CHAPTER THREE

Authority, Conventions, and the Normativity of Law

IN THIS CHAPTER I would like to complete the outlines of a plausible version of legal positivism. The chapter is composed of two parts. In the first section I will discuss some of Joseph Raz's ideas about the nature of practical authority and the implications of his views about the normativity of law. In the second section I will return to the rules of recognition and try to show that, though H.L.A. Hart is basically right about the idea that social rules are at the foundations of law, we need a theory of social conventions to articulate the requisite foundations. With these two ideas in hand—the authoritative nature of law and its conventional foundations—we will have the main building blocks needed to reconstruct a plausible version of Hart's theory of law.

Raz's main insight, as noted in the previous chapter, is that the law necessarily claims to be a legitimate authority. There are three lessons I would like to draw from this general insight. First, though Hart is right that legal philosophy should confine itself to an explanation of the normativity of law—without slipping into a moral-political account of what makes law justified or worth having—we can still do better in explaining the normativity of law than just pointing out the fact that people tend to regard legal requirements as binding. As we will see in some detail below, Raz's account of practical authorities gives some structure to the normativity of law—explaining the kinds of reasons that would make legal instructions binding and their possible relation to moral and other normative considerations.

The second lesson I will draw from the authoritative nature of law is that legal norms are basically directives or instructions

issued by an authority aiming to guide the conduct of others. In this respect, I will argue, legal norms are crucially different both from moral norms and from social norms. This will prove to be a controversial thesis, and part of its defense will be taken up only in the next chapter.

Finally, the third lesson is that, in spite of Hart's considerable efforts to detach an understanding of law from political sovereignty, these efforts went a bit too far. John Austin's command theory of law may have been too crude or too simplistic, but his basic insight—that law is an instrument of political sovereignty—is in the right direction. Raz's observation that law is essentially an authoritative institution holds these points together. It invites us to see that there is something unique to the normativity of law and in a way that ties law to political authority much more intimately than Hart's theory maintains. I will not propose any particular argument in support of this last point, but I hope that we will be able to see it as we go along.

AUTHORITY AND NORMATIVITY

Whenever the law imposes an obligation or requires you to do something, it conveys a dual message: You *ought* to do it, and you ought to do it *because the law says so*. When the law prescribes a certain mode of conduct, it purports to make a practical difference that *it is the law* that requires it. If you recall the California signposts about the hands-free mobile phone requirement, the signposts got it exactly right: We ought to use a hands-free device, they remind us, *because* "it's the law!" This is one of the crucial respects in which both moral requirements and social norms are different from law. When you are presented with a moral reason for action that applies to you under the circumstances, or are told that there is a social norm that requires you to do something (say, greet an acquaintance or bring a bottle of wine to the dinner party), it would be rather silly and pointless if you ask, "Who says so?" Nobody does, of course; it is not a relevant question. But in the legal case, it always is. It always matters that it is the law (or some particular legal authority) that says so. One of the

main challenges about the explanation of the normativity of law is precisely to explain this connection between reasons for action and the relevance of the answer to the "who says so?" question. So let us begin with some general observations about these types of reasons for action, and then see how Raz's theory of authority explains this crucial aspect of the normativity of law.

There are several types of cases in which a person may have a reason to do something because she was told by somebody to do it. Sometimes, of course, one has a reason to do what the other recommends or suggests simply because the recommendation is sound on its merits. When I ask my daughter to finish her homework before she goes out to meet with her friends, I expect her to comply with the request because I believe that she has a reason to finish her homework regardless of my asking her to do so. The purpose of my request is simply to remind her of something that she has a reason to do independently of my request. However, if I ask a friend to help me with moving a heavy piece of furniture, I expect the friend to comply with my request, in crucial part because it is my request. I am not suggesting that the friend would have an independent reason to move the furniture whether I ask him to do so or not. Let me call these latter kinds of reasons *identity related*. Somewhat loosely expressed, these would be the kinds of reasons where *A*'s reason to φ partly depends on the identity of another agent, *B*, who suggests, requests, or orders *A* to φ.[1]

There are various situations in which a reason for action is identity related. Some of them pertain to knowledge. Suppose that my broker (if I had one) recommends that I sell my shares in GM because she predicts that their value will plummet. Being ignorant about such matters, as I am, I have a pretty good reason to do what the broker recommends. And the fact that it is my broker who recommends this course of action and not, say, my department chair, is crucially relevant (for example, when asked by

[1] In the literature, these reasons are often called *content-independent reasons*; I find this expression somewhat confusing (because the reasons are not entirely content independent, only partly), and hence allowed myself to introduce the notion of identity-related reasons.

somebody why did I sell my shares, it would make perfect sense to reply that I did so because my broker had recommended that I do so). The assumption is that she just knows better what reasons apply to me and that, by following her recommendation, I am more likely to comply with the reasons that apply to me than by trying to figure it out by myself.

In other cases, however, the reason to do what another person tells you to do has nothing to do with knowledge or expertise. The example of complying with a request of a friend is one such case. The connection is not epistemic; it is not the case that your reason to comply with a request of a friend has anything to do with the fact that the friend knows better what reasons apply to you. The fact that it is your friend who asked you to help constitutes part of the reason to do what he asked because he is your friend, and because the value of friendship is such that there are good reasons to abide by friends' requests.

The law essentially purports to generate identity-related reasons for action. When the law prescribes a certain mode of conduct, it purports to make a practical difference that *it is the law* that requires it. Therefore, it is one of the main questions about the normativity of law: how to explain the rationale of identity-related reasons of the kind the law purports to generate. I think that Joseph Raz has suggested the most plausible answer: The law is essentially an authoritative institution and the reasons to comply with an authoritative directive are, by their very nature, identity-related reasons.

The main challenge facing any explanation of the nature of authority is to make sense of the idea that a person may have an obligation to do something because another person has instructed her to do it. I use the word *obligation* advisedly. There are many situations in which identity-related reasons for action make perfect sense in contexts that have nothing to do with authority. Complying with a request of a friend or acting on the advice of an expert are examples already mentioned. What makes authoritative instructions unique is not that they generate identity-related reasons, though they necessarily do that as well, but the fact that those reasons are of an obligatory nature. If *A* has legitimate authority over *B* in context *C*, then *A*'s authoritative

directive requiring *B* to φ in *C* would normally entail that *B* has an obligation to φ.[2]

Raz's main insight is that the way to justify an obligation to follow an authority's directive is by showing that there are cases in which a person is systematically more likely to comply with obligations[3] that apply to him if he follows the authority's instruction than by trying to figure out (or, act on) those obligations by himself. In other words, an authority is legitimate when it provides a service—the service of making it more likely that, in the relevant area of its authority, the subject would act as he or she *ought* to act if he or she follows the authority's instructions rather than trying to act without the authoritative guidance. Raz calls this the normal justification thesis:

> the normal way to establish that a person has authority over another person involves showing that the alleged subject is likely better to comply with reasons that apply to him . . . if he accepts the directives of the alleged authority as authoritatively binding and tries to follow them, rather than by trying to follow the reasons which apply to him directly.[4]

Admittedly, the details of Raz's account are controversial. In particular, Raz's formulations make it somewhat unclear how his account allows us to distinguish between cases in which reasons for action are identity related, as in the case of following an expert's advice, and genuine authoritative relations, where the authority's instruction constitutes not just identity-related reasons but also an obligation to comply. Furthermore, if we cannot explain how an authoritative directive generates an obligation to comply, we would also lack an account of what gives a putative authority the right to issue such directives—we would lack an account of what gives anyone a right to rule.[5]

[2] *Pro tanto* obligation, not absolute and not all things considered.

[3] Raz uses the word "reasons" in his formulation, not obligations, and he may not agree with my suggested modifications here.

[4] Raz, *The Morality of Freedom*, 53.

[5] Both of these objections have been forcefully advanced by Stephen Darwall, "Authority and Second-Personal Reasons for Acting." In a recent (yet unpublished) article, "The Role of Authority," Scott Hershovitz further develops these

My own view is that Raz's theory makes more sense if the normal justification thesis is confined to facilitating the obligations that apply to the authority's putative subjects; the only way to get the conclusion that there is an obligation to follow the directive of a legitimate authority is to assume that the role of a practical authority is to make it more likely that its subjects would comply better with the obligations that apply to them by following the authority's directive than by trying to figure it out for themselves. Needless to say, this view depends on the availability of a fairly robust distinction between reasons for action and a particular subset of such reasons that constitute an obligation or duty. I am not assuming that we have a very satisfactory account of the distinction; but we should, because it is very intuitive. The idea is that there are countless things we may have a reason to do, but only some of them we also have an obligation to do. Note, however, that if you are doubtful about the availability of such a distinction then you need not worry about Raz's account either. In that case, you will be forced to admit that it does not matter whether an authoritative directive constitutes an obligation or only a reason to comply.

Admittedly, even if we confine the legitimacy conditions of authority to cases in which following its directives makes it more likely that we will comply with obligations that apply to us, we may still lack an account of what gives any putative authority a right to issue such directives, that is, we would still lack a general framework of the idea of the right to rule.[6] I do not find this to be a weakness of Raz's theory; on the contrary, it seems to me much more plausible to maintain that nobody has a right to rule, not even a legitimate authority. Telling other people what to do may be justified under countless circumstances, but I doubt that

critical themes. For my own stab at some of these vexing issues, see my "The Dilemma of Authority" (draft posted on the Social Science Research Network, www.ssrn.com).

[6] Darwall ("Authority and Second-Personal Reasons for Acting") gives a nice example: Even if *A* is morally obliged to invest the family savings according to expert advice, it does not follow that the expert gains an authority to guide *A*'s investments.

it is something that anyone can acquire as a matter of right.[7] (It is possible that particular persons or institutions may have a right to occupy a certain authoritative role, but that is a different matter and one that is typically justified on procedural grounds.) In any case, these details, important as they may be, do not matter for our purposes here, that is, as long as we would concede that the service conception of practical authorities, the basic idea of the normal justification thesis, forms at least a necessary—though perhaps not sufficient—condition for the legitimacy of a practical authority. What we need to explain in the legal context is the rationale of the kinds of normative demands that the law purports to make—that you ought to do *X* and that you ought to do it partly because the law says so.

The first part—that you ought to do *X*—is explained by the idea that the function of authorities is to facilitate our ability to act on the reasons that apply to us anyway, that is, regardless of authorities. The second, identity-related part—that you ought to do *X* because the law says so—is explained by the service conception, or rationale, of practical authorities. The assumption has to be that the authority is somehow in a position to make it more likely that you will comply with what you ought to do by following its directives than by trying to figure it out for yourself. Even if these two conditions are not quite sufficient to explain the range of issues that any theory of practical authority would have to explain, at least they provide the core idea of what would make compliance with a legal directive rational and obligatory.[8]

[7] For a more detailed argument to that effect, see Arneson, "Democracy Is Not Intrinsically Just."

[8] Let me try to clarify a terminological point here: The word "ought" is sometimes used interchangeably with the word "obligation" or "duty." Generally, however, "ought" has a broader and looser use. For example, we often use the word "ought" as an indication of an all-things-considered reason, as in, for example, "I ought to finish this article." The word "obligation" or "duty," however, stands for a much more structured and narrower concept. To say that "I have an obligation to finish the paper" would imply that I have a reason to do it, and one that is structured in a certain way; it is a reason to finish the paper and a reason to exclude certain types of considerations to count against this reason, etc. A full account of the nature of obligations cannot be given here. See, e.g., Raz, *Practical Reason and Norms*.

If this general idea is correct, some important implications follow about the nature of law and its normative character. First, the authoritative nature of law gives considerable support to the idea that legal norms are basically instructions issued by some persons in order to guide the conduct of others. This is a very controversial idea. The objection to it is that norms can be legally valid even if they do not originate with any particular authority. There are two main versions of this argument. One argument, alluded to by Hart in the context of his critique of Austin, concerns the prevalence of legal constraints on authorities. The second argument pertains to the general claim that sometimes we can deduce the content of the law by way of reasoning or moral justification. This latter objection is rather complex and I will leave it for the next chapter, where I will discuss it in some detail. For now, let me answer the first objection.

One of the difficulties Hart raised about Austin's command theory of law concerns the legal constraints on lawmaking authority. If the law simply is the command of the sovereign, how can we explain the fact that, in countless jurisdictions, the sovereign is bound by law. There are constitutional and other legal constraints that curtail the sovereign's authority to make laws of certain kinds or in certain ways. If there are legal constraints on lawmaking authority, how can we say that all law originates with such authorities?

The underlying point of this argument seems to be that an authority cannot be self-binding in the requisite sense. I doubt, however, that the idea of self-binding authority is an absurd one. To begin with, authorities, just like ordinary persons, can make decisions that are binding themselves.[9] The making of a promise (as Hart himself mentions) would be a paradigmatic example. A person who expresses a promise thus undertakes a commitment—one that binds her and constraints her future reasons for action.[10]

[9] The idea is familiar from the literature on precommitment. See, e.g., Elster, *Ulysses Unbound*.

[10] I suspect that Hart failed to see the point here, in spite of his own reference to the promise example, because he must have subscribed to the so-called practice theory of promises, whereby promises only work on the basis of some

But one may still wonder whether Raz's service conception of authority is compatible with an idea of a self-binding authoritative decision. Remember that the idea is that the whole point of an authority is to facilitate better compliance with reasons for action. The assumption has to be that an authority's subject is more likely to comply with the reasons that apply to her if she follows the authority's directive than if she tries to figure out or act on those reasons by herself. How can this rationale apply to the authority itself? In other words, authorities generate identity-related reasons for action; but it seems that identity-related reasons cannot apply to the same agent whose identity is relevant to the reasons generated.

Actually, I think that this objection holds only with respect to the cases in which identity-related reasons rest on expertise. My (imagined) broker is an expert relative to me, for example, and therefore I have reasons to take her advice very seriously; but surely she is not an expert relative to herself. In this sense, it makes no sense indeed to say that a person is an authority vis-à-vis herself. But not all cases involve such epistemic considerations. To take a simple example: Suppose that there is a recurring coordination problem, say in circumstances *C*, that needs to be solved (that is, suppose we *ought* to solve it). Suppose, further, that no solution is likely to emerge naturally, so to speak, unless someone makes a decision and communicates it to the parties concerned. Now, you—one of the parties concerned—happen to be in a position that enables you to make the decision and communicate it to the others: "In circumstances *C we* do *X*"; and then, if everybody complies, the coordination problem is solved. This could be an example of a self-binding authority. You have made an authoritative decision that binds you in exactly the same way, and for the same reasons, that it binds the others. There is nothing inherently absurd about the idea that an authority may issue

conventional practice of promising that prevails in society. I think that Hart assumed that there must be some rules about promising in the background before any speech act of promising can gain the significance that it has. But the practice conception of promising is far from obvious and has been quite convincingly rejected by some philosophers (e.g., Scanlon in his *What We Owe to Each Other*, chap. 7). I have explained some of this in my *Social Conventions*, chap. 5.

directives that bind the authority itself. Under certain conditions, authorities can be self-binding.

I am not suggesting that this is the main rationale of most prevalent legal constraints on lawmaking authorities. The most common constraints on lawmaking authority are constitutional. In most legal systems these days, a written constitution defines the lawmaking authority of various institutions and establishes certain mechanisms for adjudicating potentially controversial cases about such matters. Nothing in the nature of constitutions poses a challenge to Raz's thesis. Constitutions are authoritative directives determining the legal powers of various institutions and often imposing various legal constraints on the exercise of those powers. There is nothing problematic in the idea that one authority can bind or constrain the power of another.

Let me conclude: The fact that lawmaking authority is often legally constrained does not, by itself, count against the thesis that law essentially consists of authoritative directives or instructions. If there is a serious challenge to this thesis, it comes from a different direction; the challenge is posed by those who argue that norms can be legally valid by reasoning about what the law ought to be. I will discuss this challenge at length in the next chapter.

Before we proceed, let me deal with this nagging doubt about our underlying assumption here: What if the law is not more than an organized group of gunmen? Could it not be the case that Austin was right, and the difference between the law and the gunmen is only a matter of scope? One gunman does not make law; but a whole bunch of them, acting in some organized fashion and sustaining control over a certain population, could well be law. Why not? And, if so, just like the gunmen who need not make a claim to be a legitimate authority, perhaps the law need not make such a claim either? It could be argued that Raz's insight that the law claims to be a legitimate authority does not necessarily hold true.[11] Perhaps it is true in most civilized societies, not because

[11] Note, however, that historical examples would not settle this question. Even the most draconian terror regimes that we know from history do not necessarily undermine Raz's argument that the law *claims* legitimate authority. Those horrible regimes and the agents who acted on their behalf tended to claim legitimacy, incredible and outrageous as such claims were. Furthermore, it is worth keeping

it is law, but because they are civilized societies in which governance is generally regarded as requiring moral legitimacy. So the question is, why would a gunmen-type control over a given population not be law?

Note that there are, actually, two different questions here. First, there is the age-old question of whether a regime of sheer terror that has no plausible claim to legitimacy can possibly count as law, or not. I will not purport to provide an answer to this question, mostly because I believe that it is not theoretically important to give one. Some forms of social control might be borderline cases of law. Certain regimes might have some features that would make them legal, and others that would not. Borderline cases are just that, borderline cases. The second and more important question is whether there is something about the nature of law itself that requires it to make a claim to be a legitimate authority. And here I think that Raz's positive answer is correct. As I understand the answer, it consists of two points. The first is that whenever the law makes a certain requirement about the conduct of its putative subjects, it purports to impose the requirement as a matter of obligation or duty to comply; it is the way in which laws are invariably expressed. The second is that the only way to make sense of this kind of obligation is by interpreting it as an instance of an authoritative directive. And here the basic insight is the same as the one noted by Hart: If I tell you that you ought to φ, I have appealed to reasons that apply to you; if I tell you that you should φ *only* because otherwise I will harm you, I have renounced a claim to reasons that apply to you—except the reason to avoid the harm that I might inflict on you. Therefore, whenever the law expresses a demand in terms of what its putative subjects *ought* to do, as the law invariably does, it appeals to reasons, albeit of the identity-related kind. And the best way to make sense of such identity-related reasons in the context of law is by interpreting them as authoritative in nature. (Remember that law's essential

in mind that in any functioning legal system, morally wicked as it may be, the law fulfills some functions that may be quite valuable regardless of the overall iniquity of the regime. And if it does not even do that, then it may be questionable that there is law in that society. It may be a borderline case of law or not law at all.

claim to the legitimacy of its authority can always fail to be true or justified, either wholesale, or in any given case.)

So now we can move to the second lesson I would like to draw from the authoritative nature of law, and it concerns law's normative character. Raz's thesis about the essentially authoritative nature of law gives us the basic structure of the kinds of reasons we may have for regarding the law as binding. It explains the sense in which a legal obligation can be an obligation to do something because the law says so. Raz's account, however, does more than that: It resolves an important aspect of the debate between Hart and Kelsen about the relations between legal obligations and moral ones. According to Kelsen, as noted in chapter 1, the difference between a moral ought and a legal ought is not a difference in the relevant nature of "ought," but only a difference in point of view. Hart's account of the normativity of law, as we saw in the previous chapter, is entirely reductive. It purports to explain the normativity of law in sociological terms, given by the details of the practice theory of rules. This reductionism leads Hart to the conclusion that legal obligations, and other normative aspects of law, can be explained without any reference to morality or moral reasons for action. Whether in any given case there is a moral reason to abide by a legal obligations is, for Hart, purely a moral question that has nothing to do with the nature of legal "ought." When we talk about a legal obligation, we basically describe a complex social reality. When we talk about a moral obligation, we express a judgment about the way things ought to be. Thus, somewhat crudely, Hart would say that, in the moral context, it makes perfect sense to distinguish between what people (in any given population) *believe* that ought to be done, and what ought to be done. In the legal context, if the relevant population believes that there is a legal obligation to φ, then ipso facto, there is such a legal obligation.[12]

You might suspect that this is just another way of expressing the question about the possibility of reduction: Hart maintains that the normative language of the law is reducible to social facts

[12] This is somewhat crude since, even in the legal context, we must make room for the possibility that anyone, including judges, can make a mistake about the law. I will discuss this problem in the next chapter.

(about people's beliefs and attitudes), whereas Kelsen seems to deny the possibility of such a reduction. To some extent, this is true. But let me suggest that Raz's thesis about the authoritative nature of law shows that both Hart and Kelsen have missed something important. Even if Hart is right that legal ought is reducible to facts about people's beliefs and attitudes, one would still have to give an account of what it is that people need to believe in order to make sense of attributing an "ought" to the content of a legal directive.

The point is that the idea of practical authority gives some structure to the rationality of such beliefs. Raz's thesis does not come in to settle the question about reductionism, at least not directly. What it shows is that the way in which we can make sense of a legal ought, or the way in which law is regarded as binding—normatively speaking—is by way of understanding the role of law as an authoritative resolution. The general conditions of an authority's legitimacy provide the framework for connecting a legal ought to a moral ought. The law is morally obligatory if its claim to legitimate authority is morally warranted.[13] At the same time, Raz's thesis proves Kelsen to be wrong as well. Legal ought is not, as Kelsen maintained, just like a moral ought from a different point of view. On the one hand, there is nothing in the structure of morality that connects an "ought" to authority or, in fact, to identity-related reasons. Legal ought, on the other hand, is essentially identity related, because it is authoritative. Furthermore, you may recall that I have complained about Kelsen's account of the normativity of law, that it leaves the choice of adopting any given basic norm to be entirely whimsical, devoid of an explanation of the kinds of reasons people may have for endorsing the basic norm. But now we can see that there are such reasons: the reasons for acknowledging the legitimacy of the relevant legal authority. If, and to the extent that, a legal authority fulfills the conditions of legitimacy, one would have a reason to regard law's

[13] I am not saying if and *only if*. There may be all sorts of moral reasons to comply with the law even if the law fails the conditions of legitimate authority. In other words, an obligation to obey the law may be present even if the law fails in its claim to legitimacy. This point is widely recognized in the literature on political obligation.

instructions or directives as morally binding (*pro tanto*, of course, and not necessarily all things considered).

Where does it leave the question about the possibility of reduction? I think that Raz's thesis about the authoritative nature of law leaves Hart's reductionist framework intact. The requirements of law, whether obligations or other normative prescriptions, are authoritative resolutions. Now, of course, whether a directive has actually been issued by an authority or not is a question of fact—a non-normative fact, that is. Therefore, as long as we can show that (1) law always consists of authoritative directives and (2) the question of who counts as a legal authority and how such authority is to be exercised is determined by social rules, we have laid the foundations for a reduction of legality to facts of a non-normative kind. Further considerations that support (1) will be discussed in the next chapter. Here I want to discuss the considerations that support (2).

The Conventional Foundations of Law

Let us assume that legal norms consist of authoritative instructions or directives. What we need, therefore, is an account of who counts as a *legal* authority in any given legal system. Hart's idea about the rules of recognition would seem to provide a very plausible answer: In each and every society that has a functioning legal system, there are certain social rules followed by the relevant population that determine who counts as the legal authority and how such authority is structured.

However, it turns out that a satisfactory account of the nature of these social rules of recognition, and the ways in which they might constitute the idea of legality, proves to be rather elusive. Some commentators have noted that there is really nothing in Hart's practice theory of the rules of recognition that would explain why people, mostly judges and other officials, are bound to follow those rules. What makes it the case that judges are obliged to follow the rules of recognition? Pointing to the fact that judges take themselves to be bound by those rules does not quite answer the question. What makes it rational for them to do so?

Some years after *The Concept of Law* was published, the philosopher David Lewis came up with a very sophisticated account of social conventions.[14] Lewis was mostly interested in the nature of language, but he offered an ingenious general theory about conventional norms. The basic idea is that conventions are normative solutions to recurrent large-scale coordination problems. A coordination problem arises when several agents have a particular structure of preferences with respect to their mutual modes of conduct—that between several alternatives of conduct open to them in a given set of circumstances, each and every agent has a stronger preference to act in concert with the other agents than his own preference for acting upon any one of the particular alternatives. Most coordination problems in our lives are easily solved by simple agreements between the agents to act upon one, more or less arbitrarily chosen, alternative, thus securing concerted action among them. However, when a particular coordination problem is recurrent in a given set of circumstances—and agreement is difficult to obtain (mostly because of the large number of agents involved)—a social rule is very likely to emerge, and this rule is a convention. Conventions emerge as solutions to large-scale recurrent coordination problems—not as a result of an agreement, but as an alternative to such an agreement, precisely in those cases where agreements are difficult or impossible to obtain.

When this novel account of conventions came to be known, some legal philosophers realized that it may provide the explanation of the nature of the rules of recognition. If the rules of recognition are social conventions, we would have in Lewis's theory both an account of how such rules emerge (like any other convention) and the rationale of following them (to solve large-scale recurrent coordination problems). Thus a conventionalist account of the rules of recognition has emerged, and one that Hart himself, years later, seemed to endorse in his postscript to *The Concept of Law*. As he put it, the rule of recognition "is in effect a form of judicial customary rule existing only if it is accepted and practiced in the law-identifying and law-applying operations of the courts."

[14] See Lewis, *Convention: A Philosophical Study*.

And he added: "certainly the rule of recognition is treated in my book as resting on a conventional form of judicial consensus."[15]

Many contemporary philosophers of law, however, think that this conventionalist turn about the nature of the rules of recognition was a turn for the worse. Ronald Dworkin, for one, argues that there are no rules of recognition at all. Others, more sympathetic to Hart's legal positivist conception of law, argue that a conventionalist understanding of the rules of recognition is fraught with difficulties, and that such a view generates more problems than it solves. To the contrary, I will argue here that with some important modifications the conventionalist account of the rules of recognition is sound.

Before I try to explain the modifications we need, let me say a few words in response to the more fundamental objection to Hart's account, raised by Dworkin. Dworkin denies that the criteria employed by judges and other officials in determining what counts as law are determined by rules, and thus he denies that there are any rules of recognition at all. But as far as I can see, Dworkin's argument is based on a single observation, which is rather implausible. He argues that it cannot be the case that, in identifying the law, judges follow rules, because judges often disagree about the criteria of legality in their legal systems, so much so that it makes no sense to suggest that there are any rules of recognition at all, or else, the rules become so abstract that it becomes pointless to insist that they are rules.[16]

The problem is this: To show that there are no rules of recognition, Dworkin would have had to show that the disagreements judges have about the criteria of legality in their jurisdictions are not just in the margins, that they go all the way down to the core. But this is just not plausible. Could we have anyone in a judicial role in the United States, for example, who seriously doubts that

[15] Hart, postscript to *The Concept of Law* (henceforth Postscript), 256–66.

[16] Dworkin, *Law's Empire*, chap. 1. The same idea is reiterated in Dworkin's recent book, *Justice in Robes*, 164, 190–96. This should not be confused with a different and much more interesting claim that Dworkin also makes: Even if there are rules of recognition, they do not settle the question of legal validity. Norms can be legally valid, Dworkin argues, even if they do not derive their validity from the rules of recognition. This argument will be discussed in the next chapter.

acts of Congress make law? Or that the U.S. Constitution prevails over other forms of legislation? More importantly, as mentioned several times before,[17] there is an inherent limit to how much disagreement about criteria of legality it makes sense to attribute to judges, because the judges' role as institutional players is constituted by those same rules that they allegedly disagree about. The role and authority of certain persons qua judges are constituted by the rules of recognition. Before judges can come to disagree about any legal issue, they must first be able to see themselves as *institutional* players, playing, as it were, a fairly structured role in an elaborate practice. Judges can only see themselves as such on the basis of the rules and conventions that establish their role and authority as judges. In short, pointing to the fact that judges often have certain disagreements about the content of the rules of recognition simply cannot prove that there are no such rules. On the contrary, we can only make sense of such disagreements on the basis of the assumption that there are rules of recognition that constitute, inter alia, the courts system and the legal authority of judges.

So let us make the plausible assumption that there are some rules, mostly followed by judges and other legal officials, determining who counts as a legal authority in the relevant legal system. Are these rules conventions? If we think that the only rationale of social conventions consists in normative solutions to large-scale coordination problems, as Lewis suggested, then the answer is probably no. But let me suggest a more general characterization of conventions that does not tie the function or rationale of conventions to the solution of coordination problems.

Two main features are intuitively associated with conventional rules. First, conventional rules are, in a specific sense, *arbitrary*. Roughly, if a rule is a convention, we should be able to point to an alternative rule that we could have followed instead, achieving basically the same purpose. Second, conventional rules normally lose their point if they are not *actually followed* in the relevant community. The reasons for following a rule that is conventional depend on the fact that others (in the relevant population) follow

[17] Hart, *The Concept of Law*, 133.

it, too. To give one familiar example, consider the almost universal convention of saying the word "hello" when responding to a telephone call. Both features are clearly manifest in this example. Presumably, there is some purpose or point in having a recognizable expression that would indicate to the caller that one has picked up the phone. But, of course, using the particular expression "hello" is quite arbitrary; any other, similar expression would have served us just as well—as long as the expression we use is one that others use, too. If the point of the convention is to have an expression that can be easily and quickly recognized, then people would have a reason to follow the norm—use the expression that others in the community follow as well. And if, for some reason, most people no longer use this expression (as seems to be the case these days), one would no longer have any particular reason to use it either.

Both of these intuitive features of conventional norms can be captured by the following definition:

> A rule, *R*, is conventional if and only if all the following conditions obtain:
>
> (1) There is a group of people, a population, *P*, that normally follows *R* in circumstances *C*.
> (2) There is a reason, or a combination of reasons—call it *A*—for members of *P* to follow *R* in circumstances *C*.
> (3) There is at least one other potential rule, *S*, that if members of *P* had actually followed in circumstances *C*, then *A* would have been a sufficient reason for members of *P* to follow *S* instead of *R* in circumstances *C*, and at least partly because *S* is the rule generally followed instead of *R*.
>
> The rules *R* and *S* are such that it is impossible (or pointless) to comply with both of them concomitantly in circumstances *C*.[18]

As we just saw, Dworkin's objection to the rules of recognition denies the truth of premise (1). But we also saw that this objection fails, so let us assume that (1) is true. Given the truth of (1), it

[18] I suggested this definition and elaborated on it in much greater detail in my *Social Conventions*, chap. 1.

would be extremely unlikely that (2) is false. If judges and other officials follow certain rules that determine what law is, surely they follow them for reasons. What those reasons, generally speaking are, however, turns out to be somewhat difficult to answer. In his original account of the rules of recognition, Hart suggested that the rationale of the rules of recognition consists in the need for certainty. In a developed legal system, Hart argued, people would need to be able *to identify* what types of norms are legally valid. In fact, he presented this advantage of the rules of recognition in providing certainty about the valid sources of law as the main distinguishing factor between "primitive," prelegal normative systems and a developed legal order.[19] Later, in his postscript to *The Concept of Law*, Hart seems to have added another kind of reason for having rules of recognition, basically of a coordinative nature:

> Certainly the rule of recognition is treated in my book as resting on a conventional form of judicial custom. That it does so rest seems quite clear at least in English and American law for surely an English judge's reason for treating Parliament's legislation (or an American judges' reason for treating the Constitution) as a source of law having supremacy over other sources includes the fact that his judicial colleagues concur in this as their predecessors have done.[20]

I have some doubts about both of these explanations. That the rules of recognition contribute to our certainty about what counts as law in our society is surely true. But is it the main reason for having such rules? This I doubt. It is like suggesting that there are some rules or conventions about what constitutes a theater performance so as to enable us to identify this form of art as distinguished from other, similar artistic endeavors. Surely, if there are some conventions that constitute what theater is, it is because there are some artistic reasons for having this form of art in the first place. Similarly, I would suggest, if there are reasons to have rules of recognition, those reasons must be very intimately linked to the reasons for having law and the main functions of law in

[19] Hart, *The Concept of Law*, chap. 5.

[20] Postscript, 267.

society. Certainty about what law is cannot be the main reason for having law. There must be some reasons for having law first, and then it might also be important to have a certain level of certainty about it. It cannot be the other way around. I am not suggesting that the reasons for having rules of recognition are the same as the reasons for having law in a society. My claim is that the reasons for having rules of recognition are closely tied to those reasons and, in some ways (yet to be specified), they instantiate them.

The coordinative rationale of the rules of recognition is even more suspect, and for reasons that are quite explicit in Hart's writings. It is true that judges and other legal agents, acting in their official capacities, need a great deal of coordination in various respects. In particular, they would need to follow basically those same rules that other officials in their legal system follow in identifying the relevant sources of law in their legal system. That the rules of recognition enable this kind of coordination in the various actions of legal officials is not disputable. But again, it makes little sense to suggest that this is the main rationale of the rules of recognition. As mentioned above, for judges to have any coordination problem that might need a solution, first we must be able to identify them *as judges*; we first need a set of rules that constitutes their specific institutional roles. In short, and more generally, first we need the institutions of law; then we may also have some coordination problems that may require a normative solution. The basic role of the rules of recognition is to constitute the relevant institutions. The fundamental rules of recognition of a legal system are constitutive rules (or conventions, as we shall see), and their coordination functions are secondary at best.

There is a rather striking confusion in some of the literature on the conventionality of the rules of recognition that connects these two points. Because the standard understanding of conventions has been the one offered by Lewis, which consists of the idea that conventions are normative solutions to coordination problems, commentators have been drawn to the idea that, if the rules of recognition are conventions, their basic rationale must be a coordinative one. But commentators have also realized that the rationale of the rules of recognition must be closely tied to the reasons for having law in the first place. And the combination of

these two points has led many to assume that the main rationale *of law itself*—the main reasons for having law in society—is also coordinative in nature.[21] This has rendered legal conventionalism, as this view came to be called, rather implausible. The idea that law's main functions in society can be reduced to the solution of coordination problems is all too easy to refute. Solving coordination problems, as complex and intricate as they may be, is only one of the main functions of law in society, and probably not the most important one.[22]

I mention this confusion partly because Leslie Green's critique of legal conventionalism, often cited as a main argument against a conventionalist construal of the rules of recognition, is based on it. Green is absolutely right to claim that the authority of law, and its main moral-political rationale, cannot be explained in terms of law's function in solving coordination problems.[23] But he is wrong to conclude that this undermines a conventionalist account of the rules of recognition. Neither the main functions of law in society nor the main rationale of the rules of recognition have much to do with solving coordination problems.

We have yet to show that the rules of recognition are conventions. The conventionality of the rules of recognition crucially depends on the third condition—on the question of whether the rules are arbitrary and compliance dependent in the requisite sense. So let us turn to examine this aspect of the rules of recognition. On the face of it, the arbitrariness of the rules of recognition is strongly supported by the following two observations. First, we know that different legal systems, even ones that are very similar in many other respects, have different rules of recognition. The rules followed in the United States in recognizing the sources of American law are very different from those followed, for instance, in the United Kingdom about the recognition of British

[21] See, e.g., Lagerspetz, *The Opposite Mirrors*, and Hartogh, *Mutual Expectations: A Conventionalist Theory of Law.* Dworkin's interpretation of what he calls *legal conventionalism* relies on a very similar idea. See his *Law's Empire*, chap. 7.

[22] Notice that coordination problems are much easier to solve than the other types of collective-action problems because there is no serious conflict of interest between the parties involved.

[23] See Green's "Positivism and Conventionalism," 43–49.

law. Second, there is a very clear sense in which the reasons for following the rules of recognition are compliance dependent in the relevant sense. This is one of the points that Hart has rightly emphasized in his postscript to *The Concept of Law*—that the reasons judges and other officials have for following certain norms about the identification of the sources of law in their legal systems are closely tied to the fact that other officials follow those same norms.

I do not think that either of these observations supporting the conventionality of the rules of recognition is really controversial. The reasons critics have for doubting the conventionality of the rules of recognition pertain to the normative aspect of the rules. Again, Green was one of those who observed this difficulty in the conventional account of the rule of recognition. As he put it, "Hart's view that the fundamental rules [of recognition] are 'mere conventions' continues to sit uneasily with any notion of obligation," and thus, with the intuition that the rules of recognition point to the sources of law that "judges are legally bound to apply."[24] So the problem seems to be this: If the rules of recognition are *arbitrary* in the requisite sense, how can we explain the fact that they are supposed to obligate judges and other legal officials to follow them?

I think that by now we have all the tools we need to answer this question. First, even if Green had been right to assume that the main conventionalist rationale of the rules of recognition is a coordinative one, the puzzle he raises about their potential normativity is easily answered. Some coordination problems are such that there is an obligation to solve them. If a conventional solution has emerged, the relevant agents may well have an obligation to follow the conventional solution. However, since I do not think that the rules of recognition are coordination conventions, I will not avail myself to this simple answer. The main answer to Green's puzzle resides in the distinction between the idea of a legal obligation to follow the rules of recognition and the separate moral or political question of whether judges (or anyone else for that

[24] Green, "The Concept of Law Revisited," 1697.

matter) have reasons to engage in the practice that is constituted by those rules.

The rules of recognition, like the rules of chess, determine what the practice is. They constitute the rules of the game, so to speak. Like other constitutive rules, they have a dual function: They determine what constitutes the practice and prescribe modes of conduct within it. The *legal* obligation to follow the rules of recognition is just like the chess player's obligation to, say, move the bishop—if one is to move it—only diagonally. Both are prescribed by the rules of the game. What such rules cannot prescribe, however, is an "ought" about playing the game to begin with. Conventional practices create reasons for action only if the relevant agent has a reason to participate in the practice to begin with. And that is true of the law as well. If there is an "ought" to play the game, so to speak, then this "ought" cannot be expected to come from the rules of recognition. The obligation to play by the rules—to follow the law, if there is one—must come from moral and political considerations. The reasons for obeying the law cannot be derived from the norms that determine what the law is.

Thus, my main response to Green's worries about the normativity of the rules of recognition is this: Once we realize that the rules of recognition are constitutive and not coordinative conventions, we can see that there is really nothing unique or particularly puzzling about the idea that judges ought to follow the rules. The sense in which a judge is obliged to follow the rules of recognition is exactly like the obligation of an umpire in a cricket game to follow the rules of cricket. Both obligations are conditional. If, and to the extent that the judge or the umpire have reasons to play the game, the rules simply determine what their obligations in the game are; they constitute what the game is. In both cases, however, we cannot expect the rules of the game to constitute the reason to play it. In other words, the internal (legal) obligation is determined by the rules themselves; the rules that constitute the game also prescribe modes of conduct within it. The external obligation (or, generally, reasons) to play the game, if there is one, is a different matter—one that cannot be expected to be determined on the basis of the normativity of the rules of the game. Whether

judges, or anybody else, would have an obligation to play the game, as it were, is always a separate question—one that needs to be determined on moral-political grounds.

Let me summarize some conclusions. Hart's reductionist project relies on the idea that the conditions of legal validity are determined by the social rules of recognition. I argued here that a plausible conventionalist construal of this thesis is available. However, I also argued that Hart's reductionism about the normativity of law is overly simplistic and needs to be modified by Raz's important insight that law always claims to be a legitimate authority. The authoritative nature of law, I suggested, supports the construal of legal norms as instructions or directives issued by legal authorities. These two points, taken together, would entail the following two theses:

(1) In every society that has a functioning legal system, there are some social conventions that determine who counts as legal authority in that society and how its authority is to be exercised.
(2) Legal norms consist of the directives or instructions of legal authorities—those authorities that are identified and constituted by the social conventions of (1).

This, I believe, is a somewhat modified version of Hart's version of legal positivism. In the next two chapters we will consider some important challenges that have been leveled at these ideas, and we will try to assess their force. The next chapter is devoted to a detailed defense of the second thesis.

Suggested Further Readings

Green, *The Authority of the State.*
Marmor, *Social Conventions: From Language to Law.*
Raz, *Between Authority and Interpretation.*
———, *The Morality of Freedom*, chaps. 1–4, 7.
Shapiro, "Authority," 382.