an intuitive elucidation of the Moral Impact Theory. Section 2 situates the theory with respect to several themes and positions associated with the natural law tradition. Section 3 addresses a common obstacle to understanding the theory. Section 4 briefly concludes.

1. Elucidating the Moral Impact Theory

In this section, I begin by clarifying two aspects of the Moral Impact Theory: the nature of the obligations to which the theory appeals and the level at which the theory's account operates. First, the Moral Impact Theory maintains that legal obligations are the *genuine* practical obligations that obtain in virtue of the actions of the relevant institutions. To be clear, it is not that there are two kinds of obligations, genuine ones and non-genuine ones. Rather, the point of the qualification *genuine* is to contrast my usage of *obligation* (and of other normative terms, such as *reason*, *requirement*, *right*, and so on) with what we might call the sociological usage. A sociologist might say, for example, that Jews are obligated not to eat meat with milk. Here, the sociologist does not intend to take any position on whether Jews are in fact under an obligation not to eat meat with milk. The sociologist is merely reporting that Judaism takes there to be such an obligation. Roughly speaking, to say that there is an obligation to Φ in the sociological sense is to say that, on the best understanding of the relevant group's practices, the group regards Φ ing as obligatory.⁵ This may not be as simple as a majority of the group's believing that Φ ing is obligatory or the group's authoritative texts' so specifying. For the group may have sophisticated practices that are best understood as recognizing an obligation to Φ , despite what members of the group believe or authoritative texts say. For convenience, rather than saying that there is an obligation in the sociological sense of the term, I will say that there is a *paper obligation*. It bears emphasizing that mere paper obligations are not obligations at all, for the mere fact that a group's practices are best understood as recognizing an obligation to Φ does not make it true that there is such an obligation.

³ See Berman forthcoming; Dworkin 2011; 2013; Greenberg 2004/2006; 2006; 2007; 2011a; 2011b; 2014; Hershovitz 2015; Sager forthcoming; Schaus 2015; Stavropoulos 2012; 2016; Waldron 2013.

⁴ See my works cited in the previous note.

⁵ A claim that there is an obligation in the sociological sense has to be relative to the perspective of a group. It is often clear from the context which group's perspective is being assumed. Also, it is worth noting that a group may take an obligation to fall on agents who are not members of the group, and need not take an obligation to fall on all members of the group.

Genuine obligations are simply real obligations – they are genuinely binding. Moral obligations are the paradigm of genuine practical obligations (though a moral skeptic could hold that moral obligations are not genuine). There may be other kinds of genuine practical obligations. For example, some theorists hold that there are prudential obligations, and, depending on how narrowly the moral domain is understood, there may be other practical obligations that are not moral, such as all-things-considered normative obligations. I understand the moral domain broadly, with the consequence that, on my view, genuine practical obligations and moral obligations are roughly interchangeable for the purposes of the Moral Impact Theory.⁶ Consequently, I will not be careful to distinguish between them in this chapter, sometimes writing of moral obligations and sometimes of genuine obligations.⁷ Readers who take genuine obligations to be importantly different from moral obligations can substitute the former for the latter throughout. I call the theory “*The Moral Impact Theory*” rather than, say, “*The Genuine Impact Theory*” because the former name more effectively conveys the central idea in a few words.

(The content of the law includes not just obligations, but also powers, rights, privileges, and so on. Just as the Moral Impact Theory takes legal obligations to be the moral obligations that obtain in virtue of the actions of legal institutions, it takes the content of the law more generally to be that part of the moral profile⁸ that obtains in the virtue of actions of legal institutions. For convenience, however, I often write simply of obligations.)

Many legal theorists and lay people alike assume that legal obligations are merely paper obligations. On this assumption, to say that there is a legal obligation in the United States to file one’s federal tax return by April 15 is to say that the relevant legal practices and materials recognize such an obligation. By contrast, the Moral Impact Theory takes legal obligations to be genuine obligations.⁹

We come now to the second clarification, which concerns the level at which the Moral Impact Theory offers an account of the determination of legal content. I said at the start that the Moral Impact Theory seeks to give an explanation of the obtaining of legal facts *at the most fundamental level*. This idea should be familiar from discussions of inclusive legal positivism. For example, on H.L.A. Hart’s positivist theory, the content of the law is determined at the fundamental level by convergent practices of judges and other officials – what Hart calls the rule of recognition.¹⁰ At this fundamental level, according to Hart, moral facts can play no role. But Hart maintained that moral facts may play a role in determining the content of the law at a less basic level, for the rule of recognition may specify that moral facts do so.¹¹

The notion of how the content of the law is determined at the fundamental level is crucial to understanding the Moral Impact Theory and its competitors. Suppose that the content of the law is determined at the fundamental level by the convergent practice of judges and other

⁶ See Greenberg 2011a: 82-84; 2014: 1306-08.

⁷ I will generally omit the qualification *practical*, as practical obligations are the only ones relevant to the discussion. Epistemic obligations are an example of non-practical obligations.

⁸ I use the term “moral profile” for the panoply of moral obligations, powers, rights, and so on.

⁹ In what follows, I will use terms such as “obligation,” “requirement,” “standard,” and “norm” for genuine obligations, requirements, standards, and norms, not merely paper ones.

¹⁰ For brevity, I will usually write “the fundamental level” rather than “the most fundamental level.”

¹¹ Hart 1994: 250-54.

officials in the way that Hart's positivism claims. To say that this is how the content of the law is determined at the most fundamental level is to say that *there is no further factor that determines that the content of the law is determined by the practice of judges*. For example, it is not that there are moral reasons that make it the case that the practice of judges is what matters. Rather, the most basic explanation of why statutes, judicial decisions, and the like have the impact that they do on the content of the law is that judges have a practice of treating those items as having that impact.¹²

That the content of the law is determined in a particular way at the fundamental level is consistent with its being determined in a different way at the surface level of ordinary practice.¹³ Suppose, again, that Hartian positivism is correct. And suppose that the practice of judges in a particular legal system is to treat statutes and other relevant materials as making the contribution to the content of the law that is supported by all of the relevant normative, including moral, factors. These normative factors therefore determine, at the surface level, the relevance of statutes and the like to the content of the law. The explanation of why the content of the law depends on normative factors is ultimately the convergent practices of judges. It is because of these practices that normative factors have the relevance that they do, not the other way around.

Of course, the converse is also a possibility. Suppose that, at the fundamental level, the content of the law is determined in the way specified by the Moral Impact Theory: roughly, the actions of legal institutions have the impact that is supported by all of the relevant values. But suppose that the upshot of these values in the circumstances of the legal system is that the actions of legal institutions should have the impact that is given to them by the convergent practice of judges (perhaps because of the importance of stability and protecting expectations). In other words, in this case, Hartian positivism is false — the practice of judges is not the fundamental determinant; rather, the fact that the practices of judges have the decisive say at the surface level is determined by the relevant values.

It bears emphasis that to say that the content of the law is determined in a particular way in a legal system (whether at the surface level or the fundamental level) is not to make a claim about what participants in the system do, such as a claim about how they in fact go about ascertaining what the content of the law is. Rather, it is to make a claim about what *makes* the content of the law obtain. A claim about how the content of the law is determined is therefore closer to a claim about how it would be legally correct to go about figuring out what the law is than to a claim about how judges actually do it. The point especially needs to be emphasized at the surface level. For example, suppose that the content of the law at the fundamental level is determined by moral values, and those moral values make it the case that, at the surface level, a

¹² There is a deep question concerning what kind of philosophical explanation, if any, there can be of why the content of the law is determined, at the fundamental level, in the way that it is. Such an explanation would have to appeal to factors that are not themselves determinants of the content of the law. For example, assuming that Hartian positivism is true, suppose that a putative explanation of why the practice of judges is the fundamental determinant of the content of the law appealed to a certain factor. That factor could not be a determinant of the content of the law, for, if it were, it would be more fundamental than the practice of judges. I hope to address the question raised in this paragraph elsewhere, but for present purposes I set it aside. I will assume throughout that the kinds of explanations that I am discussing of how the content of the law is determined are ones that appeal to determinants of the content of the law. There can be no explanation of this sort of why the content of the law is determined, at the fundamental level, in the way that it is.

¹³ There can be intermediate levels between the fundamental level and the surface level, but I will set aside this complication, as it should not affect the main points.

statute's contribution to the content of the law is determined by its public meaning. Nevertheless, judges may in practice mistakenly treat a statute's contribution to the content of the law as determined in some other way.

One consequence of the foregoing discussion is that, for at least two reasons, one cannot straightforwardly read off which fundamental theory of the content of the law is correct by looking at how lawyers and judges work out what the law is. First, judges and lawyers may be wrong about how the content of the law is determined. Second, even if their everyday practice accurately reflects the way in which the relevant factors contribute to the content of the law at the surface level, the way in which the content of the law is determined at that level may be very different from the way in which it is determined at the fundamental level.

With these clarifications out of the way, we can turn to the Dependence View. I indicated above that the Moral Impact Theory is one member of a family of theories. The underlying idea that unites these Dependence Theories is, very roughly, that the content of the law is constituted by that part of the moral profile that obtains in virtue of the actions of legal institutions.¹⁴ Different Dependence Theories develop this idea in different ways.

Most importantly, Dependence Theories differ with respect to which part of the moral profile constitutes the content of the law. The actions of legal institutions result in many moral obligations that are intuitively not legal obligations. In one kind of case, legal institutions put people in danger, thereby generating a moral obligation to help those people. More generally, legal institutional action that makes the moral situation worse may thereby generate obligations to remedy, oppose, or otherwise mitigate the consequences of the relevant action. Call this general way of changing the moral profile *paradoxical* (because the resulting obligations run in the opposite direction from the standard case). A theory that had the consequence that paradoxically generated moral obligations constituted legal obligations would be very implausible. In another kind of case, legal institutions' actions lead through an extended chain of events to moral obligations that are intuitively too far downstream to qualify as legal obligations. For example, the actions of legal institutions in setting up a tax system might lead me to promise to help you fill out your tax returns. As a result of my promise, I may have a moral obligation to help you. Although this moral obligation traces back (rather indirectly) to the actions of legal institutions, intuitively it is not brought about in the right way to constitute a legal obligation.¹⁵

The Moral Impact Theory's approach to this issue begins from an intuitive notion that the moral obligations in question are not brought about in the right sort of way to constitute legal obligations. Intuitively, for example, bringing about changes in the moral profile paradoxically is not the way in which legal systems, by their nature, are supposed to change the moral profile. Thus, the Moral Impact Theory holds that there is a distinctive way or ways in which legal institutions are, by their nature, supposed to create obligations. Call this *the legally proper*

¹⁴ For convenience, I continue to write of the *moral* profile and of *moral* obligations, but as I explained above, Dependence Theories can maintain that the relevant obligations are not moral obligations, but, for example, genuine practical obligations of some other kind.

¹⁵ The issue raised in this paragraph concerns moral obligations that come about as the result of legal institutions' actions, but are not legal obligations. There is also a converse issue. Legal obligations not to assault, murder, and rape, for example, seem to correspond to pre-existing moral obligations. So it might be objected that the relevant moral obligations do not obtain in virtue of the actions of legal institutions. For discussion of this issue, see Greenberg 2014: 1319-21.

way.¹⁶ The obligations that come about in this way constitute legal obligations. Obligations that come about in other ways do not.

In order to develop a principled account of how the law is supposed to change the moral profile, the Moral Impact Theory appeals to law's essential properties, including what law, by its nature is supposed to do. For example, as mentioned above, the Moral Impact Theory fits naturally into a background view that the law, by its nature, is supposed to improve the moral situation by solving certain kinds of moral problems. If law, by its nature, is supposed to improve the moral situation, then a method of changing the moral profile that works by making the situation worse, thereby creating reasons to undo (or mitigate) what a legal institution has wrought, is a defective way of generating obligations.¹⁷ I have not offered a complete account of the legally proper way, but I have illustrated the idea that appeal to law's nature can do the necessary theoretical work, excluding putative counterexamples in a principled way.¹⁸ Later in this section, I give examples of the legally proper way of changing the moral profile.

In a brief final chapter of his late work, *Justice for Hedgehogs*, Ronald Dworkin apparently adopts a version of the Dependence View.¹⁹ The discussion is highly compressed, but he seems to take the position that legal obligations are those moral obligations that courts and similar institutions ought to enforce with coercion.²⁰ "legal rights are those that people are entitled to enforce on-demand, without further legislative intervention, in adjudicative institutions that direct the executive power of Sheriff or police."²¹

Unfortunately, Dworkin's argument for this view – *the Judicial Enforcement Theory* – is extremely weak. He claims that, if law and morality are separate systems of norms, then any attempt to answer the question of whether the content of the law depends on the content of morality must be viciously circular. "Where shall we turn for an answer to the question whether

¹⁶ This terminology should not be taken to suggest that the relevant way of creating obligations involves taking actions that are legally permissible or legally required. Rather, "the legally proper way" is just a label for the way in which obligations have to be generated in order for them to constitute legal obligations.

¹⁷ Such a method could, in a particular case, ultimately improve the moral situation overall, for example by producing a backlash against the legal system, or even a revolution. The paradoxical method is not a reliable way of improving the moral situation in *normal* circumstances, however. In general, accounts of what makes an object or system with a point or function defective depend on a distinction between normal and abnormal circumstances (for the relevant type of object or system). For example, under certain circumstances, a heart with a leaky valve may be better at circulating blood than a heart without a leak. But a heart with a leaky valve is nevertheless defective because such a heart will not generally be effective at circulating blood under the circumstances that are normal for hearts.

¹⁸ For discussion of the legally proper way, see Greenberg 2014: 1321-23.

¹⁹ In section 2 below, I compare Dworkin's well-known "law as integrity" theory of law, presented in fully developed form in *Law's Empire*, with the Moral Impact Theory.

²⁰ It is unclear whether Dworkin means to take the position that legal obligations are simply those moral obligations that should be enforced with coercion or the position that legal obligations are those moral obligations *that obtain in virtue of actions of legal institutions and* should be enforced with coercion. (On the one hand, he offers an illustration involving a family that seems to suggest the relevance of institutional action; on the other hand, as illustrated by the quotation in the text, he does not explicitly make reference to the relevance of institutional action.) He seems to take the former position in his posthumous article, "A New Philosophy for International Law," *Philosophy and Public Affairs*, 41, 2, 12 (2013). For discussion, see Greenberg 2014, 1299-1300 note 28.

²¹ Dworkin 2011: 406.

positivism or interpretivism is a more accurate or otherwise better account of how the two systems relate? Is this a moral question or a legal question? Either choice yields a circular argument with much too short a radius” (Dworkin 2011: 403). The two possibilities that he assumes are exhaustive are that the content of the law or the content of morality specifies the relation between law and morality. He argues that either possibility is question-begging. There is, however, no reason to think that legal and moral norms are the only place to turn for an answer to the question of the relation between law and morality. Indeed, the question is not one that we would expect to be addressed by legal or moral norms at all, for it is not a question about what law or morality requires or permits. Rather, it is a metaphysical question about the relation between the two domains.²² To make matters worse, having just argued that, if law and morality are separate systems of norms, any attempt to investigate their relation must beg the question, Dworkin’s proposed way forward is to “treat law as a part of political morality” – a starting point that blatantly begs the question against his positivist opponents (Dworkin 2011: 405).

Not only is the argument weak, but the Judicial Enforcement Theory also suffers from the problem that there seem to be legal obligations that the courts and similar adjudicative institutions should not enforce, for example because of the institutional limitations of the courts, their relations with other branches of government, and the like.²³ Some obligations are better enforced by other branches of government or by more political means. Such obligations are not ones that people are “entitled to enforce on-demand, without further legislative intervention” (Dworkin 2011: 406).

Moreover, an account of law should help us *explain* why courts should enforce some rights and obligations and not others. An important part of the reason why a court should not enforce, say, the moral duty not to lie is that it is not a legal duty. Indeed, it seems plausible that the fact that an obligation is not a legal obligation is, as a rule, a sufficient explanation for why a court should not enforce it. But the Judicial Enforcement account, because it explains what is law by appeal to what the courts should enforce, excludes the possibility of such an explanation. (To elaborate: suppose, with the later Dworkin, that legal duties are simply those that courts should enforce; in that case, to offer to explain why courts should not enforce the moral duty not to lie by appealing to the fact that is not a legal duty would be to offer to explain why courts should not enforce the duty by appealing to the very fact that courts should not enforce it!)

²² Unsurprisingly, the major philosophical attempts to answer that question have not focused on the content of moral or legal norms. Dworkin elsewhere takes a confused position that seems to deny the possibility of metaphysical or philosophical claims about morality and law that are not first-order normative claims. Dworkin 2011: 40-68. If such a position were true, it would indeed be hard to see how one could answer the question of the relation between law and morality – but the problem would not be the circularity one that Dworkin sketches in the last chapter of *Justice for Hedgehogs*.

Dworkin also suggests that an analysis of the concept of law cannot answer the question “unless that concept can sensibly be treated as a criterial ... concept.” Dworkin 2011: 404. But even if we granted for purposes of argument Dworkin’s controversial ideas about concepts, he offers no argument that conceptual analysis is the only way to address the relation between law and morality. And his argument that law is an interpretive concept (Dworkin 1986: 31-113) is not a quick circularity argument for the conclusion that law is a subdomain of morality, but an elaborate and controversial argument for his earlier *Law’s Empire* position.

²³ I make this objection in Greenberg 2014: 1299-1300 n. 28. There is a large literature on the question of when courts should not enforce legal obligations. For a seminal and powerful argument that there are legal obligations that courts should not enforce, see Sager 1978.

Although an adherent of the Judicial Enforcement Theory cannot explain why a court should enforce an obligation by appealing to the obligation's being legal, she could do so by appealing to the obligation's tracing back in the right kind of way to legal institutional action. On the other hand, if the explanation of why courts should enforce an obligation is that the obligation traces back to legal institutional action, it is natural to wonder whether what makes an obligation *legal* is also that it traces back to legal institutional action, rather than that courts should enforce it. (I don't mean to suggest that the answer is obvious. One could coherently maintain that, though institutional origins explain an entitlement to judicial enforcement, it is the latter – and not the former – that makes the obligations legal.)

There are several options for improving on the Judicial Enforcement Theory. First, rather than saying that the relevant obligations are ones that people are entitled to enforce on-demand (i.e., that courts must enforce when requested to do so), we could say that the relevant obligations are those that courts have *pro tanto* reason to enforce. Perhaps better, we could say that the relevant obligations are those that, other things being equal, it is *permissible* for courts to enforce or that courts are *entitled* to enforce. Second, we could broaden the focus from adjudicative institutions to legal institutions generally, thus allowing for the possibility that, say, some constitutional rights are only enforceable by legislative action. Thus, for example, one might maintain that the relevant moral obligations or rights are those that it is permissible, other things being equal, for legal institutions to enforce.

A variant of the Moral Impact Theory makes use of Dworkin's suggestion about enforcement (with improvements just sketched). According to this hybrid view, the relevant moral obligations or rights are those that it is permissible, other things being equal, for courts – or, alternatively, legal institutions more generally – to enforce *because of the way in which the obligations or rights trace back to legal institutional action*. This theory makes room for the possibility of legal obligations that courts should not enforce. And it adopts the intuitive direction of explanation, holding that the reason courts should not enforce certain moral obligations is that they lack an appropriate legal origin.

The Moral Impact Theory and the Judicial Enforcement Theory both focus on the actual moral profile, and specify which part of it constitutes the content of the law. A second branch of the Dependence View family takes a fundamentally different approach, holding that the relevant part of the moral profile is that part of the moral profile that the (actual) actions of the legal institutions were *supposed* to yield. Let's call the moral profile that the actions of legal institutions were supposed to yield *the target profile*. There are a number of options for developing the idea of the target profile. For example, the target profile could be understood in terms of a counterfactual – as the moral profile that *would* have resulted from the actual actions of the legal institutions under ideal conditions. In order to develop a viable version of the Dependence View along these lines, it would be necessary to develop an appropriate account of ideal conditions. Space restrictions permit me to make only a few brief remarks.

If the legal system is supposed to solve certain moral problems by changing the moral profile, ideal conditions are ones under which the legal system or government is best positioned to change the moral profile. For example, ideal conditions will be ones under which the government is best positioned to make violence by private citizens morally impermissible, to

make commitments that bind all of its citizens, to set up arrangements in such a way that fairness will require all to participate, and so on. In more familiar terms, ideal conditions are conditions under which the government or legal system is maximally legitimate.

There is an intuitive way of understanding this ideal-conditions version of the Dependence View. It is natural to think that any legal system presupposes its own legitimacy. One way to make this idea more precise is this: for a legal system to presuppose its own legitimacy is for it to treat as actual the moral profile that would obtain if the system were perfectly legitimate. So an ideal-conditions Dependence Theory can be seen as a way of spelling out what it means for the law to presuppose its own legitimacy.

I have focused on the way in which different Dependence Theories specify which obligations constitute legal obligations. There are other aspects of this issue. For example, as noted above, Dependence Theories can take the obligations that constitute legal obligations to be genuine practical obligations or some more restrictive class of genuine obligations, such as moral obligations narrowly conceived. Dependence Theories can also differ with respect to whether the relevant obligations are all-things-considered or *pro tanto* obligations.²⁴

Dependence Theories can vary along several other dimensions. I will mention only that of scope. At one extreme, a Dependence Theory can purport to apply to all possible legal systems. At the other, a Dependence Theory can limit its scope to one legal system. In between these poles, the Moral Impact Theory makes a claim about a theoretically interesting class of systems that includes the legal systems of, for example, the United States and the United Kingdom. The theory makes no claim that, in ordinary usage, the term *law* or *legal system* is applied only to members of this theoretically interesting class. Nor is it part of the theory to claim that the term is best restricted to members of this class.

With these preliminaries out of the way, we can now provide an intuitive elucidation of the Moral Impact Theory. I mentioned in the introduction that the Moral Impact Theory fits naturally into a particular understanding of law. On this understanding, the law seeks to solve certain *moral problems* by changing the moral profile. In the absence of law, it would often be better if our obligations, powers, and so on were different from what they are, but this change cannot come about without some kind of law-like institution.

For example, many problems cannot be solved without complex cooperation, but cooperation will often not be morally required (or practically feasible) in the absence of law. Morality plausibly includes some obligation to help others, but it does not specify exactly what help should be provided to whom. Morality generally requires us to keep promises, but it leaves a lot of questions unresolved. (Which promises are not binding? When is the obligation to keep a promise outweighed?) Morality is arguably highly indeterminate about the conditions under which one acquires property rights and about the bundle of rights to which a property holder is entitled. Morality independent of law may allow people to acquire property rights simply by being first to stake a claim. But it would arguably be better to have a more nuanced scheme of distribution.

²⁴ For discussion, see Greenberg 2014: 1307 n. 41.

Prelegal morality is especially problematic with respect to remedies. What remedy must be provided when a promise or agreement is broken, either justifiably or unjustifiably? When one person's activity causes harm to another? Morality may also allow a wide latitude for self-help, though it would be better if there were a central source of remedies and the permissibility of self-help were quite limited.

A different kind of problem is that self-interested behavior by unscrupulous people needs to be deterred by the threat of punishment. But given the great moral importance of advance notice of punishment, the indeterminacy – or at least uncertainty – with respect to which punishment is morally appropriate for a specific wrong, and morality's unclarity about who is authorized to administer punishment, punishment is in general morally problematic without law. Another example is the need for settlement of disputes, as even well-meaning people will often disagree about the best course of action. Morality, unsupplemented by law, does not provide institutions for resolving disputes.

Legal institutions have a wide range of tools for changing the moral profile.²⁵ They can protect people from violence and punish those who violate others' rights, thereby making violence impermissible except in very limited circumstances. They can give notice of specified punishments for prohibited conduct, potentially making punishment permissible. Legal institutions have a variety of ways for making it likely that many or most people will participate in one solution to a problem. For example, they can create a mechanism for implementing that solution, threaten free riders with sanctions, and increase the salience of the relevant solution. By making it likely that many people will participate in one solution, legal institutions can create powerful reasons for participating in that solution. Democracy provides another tool. The fact that a decision is reached by a system of governance in which people have the opportunity to participate as equals has moral force. Thus, to the extent that the government is democratically constituted, legal institutions can generate democratic reasons.

Most legal theorists presumably would agree that legal institutions seek to create genuine obligations, not merely to create paper obligations. But, as noted above, many theorists nevertheless assume that legal obligations are paper obligations. For these theorists, when a legal institution creates a paper obligation to Φ (with respect to the legal system), it has ipso facto created a legal obligation to Φ because a paper obligation of a legal system *is* a legal obligation. If, in creating the paper obligation to Φ , the legal system serendipitously manages to create a genuine obligation to Φ , that is a good thing. But what makes it the case that there is a legal obligation to Φ is the existence of the paper obligation, which is independent of the genuine obligation.

The Moral Impact Theory is distinctive in claiming that the relevant genuine obligations constitute legal obligations. Let me illustrate with an example. Suppose that a community faces a problem for which there are multiple solutions, and there is no agreement on which solution would be better. As things are, though there is a moral obligation to help others, there is no obligation to participate in any particular scheme

²⁵ For discussion of these tools, see Greenberg 2014: 1310-1319 and Greenberg 2016a: 1965-67.

for solving the problem. For one thing, the efforts of many or most people are needed for a scheme to make a difference, and there is no reason to expect that enough others will participate in a particular scheme. Nothing determines which scheme is the one that people should participate in. In addition, there is no mechanism for people to join together in a common scheme.

Suppose that, in response to the problem, the legislature enacts a statute. The statute's linguistic meaning clearly designates a particular scheme, call it *scheme A*. (As there are different types of linguistic meaning, e.g., semantic content, speaker's meaning, public meaning, and so on, we can start by assuming for simplicity that all of them pick out scheme A.²⁶) Executive officials put into place an enforcement mechanism to prevent free riding, and so on. As it happens, however, the scheme that becomes salient and that a majority of people are likely to participate in – scheme B – is subtly but significantly different from scheme A. This could come about in various ways, perhaps because a private company that is a dominant player in the relevant industry misinterprets the statute early on, or because a widespread psychological tendency makes it unlikely that people would adhere to scheme A. In this kind of situation, different reasons support A and B. For example, democratic reasons may militate in favor of A, while reasons of fairness, helping others, and solving the community problem may militate in favor of B. Suppose that, in the circumstances, the effect of the enactment of the statute, all things considered, is to generate a (genuine) obligation to participate in scheme B, not scheme A.

What about *legal* obligation? According to a widespread implicit assumption, since the meaning of the statutory text picks out scheme A, there is a legal obligation to participate in scheme A.²⁷ According to the Moral Impact Theory, by contrast, there is a legal obligation to participate in scheme B.²⁸ It might be suggested that using the term “legal obligation” for the scheme picked out by the meaning of the statutory text promotes clarity. But it is not clear why this is so. In the first place, we already have a term for the linguistic meaning of the statutory text; indeed, we have various more precise terms for different types of linguistic meaning. And, second, since the mere fact that the meaning of an authoritative text corresponds to scheme A does not make it the case that there is a (genuine) obligation to participate in scheme A, it is not clear why using the term “obligation” promotes clarity. It seems at least as clear a use of

²⁶ For discussion of different types of linguistic meaning, see pp. 21-22 below.

²⁷ In section 3, I discuss this implicit assumption and its relation to the Moral Impact Theory.

²⁸ According to Hartian positivism, the relevant legal obligation depends on the convergent practices of judges in the legal system. If, for example, the practice of judges is to respect the linguistic meaning of the text, then Hartian positivism yields the same result as the assumption that the statute's contribution is constituted by its linguistic meaning. If, on the contrary, the practice of judges is to respect the on-balance upshot of the relevant values, then Hartian positivism yields the same result as the Moral Impact Theory. And, of course, the practice of judges might be to treat some other factor or combination of factors as a statute's contribution to the content of the law. If the practice does not converge on how to resolve a conflict between linguistic meaning and the upshot of the relevant values, then Hartian positivism implies that the statute does not yield a legal obligation.

language to use the term “legal obligation” for genuine obligations that obtain in virtue of the enactment of statutes and other such actions by legal institutions.²⁹

If we consider the possibility that different linguistic meanings of the statutory text correspond to different schemes, the point becomes even stronger. Why would it promote clarity to say that there is a “legal obligation” to participate in the scheme that corresponds to, say, the semantic content of the text, rather than the scheme that corresponds to the speaker’s meaning or the “public meaning”? It might be responded that there are reasons, for example of democracy, why one type of linguistic meaning rather than others should have a special status. In that case, however, we need to consider the possibility that, on balance, the relevant reasons give decisive status to something other than linguistic meaning. In any case, the argument that there are reasons for giving special status to one factor or combination of factors points in the direction of something like the Moral Impact Theory, for this argument seems to presuppose that what matters is what the relevant reasons support.

2. Situating the Moral Impact Theory

In this section, I situate the Moral Impact Theory with respect to various ideas that are associated with the natural law tradition. A preliminary point is that, although there is a tradition of natural law theorizing about morality, our concern in this chapter is only with natural law ideas about *law*.

One theme associated with natural law theories of law involves normative or evaluative claims about law. A paradigm is the claim that one is obligated to obey only legal norms that meet some standard, for example a moral or rational one. Another example is the claim that legal norms should conform to moral standards. Such claims do not purport to say anything about the nature or essence of law or legal systems, nor about what determines the content of the law; they are merely normative claims about law. These ideas therefore have little relevance to the Moral Impact Theory.

A second, more interesting type of idea involves the claim that law essentially has a moral or, more generally, normative function, aim, or point. Examples of such functions or aims include that of promoting the public good; of doing justice; of ensuring that coercion not be deployed against individuals except as authorized by past political decisions; and of improving the moral situation (to be discussed below). We can generalize this kind of idea by saying that law, by its nature, is supposed to meet a certain moral or normative standard or, perhaps equivalently, that law is defective *as law* to the extent that it fall short of some such standard. (The claim that law essentially has a certain function, understood in the relevant way, implies that law is supposed to perform that function and is therefore defective to the extent that it fails to do so.) John Finnis’s claim that internally valid law is not legally authoritative “in the focal sense” unless it is morally authoritative belongs in this category. For the claim seems to be that a legal system that does not meet a moral standard is not law in the fullest sense.

²⁹ See Greenberg 2014, 1304-05.

Another example of the idea that law, by its nature, is subject to a normative standard is what I have called *the bindingness thesis* – that legal norms are supposed to be binding (or, at least, that legal norms are supposed to be such that it is morally permissible to force people to comply with them). See Greenberg 2011a: 84-96. It is important to understand that the claim here is not that it is good for law to be binding, but rather that it is part of the nature of law that law is supposed to be binding, and therefore that law is defective as law to the extent that it is not binding. An analogy: it is good if one’s beliefs make one happy, but it is not part of the nature of belief that beliefs are supposed to make their subject happy or have the aim of doing so. By contrast, plausibly beliefs are supposed to be true, and are defective to the extent that they are not. Similarly, it would be nice if one’s clock increases in value or reduces one’s blood pressure, but it is no part of the nature of clocks that they are supposed to do so. But it is the function of clocks to tell time accurately, and an account of the nature of clocks that omits this function is unsatisfactory.

The idea that law, by its nature, is subject to certain standards – in particular that it is supposed to be binding and that it is supposed to improve the moral situation – is, as we will see, a natural part of the background to the Moral Impact Theory. But such ideas do not capture the Moral Impact Theory, for it is centrally a thesis about *what (metaphysically) determines the content of the law*.

A different kind of idea associated with natural law is the one often expressed by the slogan “an unjust law is no law at all.” Another common formulation of this idea is that there is a moral threshold test that a putative legal norm must pass in order to qualify as a valid legal norm. Such formulations imply, or at least strongly suggest, the thesis that *prima facie* or inchoate legal norms exist independently of anything moral, presumably because they have been authoritatively pronounced or otherwise posited. John Finnis seems to hold a strong version of this thesis – that the content of the law is constituted by the content of authoritative pronouncements – with the qualification mentioned above that internally valid law is not legally authoritative “in the focal sense” if it is not morally authoritative.³⁰

Even with the qualification that a (*prima facie*) legal norm must satisfy a moral threshold in order to qualify as law, as non-defective law, or as law in the fullest sense, the acceptance of the idea that a legal norm exists in some form merely because it has been posited, and thus independently of its moral or normative merit, is a basic concession to positivist legal theory. It seems to be widely believed that all natural law theories must make this concession. For example, Fred Schauer complains that a legal norm must be identified before it can be evaluated, and therefore that it promotes clarity to distinguish the existence of a norm from its moral merit:

[A]lthough the positions traditionally described as positivism and as natural law are commonly contrasted, and although the contrast is undoubtedly real in some respects, it turns out that all of those who subscribe to some version of anti-positivism, including but not limited to natural law, have a need for some form of identification of that [the content of certain social directives] which is then subject to moral evaluation. And so long as the alleged anti-positivisms engage in the process of pre-moral identification of

³⁰ See Finnis 1980: 27; 2000: 1606-10.

legal items, then it turns out that they have accepted the primary positivist premises, premises which are not at all about the proper uses of the word “law”, but which are rather about the desirability and necessity of first locating that which we then wish to evaluate.³¹

The Moral Impact Theory rejects the idea that *prima facie* or inchoate legal norms exist or can be identified independently of anything moral. Indeed, it rejects the idea that legal norms exist *even in part* in virtue of having been posited. Its rejection of positivism is thus more thorough than the natural law theories just discussed. How is the Moral Impact Theory able to reconcile the fact that human legal institutions create law with the rejection of any constitutive role for the positing of legal norms? It maintains that the role of legal institutions in the fundamental account of the content of the law is not the positing of legal norms, but the taking of actions that change the relevant circumstances, thereby changing the genuine obligations that constitute legal obligations. Remember that, according to the Moral Impact Theory, the content of the law is the moral impact of certain decisions of legal institutions. Therefore, the very idea of inchoate legal norms existing independently of anything moral is a nonstarter. It is, however, a consequence of the Moral Impact Theory that certain morally bad standards — ones that could not be obligatory — could never be legal norms. As I will explain, this is not because legal norms must pass a moral test. Rather, it simply follows from the Moral Impact Theory’s account of how the content of the law is determined.

To understand this point, it is necessary to distinguish between what morality requires *ex ante* – that is, not taking into account our specific circumstances and past actions – and what morality requires *ex post* – that is, taking into account all of the relevant circumstances and past actions, in particular the actions of legal institutions. According to the Moral Impact Theory, legal obligations are constituted by genuine obligations that obtain in virtue of certain actions of legal institutions. These obligations can be extremely different from what is required by morality *ex ante*, since actions of legal institutions can greatly alter the relevant circumstances, thus altering our obligations. Even standards that are quite morally flawed by *ex ante* standards can become law because they can become morally obligatory given the (unfortunate) circumstances. But there are some standards that can never be morally obligatory. A useful analogy is an agreement. One can agree to an arrangement that is much less than ideally fair. The fact that one has agreed can make it the case that one is obligated, even though what one is obligated to do is very far from what one would have been obligated to do *ex ante*. But some agreements, for example an agreement to murder, do not create an obligation to act as the agreement specifies (though they may still have normative consequences). In sum, the Moral Impact Theory does imply a version of the conclusion that certain morally bad standards cannot be law. But the explanation of this conclusion is not that there is any kind of moral test that putative legal norms must pass.

Another idea associated with natural law is that moral goodness is *sufficient* for a standard to qualify as law; in other words, that the moral goodness of a standard can make it a legal norm. Now, it might be thought that if a theorist rejects the idea that putative legal norms can be identified independently of morality, the theorist must accept that moral merit is sufficient for a standard to qualify as law. But this would be a mistake. The Moral Impact Theory

³¹ Schauer 1996: 43 (footnote omitted).

illustrates the point. Again, the theory holds that the content of the law is that part of the moral profile that obtains in virtue of certain actions of legal institutions. Therefore, only the actions of legal institutions can yield a legal standard – moral merit can never be sufficient. In sum, putative legal standards cannot be identified independently of morality, but moral merit is never sufficient for a standard to be law. The explanation is that legal norms are the moral consequence of actions of legal institutions. So they necessarily obtain in virtue of *both* institutional action and moral principles.

Mark Murphy has offered a characterization of the central idea of the natural law tradition: that, necessarily, law is a rational standard of conduct. On Murphy's interpretation, for law to be a rational standard is for it to be backed by decisive reasons for compliance. Murphy explicates that for an action to be backed by decisive reasons is for there to be sufficient reason to take the action and an absence of sufficient reason not to do the relevant action. So to be backed by decisive reasons for compliance is very close to, if not the same as, all-things-considered binding.³² Murphy distinguishes a weak and a strong version of his natural law thesis. On the weak version, law is defective to the extent that it is not backed by decisive reasons for compliance. According to the strong version, necessarily, all law is backed by decisive reasons for compliance.

How do these theses relate to the Moral Impact Theory and the Dependence View more generally? Murphy's weak natural law thesis evidently belongs in the category of ideas discussed above according to which law, by its nature, is subject to a normative standard. In particular, the weak natural law thesis is very close to the bindingness thesis – that law is supposed to be binding – which, as mentioned, forms a natural part of the background to the Moral Impact Theory.

Murphy's strong natural law thesis holds that, necessarily, all law is backed by decisive reasons for compliance. It does not say anything about what determines the content of the law. (Nor does the weak natural law thesis.) So it does not make the central claim of the Moral Impact Theory that the content of the law is the moral consequence of the actions of legal institutions. (Indeed, since it says nothing about how the content of the law is determined, the strong natural law thesis – and, for that matter, the weak thesis – is consistent with the extreme claim that the content of the law is entirely determined by divine will or morality independently of the decisions and actions of humans.) On the other hand, something in the neighborhood of the strong natural law thesis follows from the Moral Impact Theory.

As explained above, the Moral Impact Theory, unlike the strong natural law thesis, does not make a claim about all possible legal systems. But, for the legal systems to which it applies, it has the consequence that the law is genuinely binding. Thus, the Moral Impact Theory has a consequence that is roughly equivalent to the strong natural law thesis, but with limited scope. More important is the point that neither the strong nor the weak natural law thesis is a thesis about how the content of the law is determined and therefore neither thesis captures the Moral Impact Theory or is a competitor to it.

Ronald Dworkin's influential *law as integrity* or *Interpretivist* theory of law bears some affinity to natural law theories, as his theory gives morality an important role in constituting the content of the law. How does the Moral Impact Theory relate to Dworkin's work? Dworkin's

³² Murphy has indicated in conversation that he agrees with this point.

mature theory, as elaborated in *Law's Empire*, holds that the content of the law is the set of principles that are the best *interpretation* of the legal practices and, on some formulations, the more specific rules that follow from those principles.³³ Central to the theory is Dworkin's distinctive account of interpretation, which applies to much more than law, for example, to literature and art. On Dworkin's account, the best interpretation of a legal system is the set of principles that best fits and justifies the practices of the legal system, including the past actions and decisions of legislatures, courts, and other legal institutions.³⁴ This is not the place for detailed exposition of Dworkin, but the rough intuitive idea is that interpretation yields the principles that are implicit in the practices of the legal system. We can dramatize the idea by supposing that the past actions and decisions of the legal system had all been taken by a single, ideally coherent agent. We then ask which are the morally best general principles this agent could have been acting on. These principles, and the more specific norms that follow from them, make up the content of the law.

Dworkin's theory, like the Moral Impact Theory, is thus a theory of what determines the content of the law and one that gives an important role to moral notions. Moreover, like the Moral Impact Theory, Dworkin's theory does not accept that *prima facie* legal norms come into existence by being authoritatively pronounced, and then must satisfy some kind of moral evaluation.³⁵ But the two accounts of what determines the content of the law are very different. The central contrast is simple. The Moral Impact Theory focuses on the moral consequences of the actions of legal institutions, while Dworkin's theory seeks the principles that are implicit in such actions.³⁶ As explained above, the Moral Impact Theory takes the contribution to the law of the actions of legal institutions to be determined by all relevant normative factors. For example, according to Moral Impact Theory, the relevant question is how, in light of democratic values, fairness, the benefits of coordination, settling disputes peacefully, and any other relevant factors, a statute affects our moral obligations. By contrast, on Dworkin's theory the relevant question is what principles would have justified enactment of the statute.³⁷

³³ The point about more specific rules is that, if the law includes a general principle – say, that no one may have his or her property seized without reasonable procedures – it also includes more specific entailments of that principle, for example, that no one may have his or her car seized without reasonable procedures. This qualification will not matter to our discussion, so I will generally omit it.

³⁴ I have elsewhere argued that the dimension of fit is best understood one aspect of justification, roughly procedural justification. See Greenberg 2004/2006: 263 note 47; Greenberg 2014: 1300 n. 29, 1303 n. 34.

³⁵ See Greenberg 2011a: 58-60.

³⁶ Dworkin standardly formulates his account in terms of legal practices rather than actions of legal institutions, and his understanding of practices encompasses more than the actions of legal institutions. I am simplifying the comparison by framing his account in terms of the actions of legal institutions.

³⁷ I elsewhere (Greenberg 2014: 1301) explicated this intuitive contrast with a metaphor. I wrote that, on Dworkin's theory, the content of the law is *upstream* of the actions of the legal institutions – it is what justifies them – while, on the Moral Impact Theory, the content of the law is *downstream* of the actions of legal institutions – it is the obligations, powers, and so on that come about as a result of the actions of legal institutions. Some objected to this metaphor that, as noted above, Dworkin (at least sometimes) characterizes the content of the law as including not just the principles that justify the legal practices, but also what those principles entail. But this objection misses the mark. The entailments of the principles are not downstream of the actions of legal institutions; rather, they are simply the more specific consequences of the upstream principles. See note 35 above. We must not confuse the specific entailments of the general principles with the moral consequences of the actions of legal institutions. For example, suppose that the principle that best justifies certain distributive decisions by the legislature is that people should be rewarded according to their effort, not according to their needs or productivity. A straightforward consequence of that principle is that John, who worked extremely hard, but as a result of a poor education did not

One salient difference is that Dworkin's account makes Dworkinian interpretation constitutive of the content of the law, while the Moral Impact Theory does not give a constitutive role to interpretation of any sort. (On the Moral Impact Theory, as with most theories of law, the role of interpretation is merely epistemic; that is, interpretation is just a way of figuring out what the law is.) And of course Dworkinian interpretation employs Dworkin's notions of fit and justification, which also play no role in the Moral Impact Theory.

The Moral Impact Theory is also more thoroughly moralized than Dworkin's theory. According to the Moral Impact Theory, legal obligations are constituted by (ex post) moral obligations. Roughly speaking, what is legally required is what is morally required in light of the actions of the legal institutions. On Dworkin's theory, by contrast, it seems likely that legal obligations will often diverge from moral obligations, even ex post ones. If the actions of legal institutions are sufficiently flawed, the principles that best fit and justify them (and the entailments of those principles) will be different from what is in fact morally required in light of the relevant decisions.

Dworkin himself recognizes this point both implicitly and explicitly. The point is implicit in much of *Law's Empire*. Dworkin develops at great length the distinctive and novel account of interpretation that lies at the center of his theory. If he believed that the output of Dworkinian interpretation, i.e., the principles that best fit and justify the legal practices, necessarily coincided in content with something straightforward and familiar – that which is morally required in light of the practices – it would have been peculiar not to say so. But he never suggests that the principles that best fit and justify the legal practices must be moral principles, let alone that they coincide with what is morally required in light of the practices.

More explicitly, Dworkin often explained that, on his account, legal interpretation involves taking those candidate principles that fit the legal practices sufficiently well and asking which is morally best (Dworkin: 1986, 284-85, 387-88; 1997, 340-42). He provides no reason for thinking that the principles that fit sufficiently well are guaranteed to be true moral principles.³⁸

Dworkin also makes the point fully explicit. For example, he recognizes that the content of the law may diverge from what is morally required ex post – indeed, that the content of the law may be so morally flawed that it is “too immoral to enforce.”³⁹

produce much, should receive more than Jill, who worked less hard, but produced more and has greater needs. By contrast, in order to answer the very different question what the *moral* consequences of the relevant decisions of the legislature are, we would need to ask about, for example, whether those decisions were made democratically, whether people relied on them, how people understood the decisions, and so on. The difference between the two theories is larger than the upstream/downstream contrast suggests because, as noted above, the practices that are to be justified encompass more than the actions of legal institutions.

³⁸ Even if, as I have argued, this way of explaining the theory is merely a heuristic or expository device (Greenberg, 2014: 1303 & n. 34), given how often Dworkin relies on it, it seems fair to assume that it does not radically mischaracterize the consequences of Dworkin's theory, converting a theory on which legal obligations are moral obligations to one on which they need not be.

³⁹ *Law's Empire* at 262. In *Justice for Hedgehogs*, Dworkin characterizes Interpretivism as holding that the content of the law includes “the specific rules enacted in accordance with the community's accepted procedures” (2011: 402.) Assuming that he is using the terminology of enacting rules in the standard way, Dworkin is here taking Interpretivism to include in the content of the law the rules specified by the authoritative texts. In that case, it is

Moreover, it is difficult to see why Dworkinian principles must coincide with what is morally required in light of the actions of the legal institutions. Here is a way to bring out the intuitive implausibility of the idea that there is such a necessary coincidence. If the coincidence idea were true, then a bad decision would make it the case that one ought to follow the principles that would have best justified that decision. But this implication does not seem to be correct. Suppose that a philosophy department makes an unfair decision about how to allocate resources. That decision, once made, may well affect what it is fair to do going forward. Suppose, next, that a comparable issue about allocating resources arises at a later time. It is not at all obvious that adhering to the principle that best fits and justifies the past unfair decision is now morally justified. The past decision might have been so flawed that, even given the way in which it has changed what is morally required ex post, for example, even taking into account whatever expectations it created, it would be wrong to follow the principle behind that decision. I do not take this quick example to be a sufficient argument that Dworkinian principles need not coincide with what is morally required ex post, but just an indication of why it is far from obvious that they do coincide. Given that neither Dworkin nor anyone else has even attempted to argue for the coincidence idea, I will not address it further.⁴⁰

3. An Obstacle to Understanding – the Standard Picture and the Moral Impact Theory

This is not the place to present extensive argument for the Moral Impact Theory.⁴¹ But I want to address one common obstacle to appreciating the plausibility of the theory. A common implicit presupposition – *the Standard Picture* – is that the content of the law is just, to put it crudely, whatever the law books say. More precisely, the idea is that the content of the law is constituted by the ordinary linguistic meaning of the authoritative legal texts.⁴² More sophisticated versions, recognizing that there are different types of ordinary linguistic meaning, hold that the content of the law is constituted by a specific type of linguistic meaning of the relevant texts. If one presupposes the Standard Picture, then, at first blush, the Moral Impact Theory might seem implausible, or even crazy. For, on standard views about linguistic meaning, the linguistic meaning of a text does not in general depend on the moral or normative consequences of uttering

even more clear that the content of the law diverges from morality, for a polity can enact rules that are no part of ex ante or ex post morality.

⁴⁰ As discussed above, in *Justice for Hedgehogs* and in a posthumous article (Dworkin 2013), Dworkin seems to have converted to a version of the Dependence View that is a close relative of the Moral Impact Theory. Unfortunately, Dworkin makes no attempt to say how this briefly sketched position relates to his developed theory in *Law's Empire*. It is notable that the argument that he gives for the position is very different from the argument of *Law's Empire* and extremely weak. See p. 7 above and Dworkin 2011: 400-15; 2013: 2, 12.

I am very grateful to Nicos Stavropoulos for valuable discussion of the issues in the preceding three paragraphs in the text (although I believe he disagrees with my conclusion). Stavropoulos has proposed an understanding of interpretation according to which interpretivism about the law – the view that the content of the law is constituted by the best interpretation of the legal practices – is a Dependence Theory. See Stavropoulos 2014. Also, after writing this paper, I learned that Stavropoulos in an unpublished manuscript argues that Dworkin's mature theory of law is best understood as a Dependence Theory and claims that Dworkinian interpretation need not be redundant in such a theory. See Stavropoulos MS.

⁴¹ See Greenberg (2004/2006; 2006); 2011b; 2014.

⁴² Some adjustments and qualifications would be needed to yield an adequate account, but we can ignore such details here. For extensive discussion, see Greenberg 2011a.

the text. In this section, I will show that a quick dismissal of the Moral Impact Theory on the basis of the Standard Picture would be misguided. In fact, on the most plausible construal of the Standard Picture, it is not in direct competition with the Moral Impact Theory, and it could be argued that something in the neighborhood of the Standard Picture is actually a consequence of the Moral Impact Theory in the circumstances of familiar legal systems. Thus, the Moral Impact Theory can explain why the Standard Picture seems plausible to many.

It is worth emphasizing that the Standard Picture is a strong substantive view. The content of the law consists of obligations, rights, powers, and the like. By contrast, linguistic meaning is information represented by symbols. These two things are not even of the same general category. And there is no necessary connection between norms and information represented by symbols. After all, there are systems of norms in which texts or other representations have no constitutive role. Morality is one example. And, whatever is the truth about the actual norms of etiquette, a society could have norms of etiquette that obtain in virtue of practices not involving linguistic or other representations.⁴³

We distinguished above between an account of what determines the content of the law at the fundamental level and an account at the surface level.⁴⁴ An adherent of the Standard Picture must clarify whether the view is supposed to concern how the law is determined at the fundamental level or the surface level.⁴⁵ I discuss these options in turn.

I am not aware that any philosopher of law has ever explicitly espoused something like the Standard Picture as an account of law at the fundamental level.⁴⁶ And I suspect that philosophically sophisticated proponents of the Standard Picture, if pressed, would not maintain that it is supposed to be an account at the fundamental level. Rather, most would probably appeal to Hartian positivism – the most widely held contemporary theory of law – as their preferred account at the most fundamental level. (It should go without saying that, if taken as a theory at the fundamental level, the Standard Picture is straightforwardly inconsistent with Hartian positivism, as well as with every other well-known theory of the content of the law. According to Hartian positivism, what matters at the fundamental level is not the linguistic meaning of the authoritative texts, but rather the practice of judges and other officials.)

As I don't believe that adherents of the Standard Picture would take it to be the fundamental account of the content of the law, I will be quick with this implausible possibility. One aspect of the implausibility is the implication that matters could not be otherwise. Now, how implausible this implication is depends on its scope, i.e., to which possible legal systems it applies. On the commonly held view that the fundamental account of what determines the

⁴³ For further discussion of the point that there is no necessary connection between norms and information represented by symbols, see Greenberg MS.

⁴⁴ Again, I set aside complications involving intermediate levels, but the points that I make below about the possibility that the Standard Picture is true at the surface level could be made with appropriate modifications with respect to the corresponding possibility at an intermediate level.

⁴⁵ In part because the Standard Picture is an implicit commitment rather than an explicitly avowed position, it is typically unclear whether adherents take the Standard Picture to concern the surface level, the fundamental level, or an intermediate level.

⁴⁶ I have elsewhere criticized the Standard Picture. See Greenberg 2011a. For criticism of a specific and sophisticated version, see Greenberg 2011b.

content of the law is true of all possible legal systems, the Standard Picture as the fundamental account is especially implausible, for it would have the implication that a factor other than the ordinary linguistic meaning of the authoritative legal texts could not be a determinant of the content of the law in any possible legal system (except to the extent specified by authoritative legal texts).

I want to allow, however, that different fundamental accounts could be true of different legal systems. Therefore, consider the possibility that, though not true of all possible legal systems, the Standard Picture is the correct account of how the law is determined at the fundamental level in the U.S. legal system.⁴⁷ In that case, the Standard Picture is true in virtue of nothing more than the essential facts of the U.S. legal system – the ones that make it this very legal system. Therefore, it could not fail to be true of the legal system. Now, consider a candidate determinant of the content of the law that is not a linguistic content, for example, a legislature’s legal intentions in enacting a statute. (A legislature’s legal intention in enacting the statute is its intention with respect to what change in the law to make, or, to put it another way, what legal norm it intends to enact.⁴⁸ On no colorable theory of linguistic meaning is one’s legal intention in pronouncing a text constitutive of the linguistic meaning of the text.⁴⁹) If the Standard Picture is true at the fundamental level in the U.S. legal system, then even if the long-standing practice of judges were to give legal intentions substantial weight, and even if, say, democratic values clearly supported giving legal intentions such weight, such intentions would not be entitled to weight unless an authoritative legal text so directed.⁵⁰ (Remember that to say that the account is true is not to say that judges in fact behave in accordance with the account, but that, roughly, it would be legally correct for them to do so.) Can we really accept that legal intentions or any other non-linguistic factor could not be a determinant of the content of the law in the U.S. legal system?

It is much more plausible that, if the Standard Picture were true in a given legal system at a particular time, it would be in virtue of something more fundamental – perhaps in virtue of the practices of judges and other officials, or in virtue of the fact that, in all the circumstances, reasons of fairness and democracy support the relevance of linguistic content over other factors. In that case, if, say, the practices of judges or the morally relevant circumstances changed, then the content of the law could come to depend on something other than the linguistic content of the authoritative texts.

There is a more basic point about whether the Standard Picture’s candidate for the fundamental determinant of the content of the law is even the right sort of thing for that role. In

⁴⁷ Adherents of the Standard Picture think that it is true of the U.S. legal system or similar systems, so I will confine my discussion to the U.S. legal system, and set aside the arcane issue of whether the Standard Picture could be true at the fundamental level for some possible legal system.

⁴⁸ The legislature’s legal intention is to be distinguished from the legislature’s intention that the legal norm that it enacts apply or not apply to a particular activity – what we could call its *application intention*. For example, a legislature might enact a statute with the legal intention of requiring that people with contagious diseases be quarantined on entering the country. The legislature might have a conflicting application intention that people with psoriasis be covered by the statute, for the legislature might mistakenly believe that psoriasis is a contagious disease. See Greenberg and Litman 1998: 585-86.

⁴⁹ On the distinction between legal and linguistic intentions, especially communicative intentions, see Greenberg 2011b: 241-44 and 2016b.

⁵⁰ If the Standard Picture were true at the fundamental level, then an authoritative legal text could make legal intentions relevant by so specifying. In the interest of brevity, I will omit this qualification.

order to bring out this point, it will help to recognize that there are different types of linguistic meaning or content. Linguistic contents include, for example, semantic content (roughly, what is conventionally encoded in the text), speaker's meaning (on an extremely influential theory, roughly, what the speaker intended to communicate by means of the hearer's recognition of that very intention), what is said (as opposed, for example, to what is merely implicated).⁵¹ In addition to such notions, which are familiar in philosophy of language, legal theorists often appeal to what they call "public meaning." There is more than one way of understanding what they have in mind. For example, public meaning could be that part of the speaker's meaning that a reasonable hearer would be expected to recognize. Or it could be what a reasonable hearer would take the speaker to have intended to communicate. (On the former notion, it is a necessary, but not sufficient, condition for public meaning that the speaker have actually intended to communicate the information. On the latter, there is no such necessary condition.) Again, it could be what a reasonable hearer would take the speaker to have said or asserted. There are various dimensions that need further specification. For example, what is the reasonable hearer assumed to know about the circumstances in which the text is issued? Is the reasonable hearer supposed to assume that the text was issued by a single rational mind? Depending on how we answer such questions, we get a range of different notions of public meaning. Not all of these notions have theoretical importance in the study of language, and it may be that none of them do. In order to be as generous as possible to adherents of the standard picture, however, we can be liberal about what we count as linguistic meaning for present purposes. (If we ruled out public meaning on the ground that it is not, strictly speaking, a type of linguistic meaning, then the Standard Picture would be that much less plausible.)

Given that there are multiple types of linguistic content, the idea that the content of the law is determined by the linguistic content of the authoritative legal texts is badly underspecified. Which type of linguistic content? Once a precise type of linguistic content is specified, could it really be the fundamental truth about law or a particular legal system that the content of the law is determined by that content — say, that part of what the speaker intended to communicate that a reasonable member of the audience would recognize? It looks *ad hoc* to insist that that's just the fundamental truth about how the content of the law is determined. Why, we are inclined to ask, that particular type of linguistic content, rather than some other one? It seems that there should be an answer to this type of question, perhaps, for example, that reasons of democracy support respecting what the legislature intended to communicate, except to the extent that it is not available to a reasonably informed member of the audience. But, in that case, the fact that a particular type of linguistic content constitutes an authoritative text's contribution to the content of the law is not the fundamental truth about the way in which the content of the law is determined.

The point is about the role of reasons. A deep feature of our legal practice is that claims about the way in which factors contribute to the content of the law are supposed to be backed by reasons. Now, if we are to avoid a regress (or an infinite circle) of reasons, there must be a fundamental answer that does not itself rest on further reasons. Intuitively, however, the reason-based nature of our practice means that not just anything could qualify as the fundamental answer. I have elsewhere argued that normative facts are the sort of thing that could provide the fundamental answer about how the content of the law is determined. In other words, the fact

⁵¹ For discussion of some of the relevant distinctions, see Greenberg 2011b and Greenberg 2016b.

that, on balance, the relevant values support, say, the relevance of public meaning – and that there is no further determinant that makes those values relevant – is an answer that could satisfy the demand for reasons. And, although I have argued to the contrary, it is at least an understandable position that the practice of judges – that *this* is just how we do things in law, or in this legal system – could be the fundamental answer. By contrast, it is hard to imagine that the fundamental story, for which no further reason can be adduced, could be that the content of the law is determined by, say, what a reasonable person would take the legislature to have intended to communicate, except when that conflicts with the legislature’s clearly established legal intentions.⁵²

I turn now to the second option – that the Standard Picture is an account of how the law is determined at the surface level, not the most fundamental level. When we address the way in which the law is determined at the surface level, we need to be specific about which legal system we are considering. As noted, many philosophers believe that how the law is determined at the fundamental level cannot vary from legal system to legal system. But it is relatively uncontroversial that how the law is determined at the surface level may vary, both from legal system to legal system and over time within the same legal system. (For example, if Hartian positivism is true, then the content of the law at the most fundamental level is determined by the practice of judges, and the practice of judges – what they treat as relevant to determining the content of the law – can vary between legal systems and over time.) Thus, the claim that the Standard Picture is an accurate account at the surface level has to be a claim about a particular legal system. I will focus on the U.S. legal system because it is the one I know best and the one most written about, though I believe that my points in the next few paragraphs apply at least to common-law legal systems, if not to other legal systems as well.

If some version of the Standard Picture were true at the surface level in the U.S. legal system, then what account could be true at the most fundamental level? What fundamental story could make it the case that a particular type of linguistic meaning of the authoritative legal texts is what matters at the surface level? I suggested above that many philosophically knowledgeable adherents of the Standard Picture would, on reflection, appeal to Hartian positivism as the theory of law at the fundamental level. But, as I now explain, the claim that the Standard Picture is true at the surface level in the U.S. legal system is in fact inconsistent with Hartian positivism.

On the Hartian account, the rule of recognition – the fundamental determinant of how the content of the law is determined – is constituted by the convergent practice of judges or other legal officials. The basic point is that the contribution to the content of the law made by constitutional and statutory provisions, judicial decisions, and so on is determined by a convergent practice among at least a large majority of judges. For example, if judges: 1) regularly treat statutes’ semantic contents as their contribution to the law; 2) regard this practice as a standard to guide their future conduct; 3) are disposed to criticize other judges who fail to treat statutes’ semantic contents as their contribution to the law; and 4) regard such criticisms as justified, then statutes’ semantic contents are their contribution to the law.⁵³ To the extent that

⁵² In my 2004/2006 and 2006, I elaborate the ideas in this paragraph under the rubric of rational determination. But one need not accept the details of those arguments in order to appreciate the counterintuitiveness of the idea that the Standard Picture could be the fundamental story about how the content of the law is determined.

⁵³ Hart 1994: 55-61, 255.

there is no settled practice with respect to what constitutes statutes' contribution to the content of the law, there is indeterminacy on that issue.

In the U.S. legal system, the practice of judges does not converge on treating one type of ordinary linguistic meaning — or even linguistic meaning at all — as the sole determinant of the content of the law. There is widespread controversy among judges about how the content of the law is determined. More importantly, many judges take into account factors other than the linguistic meaning of the authoritative legal texts: the expected consequences of their decisions, widespread practices in the community, moral values such as democracy and fairness, and the legal intentions of legislators or framers. Moreover, judges do not take these factors into account because they take these factors to be evidence of ordinary linguistic meaning. Nor do they believe that the relevance of these factors derives from the linguistic meaning of authoritative legal texts. That is, it is not that they believe that there is an authoritative legal text that directs judges to treat expected consequences of their decisions or moral values as relevant to the content of the law. Rather, they treat these non-linguistic factors as independently relevant factors in deciding what the law is. Because there is no consensus to be found in the practice of judges and other legal officials that the linguistic meaning of the authoritative legal texts constitutes the content of the law, Hartian positivism is inconsistent with the claim that the Standard Picture is true at the surface level in the U.S. legal system.⁵⁴

Once we take into account that there are multiple types of linguistic meaning, the inconsistency is even clearer. As we have seen, the position that the content of the law is determined by the ordinary linguistic meaning of the authoritative legal texts is badly underspecified. One needs to specify which kind of linguistic meaning. But the practice of judges in the U.S. legal system does not converge on one particular type of linguistic content.⁵⁵

We have been considering the possibility that Hartian positivism is true, so that the fundamental determinant of the content of the law is the practices of judges. What are the implications at the surface level if the Moral Impact Theory gives the true account of how the content of the law is determined at the fundamental level? There is a range of possibilities. For example, given the circumstances of a particular legal system, it could be that the relevant values determine that, at the surface level, the contribution of statutes and other authoritative texts is determined by one specific type of linguistic meaning, perhaps their communicative content. In that case, at the surface level, values play no role. At the opposite extreme, it could be that, on balance, in the circumstances of the legal system, the relevant values do not yield much in the

⁵⁴ I have pointed out elsewhere that there are strong links between the Standard Picture and legal positivism. Greenberg 2014: 1297-99; 2011a: 60-72. For example, the Standard Picture “dovetails with – and indeed fills a gap in – legal positivism” because the Standard Picture seems to offer an easy answer – and one that makes no appeal to moral facts – to the difficult problem of how practices, decisions, and the like determine the unique norms. Greenberg 2014: 1297. Partly as a result, the Standard Picture is widely assumed by legal positivists. Greenberg 2014: 1297-99 & n.23; 2011a 60-66. In fact, HLA Hart himself seemed to assume it. See Greenberg 2011a: 60-61 & n.21, 69. My present point is that, despite the links between the Standard Picture and legal positivism – and regardless of how conveniently the Standard Picture seems to dovetail with legal positivism – in the U.S. legal system at least, the Standard Picture is in fact inconsistent with the most influential form of positivism given the actual practices of judges.

⁵⁵ A closely related point is that Hartian positivism is inconsistent with – or at least in severe tension with – any controversial view of legal interpretation. I develop this point and address potential replies in Greenberg 2016b.

way of broad generalizations at the surface level. The contribution of each statute or other authoritative text depends in a case-specific way on reasons of fairness, democracy, and so on. And there are various intermediate possibilities. For example, it could be that the contribution of statutes and other authoritative legal texts is, as a rule, determined by their communicative content, but that this default is overridden in specific circumstances.

It is notable that, when theorists and judges try to justify treating a particular type of linguistic meaning as relevant, they tend to appeal to normative considerations. That is, they often claim that it is legally correct to treat the preferred type of linguistic meaning as determining the content of the law because relevant values support doing so. For example, a familiar kind of argument is that, because of democratic values, we should interpret the Constitution in accordance with the meaning (understood in a particular way) of the constitutional text, rather than in accordance with the framers' expectations or intentions about how the provision would be applied.⁵⁶ Similarly, theorists and judges appeal to normative considerations to justify diverging from linguistic meaning. For example, a theorist might appeal to democracy to argue that when linguistic meaning conflicts with what the legislature intended to enact, we should not respect linguistic meaning. So there is some evidence that, despite their inclination to mention Hartian positivism when pressed to philosophize, theorists and judges take the relevance of linguistic meaning to be based on normative considerations.

There is a question about how to understand the status of such appeals to normative factors. Are normative factors supposed to be relevant because of some more fundamental determinant of the content of the law or are they supposed to be relevant on their own merits, i.e., at the fundamental level? Typically, appeals to democracy and other normative factors are not framed as if the relevance of those factors derives from other more fundamental determinants. For example, judges who appeal to democratic values do not argue that such appeal is justified by a consensus among judges that democratic values are determinants of the content of the law – again, there is no such consensus, nor does anyone claim that there is. Rather, those who appeal to democracy or the like seem to assume that its relevance is self-standing.

Now, if the relevance of linguistic meaning to the content of the law depends at the fundamental level on normative factors, such as democratic values, it is difficult to see why *all* relevant normative factors should not matter. For example, if one democratic reason militates in favor of public meaning, but other democratic reasons outweigh that reason, it would be hard to argue that respecting public meaning is justified on democratic grounds. Similarly, if reasons of fairness in favor of public meaning are outweighed by democratic values, it would be hard to claim on fairness grounds that respecting public meaning is justified on balance. Consequently, once one takes normative factors to be relevant at the fundamental level, one is driven to the view that the determinants of the content of the law depend on all relevant normative factors.⁵⁷ And this is, in effect, the central thesis of the Moral Impact Theory.

⁵⁶ See Greenberg 2014:1329-1330 & nn.102-104; Greenberg, 2016b.

⁵⁷ There could be a normative argument why, in practice, legal interpreters should take into account normative factors only in a limited way. But it is difficult to see how all relevant normative factors could fail to be relevant to

Moreover, the Moral Impact Theory explains why linguistic meaning has an extremely important role in determining the content of the law. In addition to democratic values and fairness, the benefits of coordination, stability, and peaceful resolution of disputes provide powerful moral reasons for respecting (some type of) linguistic meaning of the authoritative legal texts. Indeed, it is arguable that, in the circumstances of some legal systems, the Moral Impact Theory has the consequence that the contribution of a statutory provision or other authoritative text to the content of the law is always the linguistic meaning (more precisely, a certain type of linguistic meaning) of the provision. (Moreover, one could mount a related argument that, in the circumstances of some legal systems, courts are morally required not to take moral factors into account in working out what the law is.)

There are more complex possibilities as well. For example, it could be a consequence of the Moral Impact Theory that which type of linguistic meaning is a provision's contribution to the content of the law varies depending on the type of legal instrument or other circumstances. Alternatively, the Moral Impact Theory might have the consequence that a provision's contribution is its linguistic meaning, except in specific conditions. Thus, if the Moral Impact Theory is true, it arguably has the consequence that the Standard Picture or some variant thereof is true at the surface level, at least in some legal systems.

Moreover, even if no such argument succeeds – so no general Standard Picture-like position is true at the surface level – the moral impact of the enactment of an authoritative text will often correspond to (some type of) linguistic meaning of the text, or at least something in the neighborhood of that meaning. Thus, if the Moral Impact Theory is true, the contribution of an authoritative text to the content of the law will often be close to its linguistic meaning. Consequently, the truth of the Moral Impact Theory could help to explain why the Standard Picture seems plausible to many.

In sum, if one is tempted by the idea – perhaps with various refinements – that the content of the law is determined by the linguistic meaning of authoritative texts, one needs an account of the content of the law at the most fundamental level that could support such a picture. The Moral Impact Theory provides a reasonably good fit. (And, regardless of whether the Moral Impact Theory actually supports the truth of the Standard Picture at the surface level, it provides an explanation of why the Standard Picture may seem plausible). Thus, an intuitive attachment to something in the neighborhood of the Standard Picture is far from grounds for dismissing the Moral Impact Theory.

4. Conclusion

The Moral Impact Theory (and the DependenceView more generally) has recently attracted a great deal of interest. An important feature of the theory is that, in contrast with many natural law theories, it does not accept that legal norms exist even in part in virtue of being posited. It therefore is more thoroughgoing in its rejection of positivism than much of the natural law tradition. The reason is that the theory's central thesis – that legal obligations are constituted

that more basic normative argument (still assuming that the relevance of normative factors is at the fundamental level, rather than grounded in some more fundamental determinant).

by certain moral obligations – gives no constitutive role to the positing of legal norms. When legal institutions take actions that change the moral profile in certain characteristic ways, they thereby create legal norms. So it is not, for example, that a legal norm comes into existence by being authoritatively pronounced, and then must satisfy some kind of moral evaluation. The Moral Impact Theory also contrasts sharply with Ronald Dworkin’s “law as integrity,” or Interpretivist, theory. The theory rejects Dworkin’s central thesis that a particular form of interpretation is constitutive of law. Instead, the Moral Impact Theory takes the natural position that the *genuine* obligations that legal institutions succeed in creating constitute *legal* obligations. The theory is therefore more thoroughly moralized (or normativized) than Dworkin’s account, as it cannot be assumed that the best Dworkinian interpretation of legal practices consists of principles that are morally obligatory (either *ex ante* or *ex post*). Finally, finding the Standard Picture intuitively appealing should not be an obstacle to accepting the Moral Impact Theory. The most plausible version of the Standard Picture is not a direct competitor of the Moral Impact Theory, and one could make out a case that a qualified version of the Standard Picture is what we would expect at the surface level if the Moral Impact Theory is the true account of how the content of the law is determined at the fundamental level.

REFERENCES

- Berman, M. Forthcoming. "On the Multiplicity of Normative Systems," in Plunkett, D. Shapiro, S., and Toh, K. *Legal Norms, Ethical Norms: New Essays on Metaethics and Jurisprudence*. Oxford: Oxford University Press.
- Dworkin, R. 1986. *Law's Empire*. Cambridge, MA: Harvard University Press.
2011. *Justice for Hedgehogs*. Cambridge, MA: Harvard University Press.
2013. "A New Philosophy for International Law," *Philosophy and Public Affairs*, 41: 2.
- Finnis, J. 1980. *Natural Law and Natural Rights*. Oxford: Oxford University Press.
2000. "On the Incoherence of Legal Positivism," *75 Notre Dame Law Review*: 1597-1611.
- Greenberg, M. 2004/2006. "How Facts Make Law," *Legal Theory*, 10: 157-198. Corrected version reprinted in *Exploring Law's Empire: The Jurisprudence of Ronald Dworkin*, S. Hershovitz (ed.), pp. 225-64. Oxford: Oxford University Press.
2006. "Hartian Positivism and Normative Facts: How Facts Make Law II," in *Exploring Law's Empire: The Jurisprudence of Ronald Dworkin*, S. Hershovitz (ed.), pp. 265-290. Oxford: Oxford University Press.
2007. "The Prism of Rules." Available at <http://ssrn.com/abstract=1042121>.
- 2011a. "The Standard Picture and Its Discontents," *Oxford Studies in the Philosophy of Law*, 1: 39.
- 2011b. "Legislation as Communication? Legal Interpretation and the Study of Linguistic Communication," in *Philosophical Foundations of Language in the Law*. Marmor, A. and Soames, S., eds. Oxford: Oxford University Press.
2014. "The Moral Impact Theory of Law," *Yale Law Journal*, 123: 1288.
- 2016a. "How To Explain Things With Force," *Harvard Law Review* 129: 1932.
- Greenberg, M. and Litman, H. 1998. "The Meaning of Original Meaning," *The Georgetown Law Journal*, 86: 570-619.
- Hart, H.L.A. 1994. *The Concept of Law*. (2nd edition.) Oxford: Oxford University Press.
- Hershovitz, S. 2015. "The End of Jurisprudence." *Yale Law Journal*, 124: 882.
- Sager, L. 1978. "Fair Measure: The Legal Status of Under-Enforced Constitutional Norms." *Harvard Law Review*, 91: 1212.

- Schauer, F. 1996. "Positivism as Pariah," in *The Autonomy of Law: Essays on Legal Positivism*. George, R.P., ed., Oxford: Clarendon Press. 31-55.
- Schaus, S. 2015. "How To Think About Law as Morality: A Comment on Greenberg and Hershovitz," in *Yale Law Journal Forum* 123: 224-45. Available at: <http://www.yalelawjournal.org/forum/how-to-think-about-law-as-morality>.
- Stavropoulos, N. 2012. "Obligations and the Legal Point of View" in Marmor, A., ed., *The Routledge Companion to Philosophy of Law* (Routledge).
- 2014. "Legal Interpretivism," *The Stanford Encyclopedia of Philosophy* (Summer 2014 Edition), Zalta, E., ed., URL = <http://plato.stanford.edu/archives/sum2014/entries/law-interpretivist/>.
- 2016. "Obligations and Social Practice," (in Spanish translation), in Ramírez Ludeña, L. and Vilajosana, J., eds., *Convencionalismo y derecho* (Marcial Pons 2016).
- MS. "Do We Still Need Interpretation?," (unpublished ms, on file with author).
- Waldron, J. 2013. "Jurisprudence for Hedgehogs." Available at: ssrn.com/abstract=2290309.