

Dear NYU workshop participants,

The paper that I present at the workshop is the last chapter of a book manuscript which is still very much work in progress. I picked this chapter because it is the one most closely related to law and some of the familiar debates in legal philosophy. However, since it is the last chapter, I thought that it might be helpful, though not strictly necessary, to read first a brief description of the book project. (pp 1-3).

## **Foundations of Institutional Reality**

*Andrei Marmor*

### **Book description**

Human life without social practices and institutions would be life that our ancestors may have had tens of thousands of years ago. The world we live in is a world that we, collectively somehow, created ourselves. There is hardly any aspect of our daily lives which is not enabled, structured, or regulated by social practices and institutions. It is this collective social reality or, at least, some of the basic aspects of it, that I aim to explore in this book from metaphysical perspective. The book aims to explicate the foundations of our institutional reality, ways in which we can build it up, as it were, from relatively more simple building blocks. The construction metaphor, shared by many philosophers who write on these issues, is not accidental. It suggests that the philosophical project envisioned here is one of grounding or reduction. The idea is that the central aspects of institutional reality can be accounted for by way of reducing them to more foundational elements, like patterns of conduct, attitudes and dispositions, functions, motivating reasons, products of human agency, and similar elements. Therefore, part of the project here is to identify what is more foundational than something else, how to unpack this building or construction metaphor, and what is the explanatory payoff we get from it. One explanatory gain that I hope to show in this work concerns a number of epistemic considerations that are entailed by the metaphysics of our social

reality, showing how they have considerable implications for what would be an acceptable philosophical account of social practices and institutions.

The philosophical foundations laid out in H.L.A. Hart's *The Concept of Law* (1961), and John Searle's *The Construction of Social Reality* (1995) inspire the project of this book. Hart's work on the foundational role of social rules, and Searle's more ambitious project of constructing the basic elements of institutional facts from collective attitudes, functions and constitutive rules, I take to be very important philosophical insights in the right direction. My main purpose here is to offer some modifications and make a few more steps, thereby offering a more complete picture of the metaphysical foundations of our institutional world. Recent work in metaphysics about grounding, collective intentionality, and the nature of artifacts, helps me in this endeavor. I hope to show that we can use some of the recent developments in metaphysics to shed new light on these old ideas and, hopefully, develop them further.

The book begins with a clarification of some of the main concepts in play, distinguishing social facts from institutional facts, and showing why Hart was right to assume that rules are at the foundations of institutional facts. And yet I also show why Hart's own theory of social rules is problematic and crucially incomplete. Thus, one challenge is to explain the idea of a rule and what it is for a social rule to exist; this, I argue, requires an account of the grounding of collective attitudes. In chapter two I explore the idea of metaphysical grounding, explain why the hierarchical metaphysics it presupposes is particularly apt for theorizing about the social domain, and relate it to the idea of reduction. The main conclusion of this chapter is that there are two distinct types of reductive explanations: reduction as (complete) metaphysical grounding, and reduction by way of a real definition. The former is the one that is mostly relevant for social ontology. The argument of chapter 3 is that an account of social rules requires a two step reduction; the content of social rules is fully determined by the collective attitudes of the relevant population, and the existence of rules is grounded in the relevant collective attitudes coupled with certain patterns of conduct. Thus, providing a metaphysical account of collective attitudes, and one which comports with methodological individualism, forms the aim of this chapter.

Chapter 4 takes a closer look at what it means for rules to constitute practices and institutions, arguing for a functional explanation of this grounding relation, but also showing that the relation between rules and practices is only partial grounding; some minimal level of agreement in judgments about the practice's point or purpose is also practically necessary. (This chapter also calls into question Searle's famous distinction between constitutive and regulative rules, and explains why the relevant notion of function is grounded in practical use.) In chapter 5 I turn to an elaboration of the nature of artifacts, striving to show that social rules and their emergent practices are artifacts. And this is important because there is a certain epistemic privilege that is entailed by the artifact nature of objects; since artifacts are essentially products of human agency, it is not possible for an entire population to be comprehensively mistaken about what the artifact they created is. The epistemic privilege here pertains only to the content and use of the artifact in question; comprehensive errors about the ontology of artifacts are quite possible, and in some cases prevalent. The epistemic privilege about content and use of artifacts is aligned with the idea, shown in chapter 3, that the content of social rules *just is* the collective intention of the relevant population, also implying that comprehensive errors about it are not quite possible.

Chapter 6 turns to the role of a particular type of rules in the foundations of institutional practices, namely, power conferring norms. The chapter offers an analysis what normative powers are, how they work, and the foundational role of basic power structuring rules in the grounding of our institutional reality. Finally, in chapter 7, I turn to some of the methodological implications of the reductive account of social practices offered thus far, arguing that a metaphysical grounding account of a social practice is not constrained by the need to explain it in terms that would make the practice intelligible to the participants, explaining why they engage in it and what is the point of it, for them. This, I argue, is one of the main differences between metaphysical and hermeneutical accounts of sociality, and I demonstrate this difference, and its philosophical implications, by showing how it plays out in the famous debate between Hart and Dworkin about the nature of law.

## CHAPTER 7

### RATIONALIZING PRACTICES AND THE HERMENEUTIC CHALLENGE

The content of social rules, as we saw in chapter 3, is fully grounded in the collective attitudes of the population whose rule it is. Which means that a population that follows a social rule has a certain epistemic privilege with respect to the content of the rules they follow: comprehensive errors about it are not possible. Similarly, we saw in chapter 5 that the artifact nature of our social practices excludes the possibility of comprehensive errors about their contents and practical functions. Finally, as we saw in chapter 4, for a social practice or institution to exist there has to be a minimal level of agreement in judgments amongst its participants about the reasons for having the practice and the values it instantiates for them. These three independent conclusions sit comfortably together. The concern I want to address in this chapter, however, is that the direction of these conclusions might lead us to another conclusion, but one that I would like to resist. The thought might be that any philosophical account of a social practice would need to offer the kind of explanation that would rationalize the practice for its participants, make it intelligible to them why they engage in it, and what is the point of it, for them. I call this the internal rationalization constraint and will explain it shortly. My main argument in this chapter is that a reductive metaphysical explanation of a social practice is not subject to the internal rationalization constraint, and that this is one of the main differences between a metaphysical theory of social practices and an hermeneutical one.

Theories about the nature of law will be my main example here; I will try to show that H.L.A. Hart's conception of what to expect from a theory about the nature of law is metaphysical and reductive; the interpretative critique leveled against it, predominantly by R.M. Dworkin, presupposes a very different methodology, and one that is clearly guided by the internal rationalization constraint. My point is going to be that metaphysical theories about the nature of law, or any other social practice for that matter, do not need to rise to the challenge; a grounding-reductive account of social practices does not need to meet the internal rationalization constraint. Hermeneutic theories of

practices do take the rationalization of the practice for its participants as an important objective, but then they must also rely on a different, essentially normative, methodological framework.

### 1. The Internal Rationalization Constraint.

Social practices are structured by rules; rules have constitutive functions in making the practice what it is, functions that are grounded in purposeful use, which itself is grounded in patterns of conduct and collective, learnable, attitudes and dispositions. As we noted in chapter 4, social practices require, if for no other than practical reasons, a certain level of agreement amongst its participants about its main functions and values, at least for those who willingly engage in the practice.<sup>1</sup> One pervasive feature of social practices (as well as of most massive artifacts) is their multiple instantiations. Practices and massive artifacts can differ from each other as different sub-types of the same general type. There are quite a few different schools of Yoga, but despite significant differences, Yogis will tell you that they are just different practices of Yoga. Cornell university is quite different from the University of Oxford, in its teaching practices, governance structure, admission procedures, and countless other aspects, yet they are both research universities. Practices and massive artifacts allow for different versions of the same general type and they are fully grasped as such by their respective participants. This is a very pervasive aspect of social practices, they often exist in different, sometimes significantly different, versions or instantiations. But they are seen by the relevant agents as different instances of the same general practice, the same kind of social institution.

It would seem to be impossible to explain this variability if we tried to account for social practices and institutions without taking into account their overall purposes, social functions, and generally, some explanation of what the practice is there for. And this, of course, is closely tied to the fact (mentioned in ch 4) that some of the rules or sub-parts structuring a social practice are much more essential or central to it than others. What

---

<sup>1</sup> We should not lose sight of the fact that there are many social institutions which are imposed by force and maintained by violence or at least violent threats. In most cases, I think, this only means that the relevant population whose practice we need to focus on is the population that imposes the institution by force and manages to maintain it as such. The relevant participants in our institutions of incarceration, for instance, are not the inmates, but those who put them there and manage to keep them incarcerated.

makes certain rules or parts of a practice more central to it than others must have something to do with the general purposes or aims of the practice, it's main point, as it were.

When you think about a university, for example, you would think of the pursuit of knowledge, the value of research and teaching, the freedom of thought and innovation, and things like that, as essential aspects of what universities are for and what the point of being part of it is for you. Giving us a list of the rules and conventions that constitute an institution like a university, with their corresponding functions, is just not going to give us the full story of what universities are, what motivates people to take part in the practice or, as often the case, seek to change aspects of the practice to better serve its overall purposes. But precisely because we understand these things, we can easily see how different versions of a practice or institution are just different versions of essentially the same thing. In other words, the idea here is that combining the functions of rules structuring a social practice is not going to generate the overall functions and purposes of the institution or practice as a whole. We need to think about features of social practices that rationalize their existence, aspects of the practice that would explain what the practice is all about, and how it is tied to things people value.

Thus, social practices and institutions must have rationalizing features, which would be aspects of the practice that explain why it exists, and why it is valued by those who practice it and maintain its existence as the kind of practice it is. But here we need to draw a crucially important distinction: rationalizing features of a social practice can be external or internal. A rationalization is *internal* if it is the kind of explanation that the participants in the practice, those whose practice it is, can see as their own, so to speak, explaining the point of the practice, the reasons for engaging in it, *for them*. Thus, an aspect or a feature X of a practice, P, would be internally rationalizing for participants of P if they can come to see X as a reason to engage in P, or something that ties the practice to things *they* value or care about.

A rationalizing feature of a practice would be *external*, therefore, if it explains why people have the practice in the way they do, or why they engage in it, without necessarily offering an explanation that participants can see as a reason to engage in the practice or things they value about it. Thus, Y would be an externally rationalizing

feature of a social practice P, if Y contributes to an explanation of P as the kind of practice it is, even if Y is not something that participants in P would be willing to see as a reason to engage in P. Now this, of course, does not mean that externally rationalizing features of a practice cannot be internal as well; it only means that they do not have to be. Let me emphasize, however, that the distinction between internal and external rationalizing features is not the distinction between motivating and normative reasons. First, the internally rationalizing features of a practice need not be on the surface; they include factors or features of a practice that participants may not be aware of, but would be willing to recognize as their own once presented to them. More importantly, external rationalizing features do not have to be reasons in the normative sense; many social practices and institutions are such that there are no reasons to have them, or have them in the ways they are actually practiced.<sup>2</sup> That does not mean that they have no rationalizing features. There are no reasons to have racist or sexist practices and institutions, but their racism, sexism or similar features are such that they can contribute to an explanation of why people have the practice and what functions and values it serves for them.<sup>3</sup> And, crucially, such features can be rationalizing even if the participants in the practice would not tend to recognize those features as normatively significant for them.

Let me give an example or two. Consider the practice of the so called “traditional” family. It is a social practice, partly but significantly structured by social and legal rules – with some considerable local and temporal variations – collective attitudes and familiar, widely shared, patterns of behavior. Clearly it is the kind of social institution that has a great deal of meaning and personal significance for those who practice it, for many people it is the most significant institution in their lives. There are, no doubt, rationalizing features of family practices that are internal, they would be the

---

<sup>2</sup> If you are inclined to think that there are bad reasons, normatively speaking, then you can understand my point as saying that externally rationalizing features do not have to be good reasons.

<sup>3</sup> The distinction between internal and external rationalizing features has nothing to do with Hart’s famous distinction between internal and external points of view about social rules. The internal point of view is already incorporated in the Conditions-H of what social rules are, as explained in chapters 1 & 3. As I have also explained in some detail in (Marmor, 2011: 53-55), Hart’s distinction between the two (or rather, three) points of view is meant to reinforce his reductionist explanation of law, not to call it in doubt. Hart’s main point here is that we can account for the internal perspective reductively, by accounting for people’s shared attitudes, without having to endorse those attitudes or presuppose their soundness or anything else about them. The distinction is a critique of Kelsen, not mainly Austin, as most commentators hastily concluded.

kind of features that participants would recognize as their own, something that rationalizes for them why they value the institution and their reasons for maintaining it. Needless to say, there is plenty of room for disagreement here; some people value the idea of a family partly for religious reasons, other care nothing about those. But the fact is that that we can easily discern the internally rationalizing features of the institution of the family despite possible disagreements about them and different interpretations of them. There are also some rationalizing features of the institution of the family that are clearly external, such as a causal-evolutionary explanation of the family. Perhaps these practices got reproduced over millennia because they turned out to be biologically successful (in whatever evolutionary terms we want to explicate this). But there are some possibilities in the middle: neither internally rationalizing nor merely causal.

One example I have in mind, quite obviously, is the feminist critique of the family as an institution that serves to entrench patriarchal dominance of women and children by men.<sup>4</sup> Allow me to assume here that there is a great deal of truth in this feminist critique. And yet it highlights a rationalizing feature of the institution that most participants would tend to find alien; it is not something that rationalizes the practice for them. Most people, men and women, wouldn't think that male domination is a reason for their adherence to traditional family structures, something that counts in favor of getting married and raising their biological children together. In this respect, it is very similar to the Marxist critique of religion, as a form of false consciousness. It is just not the kind of perspective on religion, an explanation of its rationale, as it were, that any religious person can coherently endorse. Perhaps the feminist perspective on the family is not quite as alienating as a Marxist account of religion. If you are sincerely religious, you just cannot be a Marxist about it; but you can have a family even if you recognize the truths and insights of the feminist critique. It's not easy, and tensions are bound to come to the surface here and there, but (I hope) possible. And of course, it's also true that the more publically recognized and appreciated the feminist critiques of the traditional family becomes, social pressure builds up for change, as we have witnessed significant changes

---

<sup>4</sup> The literature on the essential patriarchal nature of the family is vast; see, for example, MacKinnon (1987), ch 3.

unfolding in the last few decades. External critiques can become internalized over time and push for social and legal changes.

But what are these external rationalizing features, metaphysically speaking, that is? When we say that the traditional family structure serves to entrench male domination and patriarchal rule, we seem to imply that there is some functional explanation in play here. The question is, however, what kind of functional explanation? As we saw in chapter 4, not every functional explanation necessarily attributes practical function, function as use, to the explanandum; functional explanations are often about causal chains. Therefore, one option here as well, is that such factors are simply elements of a causal explanation. To say that traditional family structures serve to entrench male domination may be just a way of saying that it is what the institution does, it's what the effects of it are, socially speaking. Perhaps the causal chain starts earlier; perhaps the desire and ability of men to dominate women and children is what brought about the family practice, pushed for it, as it were, in a complex causal chain over centuries and millennia. Surely there is some truth in this too; social practices and institutions come into being by a complex causal chain and they have a myriad of causal effects, whether intended or not. But the problem with this causal account is that it leaves out some of the explanatory power of what a rationalizing feature is. There are countless kinds of items or elements that figure in a casual explanation of a social practice, and many causal effects of it, but only some of them would seem to bear on the practice's intelligibility, its purpose, point and people's (motivating) reasons for engaging in it.

A second option is to think about externally rationalizing features in terms of practical functions, function as purposeful use, even if some of the purposes are not on the surface, as it were. In other words, invoking the idea of hidden functions, that we mentioned in chapter 4, seems particularly tempting here. It is very tempting to suggest that something like entrenchment of male domination is a hidden practical function of the traditional family institution. In a way, we can say that it is what the institution is used for, even if this use is not explicitly endorsed by the participants or recognized by them. This seems quite plausible to me, but there are some difficulties to overcome. One difficulty comes from the fact that when the hidden function is exposed, as it were, when you present it to the relevant participants, most of them are likely to find it alienating;

they would not exhibit a kind of “Aha” moment, “so yes, now we see what reasons we have for doing this!”. If we want to retain the idea of function as purposeful practical use, the relevant question here is whether people can actually use an X for F while sincerely rejecting the idea that X serves F for them? If the answer to this question is yes, then there should be no objection to seeing externally rationalizing features of a social practice in terms of practical function, even if, as sometimes the case, those functions are “hidden”. So let us explore the possibility. Before we proceed, however, let me emphasize that I am not claiming here that all externally rationalizing features consist of hidden functions. Far from it. Externally rationalizing features would include all those putative values, purposes and functions that would contribute to an explanation of what a practice or an institution is for, what explains its existence, and ways in which it is practiced. What I am trying to show here is that externally rationalizing features can accommodate hidden functions as well, those which would fail the internal rationalization criterion.

So now the question is this: Is it possible for *A to use X for F*, even if A would sincerely reject the idea that she is using X for F? Not to bring things too close to home, we might recall an example that is familiar to all of us, I’m sure; think of the times when you participated in a workshop and someone asked a question or made a comment mostly in order to impress their colleagues how clever they are. I presume that some would be willing to admit, at least in retrospect, that the desire to impress was the true purpose of their question; but some might resist the explanation, even quite sincerely so. Which comes to show that there is a sense in which the purpose of one’s conduct is occasionally not fully avowable by an agent. Those who research and write about racism would surely agree. Many people act in ways that reflect their racist biases and prejudices while sincerely rejecting any suggestion of a racist element in their motives and purposes. So it seems that the idea of a hidden function is not so suspect after all, even if we retain the idea of practical function as purposeful use. As I mentioned in chapter 4, the idea of a purposeful action is not confined to those cases in which the aim or purpose of the action is at the forefront of the agent’s mind. Racism, for example, would have been much less

invidious and prevalent had it been the case that people cannot act on racist prejudices unless they are aware of it.<sup>5</sup>

In other words, things like racism, sexism, xenophobia, and similar aspects of social practices can be rationalizing features even if they clearly fail the internal rationalizing constraint. They are features we can employ in an explanation of why a practice exists, what is the point of it, why do people engage in it in the ways they do, how they adapt the practice to changing circumstances, and all sorts of explanatory tasks of this nature, without assuming that the participants themselves would avow those features as rationalizing the practice for them.

Now one might worry that the possibility of including hidden functions in the externally rationalizing features is in tension with the kind of epistemic privileges I argued for in previous chapters. If it is not possible for an entire community of participants in a social practice to misconceive the contents of the rules they follow, or to be mistaken about the main purposes and functions of the artifacts they create, how can those things have practical functions that are not fully avowable by the participants? We need to tread with caution here. There are three types of epistemic privilege that emerge from the discussion in previous chapters.<sup>6</sup> The first type, concerning the prescriptive content of social rules, is the least problematic in the present context. The content of a social rule, I argued in chapter 3, just is the collective intention of the relevant population. Therefore, a population cannot be comprehensively mistaken about the content they ascribe to a rule; the rule's content is what it is in virtue of the attitudes they share about it. But notice that this epistemic privilege only applies to the content of the rule, not its origin, ontological status, or anything else for that matter. As mentioned earlier, people can think that a rule they follow is commanded by God, even if there is no God; people can think that a rule they follow is not a convention, because they are not aware of alternative rules they could have followed in the circumstances, but the rule may nevertheless be conventional. So there are all sorts of comprehensive metaphysical

---

<sup>5</sup> The literature on so called 'implicit biases' is considerable and somewhat controversial. See, for example, Banaji & Greenwald (2013); Payne, Niemi & Doris (2018).

<sup>6</sup> Bear in mind that all these epistemic privileges concern collectives or groups of people, not individuals. Any individual participating in a practice or observing it might be fundamentally mistaken about essential aspects of it. What epistemic privilege rules out, in our use of the term, is only comprehensive errors, namely, the possibility that an entire community of participants or practitioners, is mistaken.

mistakes that are possible about social rules; the epistemic privilege here is only about the rules' prescriptive content, about what the rule is.

Something very similar applies to the epistemic privilege we argued for with respect to artifacts. We have seen that comprehensive mistakes about the ontology of artifacts is not just possible, but in some contexts, quite common. People often mistake artifacts, especially of the intangible kind, for natural objects, that is, things that exist irrespective of human agency. The epistemic privilege we have about artifacts concerns their contents and main functions. We cannot come to discover that chairs are not really objects to sit on, or that Zeus was not, after all, the chief Olympian god. Artifact products are what we make of them. Furthermore, we argued in chapter 5 that practical use determines what kind of artifact it is. But here one might suspect that a serious tension emerges: on the one hand, we want to say that an artifact's nature as a product of human agency, grounded in practical use, grants the relevant population some epistemic privilege with respect to the content or function of the artifact, but on the other hand, we want to allow for functions to be hidden sometimes, not fully avowable by the relevant agents in question. And this might seem contradictory.

In response, two points to bear in mind here. First, even if there is a product of human agency that has functions that are not fully avowable by the users, nothing can have only hidden functions. Hidden functions are parasitic on functions as purposeful use in the ordinary, self-avowable sense. We do not have artifacts as products of human agency without considerable epistemic transparency about the purported use of the product, its designated or widely used function. (Remember that practical use is something that we normally learn from others.) Hidden functions, if and when they are present, are always parasitic on ordinary, epistemically transparent, functions. First we need to have an object generally used for something, and then we can also come to realize that it serves something else as well, perhaps not widely recognized as such. In other words, the idea of hidden functions does not undermine the idea of epistemic privilege, it only limits its scope.

Secondly, this limit in scope is itself very limited to a certain type of comprehensive error. It is not really the case that by discovering the hidden function of an artifact or an institutional practice we come to see that the population in question is

comprehensively mistaken about what it is that they collectively created. They may be mistaken, perhaps gravely so, about the real purpose or function of their creations, but not so much about what it is that they created. In other words, the kind of epistemic privilege we argued for is privilege about content, or identifying features, not about purposes. When a Marxist claims that religion is ideology, the Marxist need not deny that what people take their religion to be is what it is. She only denies that a complete account of the real purposes or functions of religion is what religious people assume. Similarly, by claiming that traditional family structures serve purposes of male domination the feminist critique does not claim that people can be comprehensively mistaken about what a family is. We can acknowledge that people sometimes act for purposes they are not fully aware of, or even for purposes they would not be willing to avow. Think about it in the context of language use. Suppose we come to realize that the meaning of a certain word or expression is such that it evokes certain emotions in the mind of people who use the word without being fully aware of it; in fact, most people might even sincerely reject the idea that their use of the expression evokes those feelings. And suppose we can show that this emotive aspect partly explains why and when people use the word. (If an example is needed, think of the use of the N\* word in times when it was still socially acceptable; many users would have sincerely denied that their use of the N\* word gives them a sense of superiority; surely many white people were actually wrong about that.) None of this would count against the idea that the population who uses the expression knows what it means; they cannot be comprehensively mistaken about the meaning of the word, that is, even if they are not fully aware of all the functions that the use of the word serves for them.

The question that remains open, however, is the one we started with: would we have good reasons to think that a metaphysical account of a social practice must meet the internally rationalizing constraint? In particular, would it be a valid argument against a grounding-reductive explanation of a social practice that it fails the internal rationalization constraint? In the next section I will explore this issue using the example of one of the main controversies about the nature of law in legal philosophy.

## 2. Legal Positivism, Reduction & Hermeneutics.

There are at least three main debates about the nature of law in analytical jurisprudence, answering to different questions: one is a debate about the ontology of law, the other is about determinates of legal content, and the third is about law's normative or moral, reason giving, aspects. Naturally, our focus here is going to be on the ontological debate. Which is not to deny that there are all sorts of connections between these three debates and in certain ways they may be intertwined.<sup>7</sup> Legal positivism, however, is a theory, if theory is not too big a word here, focused on the ontological aspect of law, aiming to answer the question of what makes things legal, part of the legal domain, as it were, when they are. It is a question of the general form *What makes an o an F?* where o stands for some normative or prescriptive content, and F stands for legality, in some relevant sense (yet to be defined). This is a very narrow question, and an answer to it, even if fully successful, would hardly amount to a theory about the nature of law. Nevertheless, this is the question we will focus on here.

So what makes it the case that a norm or a requirement or such is part of the law (somewhere at some point in time)? As I mentioned in chapter 2, the background against which this ontological question arises is very familiar. We often say things like “It is the law (here and now) that ... such and such”; or that “such and such is legally required”; or that “It had been a legal requirement to  $\varphi$  until a year ago but it no longer is”. We normally talk about laws as things that exist, in some sense; propositions about what the law is in a given jurisdiction can be true or false. To be sure, we do not need to assume complete determinacy here; sometimes there is genuine indeterminacy about whether something is or is not part of the law.

We probably owe this focus on the ontological question about the nature of law to the influential work of John Austin (1832). In any case, I will use Austin's understanding of the question, and the research project he envisaged about it, as my starting point. There are two main ambitions Austin had envisaged: First, he sought to provide a reductive explanation as the answer to the ontological question. Austin argued that the answer to

---

<sup>7</sup> I explored some of the ways in which these debates are intertwined in (Marmor, 2019); for my take on the normativity of law, see (Marmor, 2018) where I also explain why different views on the normativity of law do not necessarily derive from positivism or anti-positivism.

the question of what makes norms legal, when they are, can be fully articulated in terms of social facts about political sovereignty,<sup>8</sup> people's conduct and attitudes, and particular action and events that take place in the world. Second, Austin expected the answer to be such that it is capable of explaining how law differs from similar normative domains, such as morality, religion, social customs, etc.,. Thus, Austin aimed to meet two criteria for the adequacy of an answer to the ontological question: reducibility and uniqueness. An account of legality has to be reductive, and it has to show how law is different from other, similar normative domains. However, the argument about the vagueness of basic power structuring rules towards the end of the previous chapter should make us somewhat doubtful that the uniqueness desideratum needs to be met, at least as strictly as Austin had assumed.<sup>9</sup> So let us focus on the reductive ambition.

Thus, following Austin, let us assume that the question we ask is: *what makes an object, o, an F?* What is *o*, the object, in the legal case? We tend to talk about norms; norms are the kind of things, it is generally assumed, that can be legally valid or not. And that is largely true, but not entirely accurate. Particular instructions, such as "You [a person] do this...[a particular act or act type]", can be legal instructions. Laws do not have to be norms, though they usually are. And then perhaps somewhat looser prescriptive contents, such as general doctrines, even ideologies<sup>10</sup>, may also form part of the law, even if they are difficult to formulate in canonical normative terms. Bearing these reservations in mind, at least for the sake of convenience, we can assume that the question is: what makes norms legal norms? Now the *F* in question here is a bit more difficult to formulate precisely. For one thing, *F* is always relative to a particular jurisdiction at a particular time. Norms are legally valid in some legal order or other, at a given time, not in the abstract. Thus the question of what makes an *o* an *F*?, fully articulated, would need to be spelled out in a way that is relative to particular jurisdictions and time.<sup>11</sup> But of course, the philosophical question is not about the specific determinates of legal validity in particular jurisdictions; what we want to know is what

---

<sup>8</sup> I use the expression 'social facts' advisedly; one of Hart's convincing critiques of Austin is precisely the idea that Austin failed to see the distinction between social facts and institutional facts. I explain this in some detail in (Marmor, 2011: 44-59.)

<sup>9</sup> As I mentioned in ch 6, Hart shared Austin's view in this respect.

<sup>10</sup> For example, I think that it would be difficult to deny that a commitment to Capitalism forms part of US law.

<sup>11</sup> International law included, I am not doubting its legal status here.

makes it the case, generally speaking, that any given norm would form part of the law where and when it does. Which is to say that we can formulate the question conditionally: If a given norm is part of the law in Si at t, what has to be true to ground this fact?

So what are the grounding facts of legality? I will not dwell on the details of Austin's answer; what matters is the philosophical ambition, which is to provide the answer reductively, in terms of people's actual modes of conduct, the beliefs and other attitudes accompanying them, and particular actions and events that take place. In other words, Austin clearly attempted to reduce legality to patterns of conduct and attitudes exhibited by individuals.<sup>12</sup> As I explained elsewhere in some detail (2011: ch 2), H.L.A. Hart (1961) shared the same theoretical ambition. His dispute with Austin was not about the reductive nature of the explanation sought, but about the particular building blocks that Austin had offered. Hart argued that Austin's building blocks are too crude and inadequate to the task. Nevertheless, Hart clearly shared Austin's theoretical ambition to explain the elements that ground legality by way of reducing them to more foundational facts about people's actual modes of conduct, their shared attitudes and dispositions, and various actions and events that take place in the world. As we saw in chapter 1, the hallmark of Hart's theory of law is the idea that social rules are at the foundation of law, and that social rules, in turn, can be explained by way of a grounding-reduction. With the exception of Kelsen (and, perhaps, to some extent Raz), this project of trying to articulate a grounding-reductive explanation of legality is characteristic of the legal positivist tradition.<sup>13</sup>

Hart's main critique of Kelsen's theory of law is actually very instructive here because Kelsen explicitly rejected any attempt to provide a reductive account of legality; Kelsen thought that legality cannot be reduced to other types of facts, it must be theoretically presupposed.<sup>14</sup> Hart and Kelsen shared the view that basic power structuring

---

<sup>12</sup> Austin (1832) had also sought to provide a reductive explanation of law's normative aspects. In that respect, his project is widely regarded as a spectacular failure. But the truth is that there are some reductive ambitions in Hart's account of law's normativity as well, though clearly the latter is much more sophisticated. This is not an issue I will take up here at all.

<sup>13</sup> Kelsen's theory of law spans many books and articles; most of the main ideas are articulated in his *General Theory of Law and State* (1945, 1961) and earlier, in his *Pure Theory of Law 1<sup>st</sup> ed.*, translated to English as *Introduction to the Problems of Legal Theory* (1934, 2002). See also Marmor (2011: ch 1) where I explain, in some detail, Kelsen's anti-reductionism.

<sup>14</sup> Hence he called his theory A *Pure* theory of law. (Kelsen, 1934).

rules are necessary foundations of a legal system; they both thought that this normative framework is hierarchical, eventually deriving from one fundamental norm in every legal system from which others are derived in various ways. Kelsen called it the *Basic Norm*, Hart called it the *Rule of Recognition*. But for Kelsen, the basic norm was a postulate, a presupposition we need to make in order to make sense of the idea of legal validity. And this is precisely the main point that formed the target of Hart's critique. In direct response to Kelsen, he argued that there is no need for presuppositions about the basic norm, the norms in play are social rules, reducible, as we have seen, to patterns of conduct and attitudes of people whose legal system we observe. As Hart put it, facts about the rules of recognition "can be established by reference to actual practice: to the way in which courts identify what is to count as law, and to the general acceptance of or acquiescence in these identifications." (Hart, 1961: 105)

In other words, Hart's theory of legality, his answer to the ontological question about the nature of law, is explicitly reductive. What Hart sought to achieve is a full grounding-reduction of legal validity by way of a two stage grounding: legality is grounded in social rules and events in the world that gain their particular legal meanings by those rules, and social rules are grounded in patterns of conduct and their accompanying attitudes, as we saw in detail in chapter 3.

As I mentioned in chapter 2, reductive projects in philosophy are inherently ambitious because they fail if the reduction is not complete. Partial or incomplete grounding does not satisfy the conditions of a reductive explanation. It is, therefore, not surprising that objections to legal positivism came up by way of showing that an attempt to reduce legality to institutional facts is bound to be incomplete. In particular, as many have argued, truths about morality are also in play, always and necessarily so, according to Dworkin, or at least sometimes, as others claimed. I want to focus here on Dworkin's critique of Hart's reductionist project because there are two stages in which Dworkin developed his critical arguments, and they nicely capture two different critiques of Hart's grounding-reductionism. Let me call them the argument from judicial reasoning, and the argument from interpretation, respectively. Both aim to show that legality cannot be reduced to institutional facts because legality depends, at least partly, on moral truths. My aim here, however, is not to defend Hart's particular theory of legal ontology, but to show

the methodological differences between the two approaches and how they bear on the nature of a metaphysical account of a social practice like law.

**a. The argument from judicial reasoning.**

It is common ground to critics of legal positivism to argue that, at least in some cases, truths about morality, what is morally right or wrong, determine the legality of norms. The standard move here, especially among legal philosophers in the common law tradition, is to draw this conclusion from observing judicial behavior and taking seriously ways in which judges actually reason to their legal opinions. We are reminded that judges, sometimes at least, take principles of justice, fairness, and morality to bear on the legal validity of norms they rely on. The metaphysical import of this critique, however, stems from the widely shared assumption that even if facts about morality are social facts (ontologically dependent on human interactions), they are not institutional facts. More precisely, the metaphysical interest in the anti-positivist stance is closely tied to the assumption that facts about morality are not fully grounded in facts about people's patterns of behavior and shared attitudes.<sup>15</sup> Why is that? If you hold the view that whatever people believe to be morally right – individually or collectively -- is, ipso facto, morally right, then the fact that legality is partly a matter of moral truths is not going to change the building blocks that ground legality; it would just add another kind of beliefs to the inventory of the building blocks that ground legality. In order to show that a grounding-reduction of facts about legality is bound to fail because legality depends on moral truths, one needs to assume that facts about morality are not fully constituted by whatever it is that people believe to be morally right. This does not seem to be problematic because very few philosophers would adhere to a simple-minded version of moral subjectivism, whereby attitudes about moral issues are fully constitutive of moral truths. And yet this plausible assumption in the background entails that anti-positivism must concede that comprehensive errors about moral judgments are possible. After all, if what is morally right or wrong is not fully grounded in what people believe to be morally

---

<sup>15</sup> Dworkin clearly recognized this, and made an effort to substantiate his meta-ethical views in ways that align with his legal argument. (Dworkin, 1994) Finnis (1980 & 2012) also made it very clear that his views on the nature of law are closely tied with his cognitivism about morality.

right or wrong, then it is possible that we might be comprehensively mistaken about moral truths. But then, if truths about morality, sometimes at least, ground legal validity of norms, it follows that comprehensive errors about legality are also possible. And this, in my mind, is the Achilles heel of the anti-reductionist argument under consideration.

Dworkin's theory of *Legal Principles*, published in the 1970s (Dworkin, 1977: ch 2), is the most famous and perhaps still best developed argument purporting to show that a grounding-reduction of legality is bound to fail because the kind of facts that determine what counts as a legally relevant consideration is often a matter of moral truths. Ordinary legal rules may well result from deliberate enactment by various authorities, Dworkin argued, but the law also relies on principles deemed legally binding, and those principles are derived, partly but necessarily, by way of moral arguments. Judges reason to the legality of a principle by striving to articulate the best moral justification of the relevant body of law. Therefore, what is morally right may determine what is legally valid. Now the details of Dworkin's distinction between rules and principles, and the details of the argument itself, do not matter for our concerns here. I have expressed various doubts about it elsewhere.<sup>16</sup> What matters is the plausibility of the conclusion, and the conclusion is shared by many critics of positivism who do not necessarily adhere to Dworkin's theory of legal principles. Generally, then, the anti-positivist critique of Hart's reductionism relies on three main assumptions:

- (1) Facts about legal validity are sometimes determined by facts about what is morally right or wrong.
- (2) Facts about what is morally right or wrong, of the kind that (partially) grounds legality, are not fully grounded in what people believe to be morally right or wrong.
- (3) If there are facts about morality that are not fully grounded in what people believe to be morally right or wrong, then comprehensive mistakes about what is morally right or wrong are possible.
- (4) Therefore, comprehensive errors about what is legally valid are possible.

Since we are largely inclined to think that (2) & (3) are correct, anti-positivism clearly entails a commitment to (4). But the problem is that (4) is rather implausible. It makes

---

<sup>16</sup> See Marmor (2011: ch 4).

little sense to suggest that an entire community of legal experts and legal practitioners would get their laws wrong. What would it mean to get the laws wrong? There are two main possibilities here. People can get their laws wrong either thinking that a norm is legally valid, whereas in fact, it is not (because it is immoral); or else people can fail to see that a norm is legally valid (because it has a great deal of moral support) even if the norm is not recognized as legally valid by the relevant legal actors. But both of these kinds of comprehensive mistakes are implausible. If all those who practice law in the US, for example, believe that decisions of the US Supreme Court amount to the final settlement of what the law is, and the US Supreme Court holds that the law is X, then it's very difficult to see how the law in the US would not be X because the court had made a moral mistake. Furthermore, suppose the US Supreme Court makes a systematic *moral* error over a period of time – though not necessarily a grave one -- reiterated in numerous holdings; the conclusion would be that the entire legal community follows putative laws that are not really laws, norms that lack legal validity. This makes very little sense. Now of course it's a bit more complicated than this. Laws may conflict or they may be indeterminate in all sorts of ways. It's possible, for example, that one holding of the Supreme Court is in some conflict with another holding, and therefore the law on that issue might remain in some doubt. But legal conflicts and indeterminacies aside, the idea that an entire community of legal experts and practitioners would misidentify what is legally valid in their jurisdiction, in either sense of misidentification mentioned above, seems rather far-fetched.

This incredulity of comprehensive errors about legality is supported by the artifact nature of law. Since law is clearly the product of human agency, it is difficult to see how comprehensive errors about what it is are possible. Can we come to discover that laws we follow are not really our laws, or that what we have taken to be laws are not so? The difficulty is compounded by the fact that law is not just a massive artifact, a product of human agency, but also, as we noted in chapter 5, a self-validating artifact, like fiction, where the saying so of the appropriate authorities normally makes it true that so. When the legislature enacts a new law, say, that “All X's ought to  $\phi$  in C”, the enactment makes it the case that it is now true that all X's ought to  $\phi$  in C, legally speaking, that is. And the same applies to a higher court ruling; when the court holds that the law is Y, it is thereby

the case that the law is Y because the court says so. Remember, however, that the self-validating nature of an artifact does not preclude the possibility of external critique of it. Certain things are rendered true in a fiction because and in so far as the fiction says so. But once we have fictional artifacts, as it were, we can say many things about them that are not necessarily rendered true by the fiction saying so. The same applies to law; once we have laws we can identify as such, there are many things we can say about them from all sorts of critical vantage points, such morality, politics, economics, etc. But this last point is not disputed.

You might find this a bit too rough and vague, so let me try to be more precise about this idea that comprehensive errors about what is law in our system are not possible. Remember that according to Hart, the grounding of legality is a two stage affair: legality, in any given system, is grounded in the social rules which determine what makes law in the system, what counts as legally relevant act, and how. These are the basic power structuring rules of a legal system. This is one grounding stage, whereby we ground the legality of particular norms or directives in social rules, and particular acts or events taking place in the world (including, possibly, customarily so). And then, as we saw in detail in chapter 3, social rules, as such, are grounded in regularities of conduct coupled with the appropriate collective attitudes. So let me start with the more foundational stage of grounding, the basic power structuring rules that determine what counts as legal in the relevant society. As we saw in chapter 6, given the fact that basic power structuring rules, the rules of recognition in Hart's terminology, are social rules, comprehensive errors about their basic contents, their definite extension, so to speak, is not possible. But remember that the impossibility of comprehensive errors is confined to the content of the social rules, not other aspects of it. People can be mistaken about the nature of the rules, their origins, or what not. If you remember the history of Europe in centuries past, the ideology of the Divine Right of the Kings was widely shared, which means that people actually believed, or so it seems, that legal-political authority is sanctioned by God. That was a mistake, and a fairly comprehensive one at that. But even under this misconceived conception of the ontological status of the rules of recognition, the content of the rules of recognition could not have been comprehensively misconceived. There is room for some vagueness, with some borderline cases in the

margins, as we argued in the previous chapter, but that is not the point of contention here.<sup>17</sup>

None of this, however, would seem to rule out the possibility that the social rules of recognition would be such that their contents render legality, either necessarily or in some cases, depend on moral truths or moral right and wrong. Perhaps in our society, the Leader's directives are law only if they are just! Let us assume that it is possible.<sup>18</sup> But then another kind of comprehensive error is made possible. Suppose the Leader directs that the law is X, and suppose that for a while we are fine with it, we do not regard X as unjust. So for a while, we regard X as legally valid, part of our laws. Could we have made a mistake about that? Can we come to discover, perhaps decades later, that X is actually unjust? As I mentioned earlier, on any plausible view of morality and justice that anti-positivism presupposes, the answer has to be yes; we could have thought that something is just when in fact it is not. But then, if anti-positivism is correct, we would also need to conclude that X had failed to be law, it has not been part of our legal order for all this time when we thought that it is. So it may turn out, on this view, that an entire community of lawyers and legal experts, including judges, were wrong, and perhaps have been wrong for a long time, to have thought that X is law. And this is what I find very doubtful. Now you might think that I find this type of comprehensive error impossible because I have assumed, without much of an argument, that laws are necessarily artifacts, products of human agency. But some, like Dworkin perhaps, may disagree; they may think that even if some laws are clearly artifacts, or some aspects of legality share features with massive artifacts of other kinds, laws are also another type of entity, the moral type, which is not entirely a product of human agency.<sup>19</sup>

---

<sup>17</sup> In later work Dworkin (1986: ch 1) went to considerable lengths to show that lawyers and judges actually disagree about the criteria of legal validity in their jurisdiction, emphasizing that the disagreements are not about borderline cases. He took this to show that there are no rules of recognition because rule following requires agreement on the content of the rules. Few found Dworkin's argument convincing on these points; the disagreements about legality that Dworkin pointed out seem to be very much in the margins, and it is not clear that they are disagreements about rules, as they are about interpretative methods, concerning ways of changing the law by judges and adapting it to novel circumstances.

<sup>18</sup> The assumption is arguendo here, I have raised doubts about its plausibility elsewhere for reasons that should not detain us here. See, for example, Marmor (2011: 92-95).

<sup>19</sup> I think that Atiq's (2020) interesting new version of anti-positivism relies on such an assumption, it can be construed as suggesting that law is partly an artifact but partly a moral issue, combining some elements of both kinds. I find this rather mysterious, but I also acknowledge that mystery is often in the eyes of the beholder.

I do not have a metaphysical argument to disprove this possibility. Pointing to the fact, and a fact it is, that countless laws are morally wrong, even iniquitous, is not a good argument. A Dworkin type of anti-positivism can allow for the view whereby legality is a mix of the moral with the product of human agency, they may come in some conflict, and sometimes we end up with laws that are morally far from good. My argument here is not an attempt to refute the anti-positivist stance with a direct argument; it is an objection based on the price tag. And the price tag, I argue, is very high: the anti-positivist argument that allows for legal validity to depend on moral truths makes it possible for a whole community of legal experts and legal practitioners to be mistaken about the laws they have, about what is legal and what is not in their community. To me this makes very little sense. Notice, however, that even if Dworkin and other critics are right and I am wrong about this, none of it would amount to a critique of the metaphysical arguments deployed thus far; the disagreement pertains to the nature of law, not the nature of artifacts or social rules, their grounding and ontology. In other words, the anti-positivist argument under consideration does not contest the idea that if and to the extent that law is an artifact, it could not be subject to comprehensive errors; it contests the assumption that law is only a product of human agency. And this may be the point in the argument where we hit rock bottom; all I can show is what follows from the fact that law is a massive artifact, essentially a product of human agency. I'm not sure how one can show that law is nothing more than that.

#### **b. The Argument from Interpretation.**

The second stage in Dworkin's critique of reductionism in legal philosophy challenges Hart's reductive ambitions from a methodological perspective, and in a way that might have much wider implications for the metaphysics of sociality generally. It is based on Dworkin's general theory of interpretation, arguing that explanations of social practices are inevitably interpretative and therefore, necessarily intertwined with evaluative judgments.<sup>20</sup> To put this in terms more closely aligned with our previous discussion, Dworkin's idea here is that reductionism cannot be reconciled with the

---

<sup>20</sup> See, mainly, Dworkin (1986), and (2006).

internal rationalization constraint. The latter calls for an hermeneutical approach, that requires any philosophical explanation of a social practice to engage, argumentatively, with the evaluative perspective of the participants.<sup>21</sup> Let me state from the outset that I think that Dworkin is right about one thing and wrong about another here; he is right to argue that an account of internally rationalizing features of social practices are inevitably interpretative, and therefor inevitably evaluative. I think that he is wrong to assume that this amounts to an argument against grounding-reductionism; the latter is not subject to the internal rationalization constraint.

Dworkin's theory of interpretation is complex, insightful, and parts of it are controversial. Since I have written on it extensively elsewhere, I will not belabor it again here.<sup>22</sup> Suffice it to acknowledge that I am not going to dispute Dworkin's claim that interpretation is partly, but necessarily, a matter of evaluative judgments. Generally speaking, I think that Dworkin (1986: ch 2) is quite right to argue that we cannot begin to interpret anything, be it a work of art or a piece of legislation, without having a sense of what kind of thing it is and, crucially, what makes things of that kind better or worse, worthy of our appreciation, and similar views about values and merit. In other words, I will just assume here that Dworkin is quite right to maintain that once we are in the business of interpretation, we cannot avoid relying on some evaluative views about merits of objects of the kind we interpret. Interpretation is not mere description, for sure. But now, in order to see why Dworkin and others would come to think that interpretation is inevitable in the context that is relevant to our discussion, let me revert to the example of works of art, and the question of whether something like *The Leaning Mirror* is a work of art or not.

When I first encountered this work in the museum, I was accompanied by a friend who happens to be a renowned art historian, and we got into an argument about it. I was somewhat dubious that what we are looking at (just a mirror stuck in a pile of sand, remember) is a work of art. But my friend objected, saying that it certainly is, thought perhaps not a great one at that. Thus we ended up having an argument about the question

---

<sup>21</sup> Historically this theme is not new; hermeneutical methods were seen as critiques of positivism in social sciences from the start, so this is just the latest iteration of debate that has been going on since mid 19<sup>th</sup> century.

<sup>22</sup> See, for example, Marmor (2005) & (2011: ch 4)

of whether an *o* is an *F* or not; whether this object is art or not? You can easily imagine the arguments we deployed here. I pushed the point that a work of art needs to demonstrate some special talent or skills in creating it, while my friend pointed out that the talent is sometimes just the ability to see something that had alluded others, not necessarily to do something that others cannot readily accomplish. I argued that an object that can be easily mistaken for a pile of rubbish if found on the street, doesn't quite get to be art. My friend reminded me of the important contribution of Dada and its counter-art revolution, and the ways in which it changed our conception of what it is the social role of art and how to appreciate it. And so and so forth, we kept arguing for a while. But of course our arguments and assumptions were closely entangled with our views about the values of art, in general, about what makes art worthy of our appreciation, and things like that. In other words, it seems to be clearly the case that different views about values we associate with art would entail different answers to the question of whether this object is a work of art or not. More generally, it is one of those cases in which the answer to the question of the type: *is o an F?* depends on values we associate with *F*.

I hope you can see that if a similar dialectic applies to law as well, a reductive explanation of legality is in trouble. After all, it is precisely the point of legal positivism's critics that views about what makes law, or particular legal domains, valuable, and reasons, in particular, moral reasons, for appreciating various aspects of law, partly determine what counts as legal. And if that is the case, then the kind of grounding-reduction of legality we explored above is not likely to succeed.

One line of response to this challenge is to question the analogy between law and art in this context. As I argued elsewhere in some detail, there are, indeed, reasons to be very careful with analogies between law and art or, at least, with taking those analogies too far.<sup>23</sup> But this is not the main line of argument I want to pursue here. Instead, I want to show that even in the context of art, a hermeneutical approach is called for only if we

---

<sup>23</sup> See, for example, (Marmor 2019). In a nutshell, the two main differences between law and art that are relevant here are these: art is typically created (or, at least, generally perceived to be created) with the intention of becoming an object of artistic appreciation, inviting different interpretations of it. Law is something that we have for practical purposes, not for purposes of appreciation or admiration. Secondly there is some sense in which ascribing artistic qualities to an object implicates a positive appreciation of it, something that is valuable or worth our appreciation. It is far from clear that this implication necessarily applies to law. By saying that it is the law in *Si* that *X*, I do not necessarily implicate anything positive about that.

seek the kind of explanation that aims to articulate internally rationalizing features of the practice. A grounding-reduction of social practices, including art, need not be constrained by that.

Let us return to the argument I have had with my friend about *The Leaning Mirror*, whether it is a work of art or not. It was, basically, an argument about internally rationalizing features of art. What we disagreed about concerned questions about what makes art valuable for us, what makes us appreciate it as a distinct form of creative achievement, what kind of creativity we value here, and things like that. It is therefore easy to agree with Dworkin that such arguments are inevitably interpretative. They combine elements of art history, that is, facts about what people regarded as art in the past at different places and in different cultures, how these things developed over time, and evaluative views about the kind of things that make us appreciate works of art, what is it that we value about it. We were trying to make sense of what the practice is and what it means for us, those who value it and participate in it in different roles. But this is not the only kind of theorizing about art one can envisage. A sociologist, an historian, or a philosopher might ask other kinds of question as well. In particular, one might be interested in what *counts as art* in different times and places and why, on what grounds? What makes it the case that people treat certain artifacts or activities as distinctly artistic? In other words, a metaphysical inquiry into the nature of art might take a very different shape: We might be interested in the kind of facts that ground a complex (and, admittedly, rather vague) set of practices, the kind of facts that would explain what makes it the case that certain activities count as activities of an artistic type or products of it as works of art, if and where they do.

Can we have a metaphysical grounding-reduction of art, fully explaining it in terms of people's conduct and shared attitudes, complying with the constraints of methodological individualism? I see no principled reason why we could not accomplish such a task. Many facts about art are institutional facts, partly dependent on social conventions that structure different artistic genres and various general expectations about what counts as an artistic accomplishment; these conventions and collective attitudes tend

to vary with location, culture and period, undergoing changes over time.<sup>24</sup> In this respect, the analogy between law and art actually holds nicely; in both cases, we have products of human agency that are significantly structured by rules and conventions; conventions and collective attitudes play a significant role in making it the case that we identify certain types of products of human agency as products of that type, laws in one case and works of art in the other. The main difference between law and art in this respect is that the constitutive conventions of law are basic power structuring rules, because the law is an authoritative institution, extensively deploying and exercising normative powers. Art is not a practical authority, and normative powers play no role (or almost no role?) in its instantiations. The constitutive conventions of various forms of art are obviously not power structuring rules; they are rules, conventions, and general social understandings that define what counts as a product with a certain artistic qualities.

Be this as it may, I would not venture to try to work out the details of such a grounding-reductive account of art or anything similar, for that matter. The relevant question for us is whether an attempt to carry out such a project would need to comply with the internal rationalization constraint? I don't see why it should. The task of a grounding account of art, or law for that matter, is not to explain the participants' normative reasons to engage in the practice. As I mentioned earlier, there are practices and institutions we have no reasons to have; they may serve functions better not served at all, such as entrenching sexism, xenophobia, or racism. But this does not mean that they are practices we cannot fully explain; whatever grounds their existence and functions would not be categorically different, metaphysically speaking as it were, from whatever grounds noble and admirable practices such as the arts or sports or law for that matter. To take an extreme, but no irrelevant, example, I think that in principle, we can give a full account of the belief systems and collective attitudes prevalent in the ancient Maya culture to explain their extensive use of human sacrifice in their religious rituals, even if there is almost nothing about that belief system that we can come to share or normatively appreciate. Which also means that it is quite possible, more than likely, actually, that our

---

<sup>24</sup> In previous work (Marmor, 2009: ch 3& 6) I argued that these types of conventions tend to come in two layers: we may have deep conventions determining what counts, generally, as art, instantiated by various surface conventions about particular genres that determine ways of creating particular types of art. And I argued that something similar applies to law as well.

explanation of human sacrifice in the ancient Maya culture would not be the kind of explanation that meets the internal rationalization constraint; the Maya people would not have been able to see our explanation as one that rationalizes human sacrifice rituals to them. That, in itself, would not make our explanation deficient in any metaphysical sense.

One might immediately object here, saying that even if we do not share, normatively speaking, the values and purposes that rationalize a practice for its participants, any philosophical account of the practice must articulate its rationalizing features, what makes the practice intelligible and desirable for those who engage in it. That is true, up to a point, but it does not entail that the rationalizing features of the practice must be of the internal kind; for the purposes of a grounding-reduction externally rationalizing features are quite sufficient. We must explain the rationalizing functions of the practice, what makes participants engage in it, but not necessarily in terms that the participants would come to recognize as their own, as their own reasons to engage in the practice or value various aspects of it. What we need for the purposes of a metaphysical explanation of a social practice is a full account of the collective attitudes and rule following behavior of the relevant population, and the functions that those rules, and the practice as a whole, serve for them. A grounding-reduction of such elements in the explanation does not have to make sense, normatively speaking, to the participants themselves. Only an interpretative account of a social practice, one that initially seeks to rationalize the practice for the participants, striving to illuminate their own self-understandings of the point of the practice and their reasons to engage in it, would have to engage with those self-understandings argumentatively, on the basis of normative reasons and values.

I think that this is precisely how Dworkin saw his own theory of law, and where he found the failures of Hart's legal positivism. Dworkin argued that the latter cannot give us an adequate account of the internally rationalizing features of law, at least not in the common law tradition. In particular, Dworkin (1986) argued that legal positivism does not have the tools to explain what is the moral-political significance for judges and other legal actors to adhere to past legal decisions, why should they care about the constraints of the past and to what extent they should. And perhaps Dworkin was right

about that.<sup>25</sup> But if I am correct to argue that a grounding-reduction theory of law should not be expected to meet the internal rationalization constraint, then Dworkin's argument misses the target, it just answers a different kind of question. From the perspective of internally rationalizing features, it may well be the case that whether an *o* is an *F* or not depends on values participants associate with *F*; that is true, and it may well be true in the law as well. However, I hope it is also clear that a metaphysical reductive account of law, or art for that matter, strives to account for a different kind of question; it purports to establish why something *counts as an F* in a given society *S*, and to the extent that we need, as we usually do, to include in the explanation various aspects of *F* that rationalize it for members of *S*, those rationalizing features need not be internal; they do not necessarily have to make sense, normatively speaking, to the participants in *S*.

In other words, the argument of this chapter shows that a suspicion many philosophers have had over the years<sup>26</sup>, that Hart and Dworkin were arguing past each other, may well be quite right. A metaphysical reductive account of law aims to achieve one thing, an interpretative account of law aims to achieve something else. Which does not mean, however, that the projects are entirely unrelated, far from it. Any theory about aspects of social practices and institutions has to make some metaphysical sense. However, within the bounds of metaphysical foundations and epistemic constraints that follow from them, there is ample room for interpretative disagreements, amongst the participants themselves, about various aspects of any social practice or institution; metaphysics, by itself, cannot settle those disagreements, they are, as Dworkin rightly claims, partly evaluative in nature. Finally, let me conclude by reminding us that law was meant to be only an example here. The differences between metaphysical and hermeneutical theories of social practices apply much more broadly, they pertain to any theoretical explanation of our institutional realities.

---

<sup>25</sup> Interestingly, Raz probably sides with Dworkin on this point; he also held the view that a theory about the nature of law should make it fully intelligible to those whose practice it is, and in ways they can come to recognize as rationalizing the practice for them. His disagreement with Dworkin is mostly about what it is that rationalizes the practice, in Raz's view, it is law's essential authoritative nature, and conclusions about legal validity that follow from it. See Raz (1994: chs 8-10, and 2009).

<sup>26</sup> Including Hart himself in the Postscript to *The Concept of Law* (Hart, 1994: postscript).