

# Taking Institutions as They Are, or as They Ideally Ought To Be?

*Andrea Sangiovanni*  
(King's College London, Philosophy)

June 30, 2013

DRAFT – PLEASE DO NOT CITE OR CIRCULATE WITHOUT PERMISSION

Note to the reader: The paper is about 12,500 words. If you are pressed for time, you can skip the last section on 'Interpretation', which brings it down to 8,000.

Imagine you were to ask the following question. What principles of social justice, if any, should apply to the European Union? Do European citizens have obligations of social justice that cross the borders of member states? If so, what are their grounds? In this paper, I want to explore a foundational issue that must be addressed *before* we can even begin to answer a question like this. The issue is this: To what extent, if any, should the justification and formulation of principles of social justice be responsive to facts about social relations and practices (e.g., facts about the EU)? If the justification of such principles should *not* be responsive to facts about the social relations and practices they are meant to govern, then European integration would only provide an instrument for realizing aims that are not EU-specific, and that could apply, in principle, to any political and social institution or system of institutions. But if, on the other hand, principles *are* meant to be responsive to such facts, we

need to know both *why* and *in what sense* they ought to be ‘responsive’. The aim of this paper is to show how and why principles of social justice can be, as I shall say, *practice-dependent*.

To show the importance of the question raised, let me reformulate it in more concrete terms. Imagine you are a *globalist cosmopolitan*. You believe that all persons deserve equal respect and consideration regardless of their place of birth, sex, age, race, or nationality. You also believe that the scope of all solidaristic obligations should, as a matter of political morality, be global in reach, and that the scope of the obligations does not depend on the existence of any social interaction. Such obligations hold among persons *as such*, and hence would hold even, let us say, in a pure state of nature (in which individuals do not stand in any social relationship). You therefore oppose those who believe that obligations of social justice only hold among citizens and residents of states, and *a fortiori* you also oppose those who believe they only obtain among fellow members of a nation.<sup>1</sup>

It follows that you would conceive of the EU as an instrument—along with all other international social and political institutions—for realizing a unitary globe-encircling pattern of distribution, whose principles we can know independently of any specific knowledge about the EU. On this view, the EU serves the ideal when it helps to bring us closer to the globally preferred pattern (primarily by functioning as a model to other regions for how to expand the scope and depth of solidaristic obligations beyond the state), and undermines it when it props up the interests of Europeans at the expense of those globally worse off (think of agricultural subsidies under the Common Agricultural Policy).

Now imagine you are a *statist cosmopolitan*. While you also believe that all persons deserve equal respect and consideration regardless of their place of birth, sex, age, race, and

---

<sup>1</sup> E.g., Simon Caney, *Justice Beyond Borders* (Oxford: Oxford University Press, 2005); Kok-Chor Tan, *Justice without Borders* (Cambridge: Cambridge University Press, 2004); Richard Arneson, ‘Do Patriotic Ties Limit Global Justice Duties?’, *Journal of Ethics* 9 (2005): 127-50; Derek Parfit, ‘Equality and Priority’, *Ratio* 10 (1997): 202-21.

nationality, you believe this entails, at most, a commitment to human rights and a general duty to assist the global poor.<sup>2</sup> However, you do not believe the idea of equal respect and consideration entails that fundamental principles of social justice more demanding than humanitarianism must be global in scope. This is because you hold that obligations of social justice are only triggered in the presence of the kinds of extensive social interaction present among citizens and residents of states. One prominent representative of this view, Thomas Nagel, contends that obligations of social justice are only triggered among those who share in upholding and imposing a comprehensive system of societal norms backed by coercion.<sup>3</sup> Because international law—and, indeed, European law—is not backed by a centralized system of coercion, principles of social justice do not apply there. Statists of this kind need not be euro-skeptics; their position only commits them to the thought that cooperation among EU member states raises no distinctive issues of justice. As long as the EU does not undermine the capacity of states to secure domestic commitments to solidaristic redistribution, then the EU is, as it were, justice-neutral.

The position I have defended elsewhere is a version of *internationalism*.<sup>4</sup> Along with both statists and globalists, internationalists are cosmopolitans insofar as they believe that all persons deserve equal respect and consideration. They also share with statists the position that obligations of social justice are only triggered in the presence of relevant forms of social interaction. But, contrary to statists, they do not believe that international relations are, beyond a human rights/humanitarian floor, a justice-free zone. Obligations of social justice *do* apply

---

<sup>2</sup> E.g., John Rawls, *The Law of Peoples* (Cambridge: Harvard University Press, 1999); David Miller, *National Responsibility and Global Justice* (Oxford: Oxford University Press, 2007); Thomas Nagel, 'The Problem of Global Justice', *Philosophy & Public Affairs* 33 (2005): 113-4; Mathias Risse, 'What to Say About the State', *Social theory and practice* 32 (2006): 671-98.

<sup>3</sup> Nagel, 'The Problem of Global Justice'.

<sup>4</sup> On the EU, see Andrea Sangiovanni, 'Solidarity in the European Union', *Oxford Journal of Legal Studies* forthcoming, 2013); on the state, see Andrea Sangiovanni, 'Global Justice, Reciprocity, and the State', *Philosophy & Public Affairs* 35 (2007): 2-39.

at the international level. The key element that distinguishes internationalist views is that the content and scope of fundamental principles of social justice *varies with the type and extent of social interaction involved*. Statism has a binary structure: either the relevant relations hold, and the full panoply of social justice obtains, or the relevant relations do not hold, and then only the humanitarian/human rights floor applies. Internationalism has a multinomial structure: different principles of social justice apply to different types of social and political interaction, depending on the kind of interaction that the institutions instantiate.

As a result of its more complex structure, internationalism faces a methodological challenge: How does one go about identifying the correct principles for different political and social institutions (such as the nation-state, or the WTO, or the EU)?

## **I. The central challenge**

We need principles because we need guidance. And we need guidance because we face a particular set of practical problems in the here and now.<sup>5</sup> To what extent should we conceive of those principles as dependent on the structure and nature of the problems? No one disagrees that such principles will ultimately have to be applied and implemented and adapted in response to the actual, contingent circumstances we face here and now. And no one disagrees that, in applying and implementing and adapting those principles, we won't be able to realize their demands to the fullest possible extent. There will have to be tradeoffs and sacrifices of all kinds, even in the best of cases (hence the need for what is referred to as 'non-ideal' theory). It is often said that things are very different, however, with respect to the *content*, *scope*, and *grounds* of the principles *themselves*. Principles and values<sup>6</sup> give us a critical

---

<sup>5</sup> G. A. Cohen does not deny these claims. He does not deny that we need principles for guidance; rather, he denies that this fact affects their content, scope, or grounds in any relevant sense. This is an important and often overlooked distinction. More on this below.

<sup>6</sup> In the rest of the paper, I will no longer refer to principles *and* values. I will simply refer to principles. The distinction between the two is not important to the points made here.

perspective on which to judge the problems, whatever they are. To think of them as somehow constrained or otherwise dependent on the way the problems present themselves would vitiate and artificially circumscribe their very point. And, worse than that, it would entrench an arbitrary and unfounded status-quo bias into the very heart of moral and political theory.

This is clearest when the problems we face are occasioned by the unfolding of a practice. A practice is a system of conventional norms that assigns roles, positions, offices, and the relations among them. Marriage, friendship, the state, promising, boxing, the novel, games, parliamentarism, slavery, foot-binding, helotry, jury duty, human rights, and academic assessment are practices. How should we evaluate a practice? Evaluation comes in four different types. Evaluation can be *esthetic*—is the practice in question beautiful? Does the practice enable the production of beautiful objects and relations?; *pragmatic*—does it achieve its ends efficiently and effectively?; *ethical*—does it contribute to a good life?; or *moral*—can it be justified to those who would rather change or end it? In this paper, we are concerned primarily with evaluation from a moral point of view (although we will trace an important analogy to other forms of evaluation below). Imagine that we query a practice like traditional marriage and ask: Can the restriction of traditional marriage to heterosexuals be justified to homosexual couples who would like themselves to be married? And suppose someone claims in response that the principles appealed to in resolving the question should depend for their content, grounds, or scope on the nature of traditional marriage itself. Surely one could reply: But how can it be that the mere existence of traditional marriage can give it any normative standing? Wouldn't starting in this way arbitrarily bias the question against those who want to criticize the practice? Isn't the whole point of critical evaluation to allow us to step back from the practice?

One might then infer that the content, grounds, and scope of the principles one ultimately appeals to in evaluating *any* practice must be independent of the practice for just the

same set of reasons. To constrain the content, scope, or grounds of a principle to a practice is to give the mere existence of the practice a normative power that it could never have on its own. No principle or value can convincingly be defended by appeal to just the existence of a practice. There must be some further reason why the value or principle or norm that is part of the practice itself promotes some good end, or one we are bound to follow. Again it is worth recalling the distinction between the content, scope, and grounds of principles and their application to the ‘real world’: while of course such principles, once we know what they are, might be *applied* to practices like marriage, their normative standing lies, as it were, upstream. Call this the *External Principles Thesis* (EPT).

This general picture supports a distinction between, on one hand, *fundamental* principles whose basic formulation and justification must proceed entirely independently of facts about practices, and on the other, *rules of regulation* which are formed by applying fundamental principles to practices. On this picture, facts tell us ‘how much’ justice can be achieved in different situations, but they do not inform, in any relevant sense, what justice *is*.

Cohen has done the most to explore the implications of this distinction. He writes,

if a fact F enables more of J to be brought about than not-F does, it will then prove true to say that more justice can be achieved if F is true than if not-F is true. F and not-F will not determine what is (straightforwardly) just in different situations, but how much justice can be achieved in different situations. Facts [including facts about practices—A.S] thereby make a difference to what are the right rules of social regulation, rules the effect of whose operation may justifiably deviate from justice itself.<sup>7</sup>

Rules of regulation tell us what kinds of social relations and practices we ought to adopt for our world; such rules are devices for having certain effects given background empirical circumstances. It is natural to think that the correct rules of regulation should therefore depend on the facts. The set of correct principles of *justice*, however, must be independent and prior to the set of correct rules of regulation, precisely because they provide one among many standards

---

<sup>7</sup> G. A. Cohen, *Rescuing Justice and Equality* (Cambridge, MA: Harvard University Press, 2008), p.???

for those rules to live up to. It is important to emphasize how expansive Cohen's arguments are: by 'justice' Cohen means to include nothing less than all the moral and political principles comprehending the moral domain that Scanlon has identified as 'what we owe to each other'.<sup>8</sup> These claims are then buttressed by a regress argument, which I will refer to as the *Fact-Principles Thesis* (FPT), that purports to show that any affirmation of a fact-sensitive principle, P1, necessarily commits one to a fact-insensitive principle, P2, that explains why the facts ground principle P1. If we take the facts to refer to practices, then this argument also entails the conclusion that we reached earlier, namely that the mere existence of a practice cannot by itself give us a conclusive reason to do anything (unless there is some higher-level, external principle that tells us to do what the practice says).

The next stop on this train of arguments is the oft-mentioned division of labor between moral-political philosophy and the social sciences. Adam Swift and Stuart White are representative when they write,

It is for the empirical, descriptive/explanatory, social scientific disciplines to (try to) tell us which states of affairs can be realised by what means ... given where we are now. But it is for philosophy to tell us which of those states and means of achieving them are better and worse than one another.<sup>9</sup>

The idea is that any non-question-begging normative political and moral philosophy should at some point aim to justify and formulate *basic* or *fundamental* principles, namely principles which are not derived by application from any further principle or fact.<sup>10</sup> While of course political philosophers may themselves *apply* such fundamental principles to specific circumstances (such as, in Swift's case, the family), the normative standing of the higher-level principles *themselves* lies upstream from whatever they are being applied to. On this view,

---

<sup>8</sup> Cohen, *Rescuing Justice and Equality*, p.???

<sup>9</sup> Adam Swift and Stuart White, 'Political Theory, Social Science, and Real Politics', in *Political Theory: Methods and Approaches*, eds. D. Leopold and M. Stears (Oxford: Oxford University Press, 2010), pp. 49-69 at p.???

<sup>10</sup> They may, however, be supported by networks of further values or facts as long as they are not derived from them. For example, they may gain support by being consistent and coherent with other fundamental principles and convictions one holds. See footnote xxx, below, for more on this point.

philosophers set out the fundamental questions and structure of values that ought to govern decision-making within any institution or practice or relation, and it is a further task, mostly left to practitioners or the ‘explanatory/descriptive’ sciences, to figure out how best to implement them.

Summarizing, we have the following set of claims:

- (1) The content, grounds, and scope of the principles that we use to evaluate a practice from a moral point of view cannot themselves be ultimately constrained or bound by the contingent properties of the practice itself. Otherwise, our affirmation of the principles would reflect an arbitrary status-quo bias. (The EPT)
- (2) Therefore, there must be *fundamental principles* whose content, scope, and grounds are independent of any practice, and *rules of regulation* that are the result of applying such higher-level principles to specific practices.
- (3) (1) and (2) are supported by the FPT, which holds that fundamental principles must be fact-insensitive, i.e., they must survive the denial of any facts (including facts about the practices to which they may be applied).
- (4) From (1), (2), (3): Facts (including facts about practices) tell us how much, say, justice (or freedom, legitimacy, equality, solidarity, etc.) can be achieved given the way the world is, but they do not tell us what justice *is*.
- (5) (1), (2), (3), and (4) support a division of labor between moral and political philosophers, whose primary task is to develop, formulate, and defend fundamental principles independently of the facts (including facts about practices), and practitioners and social scientists, who are meant to apply such principles to the ‘real world’ in full view of the facts.

From an uncontroversial starting point regarding the purpose of any critique of a practice, namely the EPT, we arrive to a sweeping claim about the basic nature and structure of normative moral and political philosophy. It is this overall pattern of argument that I want to challenge. I shall contend that (1) and (2) are innocuous, (3) is trivial, (4) doesn't follow and is in fact false, and (5) is misleading. As we shall see, the summarized pattern of argument overlooks the diversity of ways in which 'application' of higher-level principles and value-concepts enters into any plausible, substantive moral and political philosophy. Our main aim is to draw attention to this diversity. In the second part, I then seek to draw some methodological implications of our discussion by demonstrating the importance of *social interpretation* to any moral or political philosophy that takes such diversity seriously. In the process, I will show how facts about practices can enter into the justification and formulation of moral and political principles, and so answer the question with which we started.

## II. Three ways in which principles can depend on practices

Before turning to our main line of reasoning, we first need to be clear about three ways in which a principle might 'depend on' or be 'responsive to' facts about a set of practices. I make these distinctions because they are essential in locating the challenge posed by our initial train of arguments.

The general form of the practical principles we will be interested in is this: 'For some  $z$  and  $R$ , and for all  $x$  and  $y$ ,  $x$  and  $y$  ought to  $z$  if (or only if or iff<sup>11</sup>)  $xRy$ ', where  $R$  denotes a variable ranging over kinds of social relation. A social relation exists, in the sense I am using the term, if and only if individuals share either ongoing, relatively stable forms of interaction

---

<sup>11</sup> The relevant social relation could state either a sufficient or a necessary and sufficient condition for the application of the principle. For example, one might hold that egalitarianism applies to all those who share a cultural identity and leave it open whether it might also apply among those not so related, or one might say that egalitarianism applies all *and* only to those who share a cultural identity. Both cases count as asserting principles with relational scope.

outside of any practice, or are joint participants in a practice (and such participation is common knowledge). Examples of the former include sporadic trade or loosely organized mutual protection. Examples of the latter include sharing a cultural identity, or cooperating in a system of mutual production, or being subject to a system of ongoing coercion, or being married, or being friends. The universal quantifier ‘for all  $x$  and  $y$ ’, in turn, specifies a particular *domain*. For *non-relativists*, this domain is usually taken to include all and only human beings, or all and only persons, or all and only sentient beings. For *relativists*, on the other hand, the domain will be bound by a particular, time-indexed cultural group. This is the first way a principle may depend on facts about practices: If one is relativist, the principles that apply will depend on what the members of that group (and hence participants in the set of practices constitutive of the group) happen to value. But there are many well-known problems with relativism, so I leave it aside for the rest of this paper. The more interesting question is how *non-relativist* principles might depend on practices. For the rest of the paper, I will therefore assume the domain for all the principles discussed to include all and only human beings, and so assume *arguendo* that relativism is false.

By *content*, I mean whatever particular actions or attitudes are specified by  $z$ . By *scope*, I will refer to the range of the social relation ‘ $xRy$ ’.<sup>12</sup> For example, if  $R$  takes cultural identity as a value (as in ‘for all human beings  $x$  and  $y$ ,  $x$  and  $y$  ought to distribute goods equally iff  $x$  and  $y$  share a cultural identity’), this is equivalent to saying that the scope of egalitarianism ranges over all and only those who share a cultural identity. This is the second way in which principles can depend on facts about practices: if a principle is triggered by the existence of a practice, then

---

<sup>12</sup> The *extension* of a principle would, in turn, refer to the range of individuals in the actual world who fall under the principle’s scope. This last category is important in the global justice debates, since it will determine whether the principle has a global extension or whether the principle only applies to smaller subsets of the world’s population, such as only among citizens and residents of states. If, for example, all individuals across the globe shared a single cultural identity, then the relational principle just mentioned would have a *global* extension. I leave this fourth category aside in what follows since it specifies a way in which a principle has implications for practices rather than being dependent on them.

the principle depends, in this sense, on facts about the practice. We could then say that the practice provides the conditions in which the principle applies. I shall refer to principles that are only triggered in the presence of a social relation (whether practice-based or not), *relational*, and those that apply to human beings as such, independently of the social relations in which participate, *nonrelational*.<sup>13</sup>

By *grounds*, I refer to whatever justification is offered for a principle. A principle of egalitarian social justice with a relational scope might be justified by appeal to the existence of a set of social relations in conjunction with a higher-level moral principle or value. For example, the culturally specific relational egalitarianism mentioned above might be justified by appeal to a higher-level moral principle that instructs one to respect aspects of a person's cultural background that are constitutive of their identity, plus an account of how socioeconomic egalitarianism is essential to the self-understanding and cultural practices of members of a particular group. In this case, the grounds of the principle would (partially) depend on facts about the practice, namely facts about how socioeconomic egalitarianism is central to the cultural group's self-understanding. This is the third way in which principles might depend on facts about practices. As we will see in a moment, it is with regards to this type of practice-dependence that most of the action lies.

### **III. The *Facts-Principles Thesis***

The form of practice-dependence reflected in the distinction between relational and nonrelational principles is uncontroversial: no one denies that there are at least some relational principles. The more difficult question is whether higher-level relational principles can serve as final bases for lower-level principles, or whether all justification must ultimately bottom out in a

---

<sup>13</sup> Notice that nonrelational principles can drop the 'xRy' from the specification of their principles (as in utilitarian principle: 'for all sentient beings *x* and *y*, *x* and *y* ought to act so as to maximize the satisfaction of preferences').

*nonrelational* principle.<sup>14</sup> This is a question about the *grounds* of a principle and the place of facts (and, in our case, facts about practices) within it, and so it touches on the third way in which principles might depend on practices.

I now want to argue that the *Fact-Principles Thesis* (FPT) does nothing to establish the conclusion that all relational principles must ultimately presuppose a nonrelational principle from which they are derived. I don't in fact believe that Cohen would disagree, but it is important to clarify this point since much more follows from it than one might at first expect, and since it might appear, at first glance, that Cohen's argument can be used for just this purpose.

Let us unpack Cohen's argument a bit more carefully. The FPT states that whenever a fact supports a normative principle, it does so in virtue of a more ultimate principle that is not supported by any facts.<sup>15</sup> A fact counts as supporting a principle, for Cohen, when the denial of the fact would entail the denial of the principle. I will refer to this special kind of fact-dependence as fact-dependence<sub>c</sub> to remind the reader of its special sense. Suppose, for example, that the principle P, which states that we should keep our promises, is supported by a fact F, namely that only when promises are kept can people pursue their projects. From this it follows that if F is false—if, that is, people can pursue their projects whether or not people keep their promises—then the principle is false. The principle is thereby said to 'depend<sub>c</sub>', since the denial of the fact leads to the denial of the principle.

Cohen's thesis is then that the truth of any such fact-sensitive<sub>c</sub> principle must depend on some further, higher-level fact-independent<sub>c</sub> principle that explains *why* F supports P. In

---

<sup>14</sup> The idea of 'bottoming out' in some principle or other seems to presuppose a foundationalist picture of justification. Cohen claims to be neutral with respect to the debate between foundationalism and coherentism or holism about justification, and I concede the point. My challenge to Cohen, in other words, doesn't depend on the truth of coherentism or holism (just as Cohen claims his does not depend on the truth of foundationalism).

<sup>15</sup> A normative principle, in turn, is a general directive that tells us what we ought (or ought not) to do.

Cohen's example, the higher-level principle is this one: we ought to promote people's projects. That further principle survives denial of fact F, namely that promise-keeping promotes people's projects. We ought, that is, to promote people projects whether or not promise-keeping is a good way of doing so. If the regress ends there (Cohen says, and I agree, that it must end somewhere), then the ultimate grounds for the principle of promising are not dependent on any facts. I believe the argument is sound.

But there is a crucial ambiguity at this point in the argument. Take the principle 'one ought to promote people's projects'. Is that principle relational or nonrelational? To answer this question, we need to know whether it applies to human beings as such, independently of any social relations they might participate in. Let us say that the principle is relational, and only holds among those who stand in some social relation R. If you do not stand in this relation with respect to me, I have no obligation to promote your projects. Must Cohen then say that there must be some still higher-level, *non*relational principle that explains why the principle only holds among those who stand in R? No, and for good reason: conditions of applicability are one thing, fact-dependence another. Cohen never denies that fact-insensitive, fundamental principles can come with scope restrictions. Cohen writes,

The "lovers of sights and sounds" in Book V of Plato's Republic think it suffices for saying what justice is to say what counts as just within the world of sights and sounds. They scarcely recognize the question: What is justice, as such? In a world where the facts are F, they believe that P constitutes justice, and they do not abstract even so far as to see that they believe, independently of the facts, principles of the form: if F then P. Plato thinks, and I agree, that you need to have a view of what justice itself is to recognize that justice dictates P when F is true. That is how justice transcends the facts of the world.<sup>16</sup>

This is exactly right. Principles of the form, 'if F, then P' are *fact-independent*. They are fact-independent because the truth of 'if F, then P' *survives the negation of F*. Negating the antecedent, after all, doesn't make a conditional false. Because the assertion of a conditional

---

<sup>16</sup>

doesn't require affirming the antecedent, the principle 'if F, then P' would be true even in worlds where the facts F never materialize.

But if that is true, then one can easily make any superficially fact-sensitive<sub>c</sub> principle into a fact-insensitive<sub>c</sub> principle by simply rephrasing it as a conditional. For example, the fact-sensitive<sub>c</sub> principle, 'one ought to keep one's promises because it promotes people's projects' can simply be rephrased, without loss of meaning, as 'if promise-keeping promotes people's projects, then one ought to keep one's promises'. The latter principle is entirely fact-insensitive<sub>c</sub>: it would be true even if promise-keeping *doesn't* turn out to promote people's projects. There is also no sense in which the principle necessarily commits one to the further principle that 'one ought to promote people's projects'. The truth of that further principle can be left open. This is of course also true of the higher-level principle 'you ought to promote someone's projects only if social relation R holds between you and this other person'. This principle is true (*arguendo*) even if there are no (and will never be any) social relations R. And *a fortiori* it is also true of *any* principle that takes a relational form. And, again, it does not necessarily commit one to the principle that, say, 'one ought to seek out relations R'. That is left entirely open (and, indeed, may be false consistent with the affirmation of the relational demand). It may, for example, be entirely morally neutral whether one ought (or ought not) to seek out relations R.

To my mind, the possibility of simply recasting any fact-sensitive<sub>c</sub> into a fact-insensitive<sub>c</sub> principle makes the FPT trivial.<sup>17</sup> But even if the FPT isn't trivial, it is compatible with the claim that relational principles can function as *final* bases for other principles—as grounds, that is, that do not presuppose some further nonrelational principle. This has important and

---

<sup>17</sup> A similar conclusion is reached by Thomas Pogge, 'Cohen to the Rescue!', *Ratio* 21 (2008): 454-75.

surprising implications for the third form of practice-dependence I isolated above, to which I now turn.

#### **IV. Two kinds of judgment, and the direction of justification**

So far we have distinguished three ways in which principles can be practice-dependent, argued that practice-dependence with respect to the *grounds* (but not *scope*) of a principle is controversial, and established that Cohen's FPT does not demonstrate that all relational principles must bottom out in a nonrelational principle. Where does this discussion leave us with respect to the *External Principles Thesis* with which we began?

A principle is *internal* to a practice (or type of practice) when its normative force depends on the existence of a particular practice. A good example is the culturally specific principle of socioeconomic egalitarianism: You ought to distribute equally because you are Swedish (and egalitarianism is a constitutive feature of Swedish national identity). It counts as internal because its normative force depends on the specific character and content of Swedish political culture. Had egalitarianism not been essential to Swedish identity and self-understanding, egalitarianism would not have applied. A principle is *external* to a practice when its normative force does not depend on the existence of a practice. A good example of an external principle is the higher-level principle that might justify the internal principle just mentioned, namely 'you ought to distribute equally only if you are a member of a political culture that is supportive of egalitarianism'. Note that this principle is relational but external. While its normative force does not depend on the existence of any specific cultural practice, its content (egalitarianism) is only triggered in the presence of the right kind of social relation. And general principles of reciprocity, about which we will have more to say below, such as 'one ought to give a fair return to those from whom one has voluntarily accepted benefits or from whom one has received benefits that one could not reasonably decline', are also relational but

external. Again, this is because the normative force of the principles does not depend on the existence of any one specific practice of reciprocity.

With these distinctions in hand, we can now restate the EPT: All *internal* principles must be ultimately justified by reference to an *external* principle. If we allowed internal principles to stand as the basis of a chain of reasoning, then we would have to allow that the mere existence of a practice, such as a specific culture, had a normative power it could never have. To the question, ‘Why ought socioeconomic egalitarianism only be taken to apply among members of *this* rather than *that* cultural group?’, we would either have to take a relativist stand on the domain of justice, or simply say that the internal principle applies because the practice says so. But why take the practice so seriously? The EPT urges us either to explain why (with reference to some external principle), or be saddled with an implausible and unmotivated conservatism.

The EPT strikes me as correct: all internal principles must ultimately be supported by an external principle. And I am also happy to grant an implication of the EPT, namely that all fundamental principles are external, and that we can call (if you like) all internal principles ‘rules of regulation’. But conceding the EPT as well as the distinction between fundamental principles and rules of regulation does nothing to establish that practices (or other facts) only matter in determining *how much* justice (or liberty, equality, etc.) can be achieved, but not in determining what justice *is*. On this view, practices (or other facts) are merely instruments or tools for realizing the states of affairs, relations, actions mandated by principles. This further claim ([4] above) not only does not follow from anything we have said so far but is also, I will now argue, false.

In making this point, I want to contrast two very different ways in which one might apply an external principle to a practice. Only the first is compatible with claim (4). The second

will allow us to demonstrate that claim (4) is false, and to clarify in what way practices can enter into the justification and formulation of moral and political principles.

I draw the distinction via a running example, namely filial obligations. It is often said that the existence of relationships between children and their parents strengthens (preexisting) negative and positive obligations, and sometimes generates new ones. Some examples include: It is worse, all else equal, to physically assault your father than a stranger; given a choice between saving your father and a stranger, one has an obligation to save one's father rather than a stranger; one has an obligation to take one's father in when he can no longer care for himself, but not to take in strangers; and so on. Call the set of all such lower-level filial obligations, O. What might justify the existence of such obligations?

Notice that to reply 'Because current cultural understandings of the family require one to discharge such obligations' would fail the EPT. If we believe that we ought to respect such obligations, we need to explain why. Is it because they are part of our identity? But then how do 'identities' generate obligations? Is it because of some other feature of the practice? What external principle, in other words, warrants our affirmation of such filial obligations?

*Instrumental.* *Instrumental* claims that acting according to O turns out, in the long run, better to promote realization of the utilitarian principle, 'one ought to maximize the aggregate well-being of all persons', than attempting to follow principle of utility directly. In Goodin's terms, the obligations O are assigned to children merely as an 'administrative device for discharging our more general duty [to maximize well-being] more effectively'.<sup>18</sup> This explanatory strategy is *instrumental* because the relations are means to realizing some independently specified end. There are no noninstrumentally relevant features of filial relationships that play a role in explaining why we have obligations O. The principle 'O if we

---

<sup>18</sup> Robert E. Goodin, 'What Is So Special About Our Fellow Countrymen?', *Ethics* 98 (1988): 663-86 at p. 685.

stand in a relevant filial relationship' is just a means (in our world) of promoting aggregate well-being, and derives all its justificatory force from that fact.

Summarizing:

- (6) One ought to maximize the aggregate welfare of all persons.
- (7) If one stands in a filial relationship, acting according to obligations O better promotes aggregate welfare than the alternatives.<sup>19</sup>
- (8) One is a child in a filial relationship.
- (9) One has obligations O towards one's parents.

The scheme here is fully compatible with claim (4). Notice that the first premise is an external principle, the second a factual claim regarding the effect on aggregate welfare of following the internal norms of a practice, and the fourth a lower-level principle regarding what one, as a child, ought to do with respect to one's parents. (7) tells us 'how much' aggregate welfare can be achieved via various feasible options given the world as it is. The next schema we will consider works in very different way.

***Mediated deduction.*** *Mediated deduction* shows that O follows from a fundamental (and hence external) principle in conjunction with what I will call a 'mediating' principle and some empirical facts. A good example of this schema appeals to a higher-level principle of reciprocity:

- (10) If one has engaged in long-standing, mutual, and beneficial interaction, then one has an obligation to give a fair return to those involved;
- (11) Practice-mediated relations between children and parents in modern families normally give rise to long-standing, mutual, and beneficial interactions;
- (12) Given the nature and structure of such filial relations, the demand to give a fair return for benefits received generates and grounds a specific set of obligations O;
- (13) One is a child in a practice-mediated filial relationship of this kind.
- (14) One has obligations O towards one's parents.

---

<sup>19</sup> This conditional obligation takes narrow scope (unlike normative requirements). The reason one can't wide-scope a conditional of this kind is straightforward: It would imply that if we ought to relate in way R, then we ought to act according to S (even if we don't *actually* relate in way R). This follows directly from the application of the distribution axiom common to both standard modal and deontic logic  $[O(p \rightarrow q) \rightarrow (Op \rightarrow Oq)]$ .

There are at least three ways in which this schema is different from *instrumental*. First, the facts stated in (11) do not tell us ‘how much’ reciprocity can be achieved given a set of feasible options. Rather, they state how modern family relations (normally) instantiate the types of relations mentioned in (10). Second, (12) states a further ‘mediating’ principle that specifies what a fair return requires among children and parents given the practice-mediated relations mentioned in (11). Third, the relation between the general principle of reciprocity stated in (10) and the particular, contextually fixed principle of reciprocity stated in (12) is, in *mediated deduction*, the relation of a genus to a species, rather than the relation of a means to an end, as in *instrumental*. (12) gives us an account of what reciprocity requires *for* a given practice, given the kinds of relations the practice creates and sustains.

The possibility of schemata like *mediated deduction* shows that claim (4) is false. Although *mediated deduction* is an instance of the application of a higher-level, external principle to a specific practice, the practice-mediated relations referred to in (11) and (12) are not in any sense *tools* or *instruments* for the realization of some states of affairs, or actions, or relations mandated by (10). They are not means for the realization of ‘more’ reciprocity, and they do not determine, in any straightforward way, ‘how much’ reciprocity can be achieved. Rather, the mediating principle (12) just tells us what reciprocity *is* given some specific practice-mediated set of relations.

Specific practices and the relations they constitute and regulate therefore have a crucial role to play in specifying fundamental moral and political principles. To see the point, consider some further examples of the schemata. From (10), for example, we could derive principles of reciprocity for friends, university colleagues, citizens, states, team members, etc. Consider:

- (10) If one has engaged in long-standing, mutual, and beneficial interaction, then one has an obligation to give a fair return to those involved;

- (11') Practice-mediated relations between long-term residents of a state normally give rise to long-standing, mutual, and beneficial interactions;
- (12') Given the nature and structure of the state, the demand to give a fair return for benefits received generates and grounds a specific set of obligations O';
- (13') One is a long-term resident of such a state.
- (14') One has obligations O' towards other long-term residents.

There are of course many reciprocity-based schemata like this (as many as there are different types of mutually beneficial, on-going relationship). The key point is that the obligations specified in variants of (12) will vary according to the nature and structure of the practices involved. What reciprocity requires among friends requires very different things than reciprocity among citizens, family members, and so on, precisely as a result of the different nature of the practice-mediated relations. We return to this point below, when we consider how to generate context-dependent versions of higher-level principles.

So far, we have focused on fundamental *relational* principles, where the relevance of practices will be most obvious. But there are also fundamental *nonrelational* principles whose application must adopt the pattern exemplified by *mediated deduction*. Consider:

- (15) One owes respect to all *x*;
- (16) Given the nature and structure of the (practice-mediated) doctor-patient relationship, the demand to show respect generates and grounds a specific set of obligations D (if one is a doctor) and P (if one is a patient);
- (17) One stands in a practice-mediated doctor-patient relationship (as a doctor);
- (18) One has obligations D.

(15) is a nonrelational principle (if anything is). It tells us to respect all human beings whatever our relation to them. But what respect requires will vary according to the types of social

relations one stands in with others. Again, what respect requires among doctors and patients is different from what respect requires among human beings as such, colleagues, friends, lovers, citizens, and so on. Each of these practice-mediated relations, furthermore, is not a means or tool for realizing ‘more’ respect. They are simply different relations in which respect plays a varying role (in the same way that reciprocity played a varying role in the previous examples).

At this point, an advocate of the argument with which we began might object: ‘But no one disagrees that *judgment* must be used in applying higher-level principles! *Mediated deduction* is just an instance of such judgment applied to specific practices. But the higher-level principles themselves aren’t internal to the practices in any way, which you have conceded!’ This response squarely misses the point. Of course one must exercise judgment in moving from (10) to (12), or from (15) to (16). The point is that the kind of judgment at stake in *mediated deduction* is very different from the one required in *instrumental*. Seeing how and why opens up a diversity of ways in which practices don’t just serve as tools or instruments for the realization of principles and values, but condition their very formulation and justification. The problem with the argument with which we began is that it leads us to overlook this diversity in kinds of judgment, and hence to treat ‘application’ or ‘judgment’ as a kind of black box, of lesser interest to the philosopher.<sup>20</sup> Judgment becomes something we employ either as practitioners, political and social actors, social scientists, or at the point we *cease* doing philosophy, and start doing something else.

That picture of the so-called division of labor (captured in claim [5]) couldn’t be more wrong-headed. Consider that very few contemporary moral and political philosophers actually spend much time focusing on the formulation and justification of what we have called, following Cohen, *fundamental* principles. Such fundamental principles are usually taken, in a

---

<sup>20</sup> Oddly, the problem of judgment is seen as central to other areas of philosophy, but not so much in substantive moral and political philosophy.

very abstract form, as uncontroversial starting points for further reflection (or as summary conclusions). They are more like value *concepts* that only become *conceptions* of that value in the presence of particular contexts of action. This is obviously true of Rawls, where higher-level principles of reciprocity, impartiality, fairness, fraternity, and so on, are articulated, justified, and connected in terms of how they ought to regulate a well-ordered basic structure. But it is also true of Dworkin, for whom the fundamental values of equal concern and respect are interpreted in light of their role in a coercive order.<sup>21</sup> For Raz, similarly, the value of autonomy is explicitly defended as a value for a specific kind of modern pluralist society (hence the ‘social forms’ argument).<sup>22</sup> In his case, the higher-level value is something like human flourishing; autonomy then becomes a constitutive part of human flourishing only against the backdrop of a modern society. Feminists also fit the bill. Consider that many feminists defend higher-level values of non-domination or equal recognition in the context of, say, the family, or the modern workplace, or, alternatively, focus on how practices like patriarchy or pornography make impossible the realization of such values.<sup>23</sup> And finally we can say much the same thing of political realists, genealogists, and ideology critics, who can be understood as demonstrating how certain higher-level values that are championed by modern societies, such as freedom, are in fact subverted and undermined by the very practices that claim to realize them.<sup>24</sup> The latter two cases are what we might call *negative instances of mediated deduction*. They can be seen as showing *not* how higher-level values and principles ought to regulate a given set of relations but

---

<sup>21</sup> See Ronald Dworkin, *Justice for Hedgehogs* (Cambridge: Harvard University Press, 2011); Ronald Dworkin, *Sovereign Virtue* (Cambridge: Harvard University Press, 2000), p. 6.

<sup>22</sup> Joseph Raz, *The Morality of Freedom* (Oxford: Clarendon Press, 1986).

<sup>23</sup> See, e.g., feminists as diverse as Catherine A MacKinnon, ‘Sexuality, Pornography, and Method: Pleasure under Patriarchy’, *Ethics* 99 (1989): 314-4; Nancy Fraser, *Unruly Practices: Power, Discourse and Gender in Contemporary Social Theory* (Minneapolis: University of Minnesota Press, 1989); Susan Moller Okin, *Justice, Gender, and the Family* (New York: Basic Books, 1989).

<sup>24</sup> See, for example, Raymond Geuss, *The Idea of a Critical Theory* (Cambridge: Cambridge University Press, 1981); David Owen, *Nietzsche’s Genealogy of Morality* (London: Acumen, 2007); Brian Leiter, *Nietzsche on Morality* (London: Routledge, 2002); Michael Rosen, *On Voluntary Servitude: False Consciousness and the Theory of Ideology* (Cambridge: Harvard University Press, 1996).

rather how such higher-level values and principles cannot possibly do so, at least not given current conditions. We return such negative forms of *mediated deduction* below.

The striking fact is that this is even true of Cohen himself (and indeed of all plausible forms of luck egalitarianism). Take, for example, the luck egalitarian axiom (which Cohen would endorse) that no one ought to do worse than anyone else through no fault or choice of their own. First, and perhaps surprisingly, notice that it is not plausibly regarded as nonrelational. This is easily demonstrated. Take friendship. Could anyone reasonably argue that one has an obligations to make friends with people who have fewer friends through no fault or choice of their own? What about lovers? Or the classic example of the ugly. Is it some kind of injustice if the ugly aren't compensated for their misfortune? Examples like this show that the axiom comes with applicability conditions. Indeed, the axiom works best (if it works at all) in the explicitly *socioeconomic* and *political* contexts typical of modern societies. This is, to my mind, what gives Cohen's own discussions in Part I of the 'incentives argument' so powerful. Principles like 'one ought not, *modulo* a personal prerogative, take more wages than the worst off save where such wages are required to compensate for special burdens'<sup>25</sup> are applications of the higher-level luck-egalitarian axiom to a specific context, namely the context of a modern wage economy. But such application fits *mediated deduction* not *instrumental*. The argument shows us what the luck egalitarian axiom requires for the specific socioeconomic context for which it is primarily tailored. It does not show us 'how much' justice can be achieved here and now. And it does little to justify the axiom itself (other than trying to convince the reader that she already accepts the axiom). Indeed, one way to read the argument is as a kind of ideology critique: there is nothing in the point and purpose of a modern wage economy or society (despite its paeans to freedom of contract) or in liberal left-leaning

---

<sup>25</sup> See Cohen, *Rescuing Justice and Equality*, p. 401.

egalitarianism more generally that precludes Cohen's stringent interpersonal principle. Although I cannot demonstrate this conclusively here, I believe it is also true of Cohen's defense of the 'personal prerogative', which is clearly designed to allow for variation in the stringency of the principle's demands, especially in the context of personal pursuits and relationships. Explaining both the reasons for the personal prerogative as well as its variation over various domains would require some understanding of what makes personal relationships and pursuits special and worth protecting, and that requires, in turn, some understanding of the practices (such as the nature of modern occupational life and the family) in which they figure.

There are exceptions of course. Some examples include Singer's utilitarianism, Parfit's 'Triple Theory', Kamm's 'intricate ethics', Korsgaard, and Scanlon's *What We Owe To Each Other*.<sup>26</sup> In each of these cases, *mediated deduction* plays little or no role. Singer's practical ethics, for example, is a challenging, bullet-biting application of the principle of utility that relies exclusively on patterns of judgment like *instrumental*. It is also interesting that Parfit, Kamm, Korsgaard, and Scanlon's theories all verge on the metaethical. This is clearest with Korsgaard, for whom a major aim is to defend her constructivism against realist rivals, and, of course, Parfit spends more than three quarters of the two volumes discussing metaethical questions and the nature of reasons. This isn't a criticism, but it does show the very different kind of project that they are engaged in. Even with respect to a main area of substantive disagreement between Parfit/Raz and Scanlon/Kamm, namely aggregation of burdens across persons, the focus is more on how to incorporate intuitions about aggregation into the broader domain of morality

---

<sup>26</sup> Derek Parfit, *On What Matters*, vol. 1 (Oxford: Oxford University Press, 2011); Frances Kamm, *Intricate Ethics: Rights, Responsibilities, and Permissible Harm* (New York: Oxford University Press, 2007); Christine Korsgaard, *The Sources of Normativity* (Cambridge: Cambridge University Press, 2004); Peter Singer, *Practical Ethics*, 2nd edn. (Cambridge, UK: Cambridge University Press, 1993); Thomas Scanlon, *What We Owe to Each Other* (Cambridge: Harvard University Press, 1998).

rather than about whether to do so.<sup>27</sup> It is unclear, furthermore, what kind or what degree of aggregation would be permitted in the kinds of real-world political and social contexts we might need such theories to illuminate. And surely the way such aggregation would enter would vary by context. Contrast how aggregation might enter into decisions regarding one's own children (assuming we have many!) and how they might enter the organization of a health service. Similarly, Parfit's own major contribution to political philosophy, namely prioritarianism, is also radically underspecified, since we know nothing about the weights we ought to assign improvements to the badly off in any calculation about what to do. All we know is that we ought to give *some* greater, noninstrumental weight to improvements to the badly off than to the well off. Even Scanlon, as is well known, leaves almost completely open the moral reasons that can serve to reject some principle when thinking in a contractualist way. Scanlon gives us a (non-exhaustive) overview of the general types of concerns that can enter the reasonable rejection test, including concerns about fairness, priority, and responsibility, but we aren't really told how they might enter in our deliberation about any specific principle. Again, these are not intended as criticisms. My aim here is simply to highlight the distinction between what I have called *theoretical* and *substantive* theories.

So we might draw a tripartite distinction in modern ethical theory. There are *metaethical* accounts that focus on the semantics, logic, meaning, and ontology of moral claims; *theoretical* accounts (such as the ones discussed in the previous paragraph) that aim to provide a first-order characterization of the moral domain and the procedures of reasoning and deliberation central to it; and *substantive* accounts that aim to specify what principles and values there are. My claim is that substantive accounts (with the possible and interesting exception of

---

<sup>27</sup> See, e.g., the exchange between Scanlon and Parfit in *Ratio*, and Raz.

utilitarianism and perhaps pure libertarianism<sup>28</sup>) will need to employ, at some point, *mediated deduction*. As we have seen, this means that specific social relations and practices will play a central role in the justification and formulation of the principles at the heart of such substantive accounts.

Before turning to the methodological implications of this claim, I want to make a final point about the justification of fundamental, external principles. So far, I have accepted the EPT, and hence the claim that all internal principles must be supported by an external principle that explains, in part, why the internal principle binds participants in the practice. But what about the *justification* of the external principle? Consider that what we might call the *order of explanation* and the *order of justification* can come apart. So far I have assumed that the two go hand-in-hand: the justification and explanation of a judgment or a lower-level principle are the same. But this need not be the case.<sup>29</sup> In an inference to the best explanation, for example, some set of facts is taken as given, and we select the theory or account or claim that *best explains* the existence of the facts.<sup>30</sup> My wallet is gone. It could either be that someone has broken into my house and stolen it, or that my children have hidden it, or that I have misplaced it. Upon further reflection, I conclude that it must be that I have misplaced it (since my children have been gone, and there is no evidence of a break-in, and I am often absent-minded). In inferences like this, the facts function as premises, and the explaining theory is taken as the conclusion. Inference to the best explanation is pervasive in substantive moral and political philosophy. We often take some considered conviction or intuition as given (or given at least provisionally), and then seek the combination of values or theories or principles that

---

<sup>28</sup> Though see Robert E. Goodin, *Utilitarianism as a Public Philosophy* (Cambridge, UK: Cambridge University Press, 1995), who argues that utilitarianism makes most sense in explicitly *public* contexts.

<sup>29</sup> Eva Erman and Niklas Möller (2013), 'Practice-Dependence and Practice-Independence: A False Dichotomy?' (unpublished) provides an instructive discussion of this point.

<sup>30</sup> Peter Lipton, *Inference to the Best Explanation*, 2nd edn. (London: Routledge, 2004).

best explains why we might think so-and-so.<sup>31</sup> We think slavery is wrong, or that we have a human right to bodily integrity, or that no one should die of starvation in a modern constitutional democracy, and we seek the principles or values that best explain such judgments given everything else we believe. In cases like this, the considered convictions are what we have called throughout this essay ‘lower-level’ principles, and the *explanans* ‘higher-level’ principles. While the order of support or justification flows ‘upwards’, the order of explanation flows ‘downwards’. Justification and explanation come apart.

This need not always be the case of course. Many inferences we make in substantive moral and political philosophy will be straightforward deductions (including *mediated deduction!*) from higher-level principles directly to judgments or lower-level principles that we have no independent reason to think are true or appropriate. This will often be the case when we extend some core value or principle to a broader and more controversial set of questions. We might argue, for example, that accepting a higher-level principle of fair reciprocity commits us to a much more stringent egalitarianism among fellow residents of a state than we might have otherwise thought was appropriate. But the point is that a more comprehensive substantive moral and political philosophy will also contain many abductive inferences of just the kind discussed in the previous paragraph. Although it would be fool-hardy to try to argue the point here, it seems clear that the vast majority of inferences in any major work of substantive theorizing are of precisely this abductive kind. And, indeed, it also seems clear that much of our confidence in what I have called ‘higher-level’ principles in fact stems from our much more secure confidence in the contextually grounded, ‘lower-level’ principles that are explained by them. A good example is the higher-level principle of reciprocity captured in (12). If anything supports the higher-level principle, it is the fact that so many human practices that we feel

---

<sup>31</sup> See Norman Daniels, *Justice and Justification* (Cambridge: Cambridge University Press, 1996) on the role of inference to the best explanation in reflective equilibrium.

deeply committed to are shot through and through with particular and context-dependent structures of reciprocal exchange and recognition. Given the ubiquity of such practices, we infer the existence of some version of the higher-level, external, fundamental principle of reciprocity captured in (12): it is this principle that *best explains* our commitment to the reciprocal exchanges that are part of the practices. If I am right, this is yet another way in which the justification and formulation of fundamental moral and political principles can depend on practices.

In the next section, we trace the methodological implications of these arguments for the role of practices in the formulation and justification of principles. *Social interpretation*, once popular among communitarians like Walzer and Taylor, is due for a non-communitarian, quasi-Dworkinian resurgence.

## V. Interpretation

Judgment is needed when the rules run out. Not all conclusions (or intermediate steps on the way to a conclusion) can themselves be established deductively, and no rule can be self-applying (or have determinate, higher-order rules for its application). In the schemata discussed above, judgment is needed in assessing feasible options for their utility-quotient in *instrumental* and in moving from the higher-level principle of reciprocity to its contextually particular counterpart in *mediated deduction* (hence the ‘mediation’). Even if there are no determinate rules for judging correctly, there is *better* and *worse* judgment. Exercising our judgment is not the same as choosing from a set of options where the choice of one is as good as the choice of any other; judgment is meant to be non-arbitrary. Assuming that there is a true or otherwise more appropriate answer, better judgment is more *accurate* judgment. So while we cannot devise rules for getting the right answer, we can produce *methods* or *methodologies* that are more likely to produce such an answer than an untutored guess. These methodologies will

vary, of course, according to the types of inferences in which they will figure. In this section, I shall argue that, given some instance of *mediated deduction*, *social interpretation* is required to ‘get from’ higher-level, external principles to the lower-level internal principles appropriate for a practice. If I am right about how ubiquitous *mediated deduction* is in substantive moral and political philosophy, then this discussion should set a kind of program. In making explicit a method that is already implicit, I hope to highlight the significance of interpretation, and so to bring it back onto center stage. As we develop this argument, we will be able to give a fuller characterization of the way in which practices play a role than we have been able to do thus far.

What is interpretation? There is no such thing as interpretation *simpliciter*. All interpretation is relative to a point of view. Any interpretation requires some (implicit or explicit) understanding of what the interpretation is ultimately *for*. Suppose we are aiming to provide an interpretation of a painting. Why do we need the interpretation? Do we need it in order to inform the organization of an exhibition such as a retrospective? Do we need it as a part of a painter’s biography? Will the interpretation serve in a psychoanalytic study of the painter’s personality? Is it part of an anthropological study of the broader culture in which it figures? Do we want to understand the meaning of the painting as it was intended by its author? Or its meaning in a broader sense, as revealing a perspective on the world and our place in it (whatever the intentions of its author)? Most of the time the point of taking up the interpretive perspective is merely unreflectively assumed; there is rarely a need to reflect on it. The background context in which we search for an interpretation will structure our responses without our being aware of it. But once we become conscious of the range of questions we might ask, it becomes clear that, in each of these cases, what we ‘see’ in the painting will differ, sometimes quite dramatically. Depending on our interpretive perspective, certain properties,

materials, and aspects of the context of production will stand out, others will recede into the background. There is no such thing as interpretation *independent of a particular perspective*.

Knowing why we take up the interpretive perspective will point us in a direction and orient us, but it is not enough. If we reflect carefully on the process involved, we will notice the need for a *further*, higher-level layer of interpretation (which, again, in most cases we may remain unaware of). Return to our previous example: We take up the interpretive perspective because we are aiming ‘to write a biography’, or ‘compose a psychoanalytic study’, or ‘plan a retrospective’, and so on. But how do each of those aims constrain our interpretation? To answer this question, we need some further, more general understanding of *biography* as a genre, including its characteristic modes and styles and history, or an understanding of the special role of *retrospectives* as exhibitions of a particular kind, or a *psychoanalytic theory* within which to make sense of the artwork as part of a reconstruction of a painter’s personality. We then use that higher-level, general interpretation to generate the lower-level, particular interpretation. The higher-level interpretation aids us in picking out relevant features and discarding irrelevant ones; it helps us to organize the jumble of innumerable intrinsic and contextual properties, relations, and materials that *make up* the painting into a coherent narrative *about* the painting. Without that higher-level interpretation, the lower-level interpretation can’t ever begin.

In substantive moral and political philosophy, we take up the interpretive perspective with respect to a practice or set of practices, or an action, or set of relations. Our aim is explicitly *evaluative*. As discussed above, our aim in this paper is evaluation from a *moral* point of view, which requires setting out to examine whether we can justify some action, or practice, or set of relations to those who would like to change or end it. If we can, then we can proceed with the action, relation, or practice. If we cannot, then we must cease. But just as we needed a

characterization of, say, biography or psychoanalysis, etc., we need a characterization of the moral *values* (such as, say, justice) that we will use to structure our evaluation. Is the practice just, or legitimate, or fair, etc.?<sup>32</sup> The selection of such values will, just as in the case of esthetic interpretation, structure what properties of the practice will become relevant in our evaluation, and which can be disregarded.

This two-level structure is a general feature of all interpretation. Consider the interpretation of text or work of art in terms of an *emotion*. Gauguin's *Ondine I*, we might say, represents the tragedy of longing for something unattainable, or the destructive but overwhelming pull of desire, or fear of the unknown.<sup>32</sup> The painting is seen as falling under an emotion (or complex of emotions), as making sense in light of the emotion. The emotion functions as a kind of interpretive key in which to 'read' the painting. But what is an emotion, and how can it function as an interpretive *key*? Following Ronald de Sousa, we can understand emotions as kinds of narrative structured by what he calls 'paradigm scenarios'. Paradigm scenarios are the scenarios through which we understand what it is to feel an emotion (longing, love, fear, anger, and so on).

We are made familiar with the vocabulary of emotion by association with paradigm scenarios. These are drawn first from our daily life as small children and later reinforced by the stories, art, and culture to which we are exposed. Later still, in literate cultures, they are supplemented and refined by literature. Paradigm scenarios involve two aspects: first, a situation ..., and second, a set of characteristic or "normal" to type responses the situation, where normality is first a biological matter and then very quickly becomes a cultural one.<sup>33</sup>

Our repertoire of emotions is not just built up from adaptive agglomerations of primitive biological stimulus/response mechanisms. Formed in our past through our upbringing, refined and elaborated by the cultures and social contexts through which we navigate, the paradigm scenarios have an irreducibly social and developmental dimension. But they also have a *temporal* dimension, which is important for understanding the way in which such scenarios are

---

<sup>32</sup> [http://2.bp.blogspot.com/\\_Nuta\\_CQvImI/TOuswabORgI/AAAAAAAAEWE/zle\\_2A4-xjY/s1600/Paul-Gauguin-In-the-Waves-1889.jpg](http://2.bp.blogspot.com/_Nuta_CQvImI/TOuswabORgI/AAAAAAAAEWE/zle_2A4-xjY/s1600/Paul-Gauguin-In-the-Waves-1889.jpg)

<sup>33</sup> Ronald De Sousa, *The Rationality of Emotion* (Cambridge, MA: MIT Press, 1990), p. 184.

like narratives. Scenarios have a beginning, middle and an end.<sup>34</sup> They have a characteristic evolution through and relationship to time, a dominant mood, rhythm and tempo. Love, for example, is not *love* if it is fleeting and momentary (but must love never ‘alter when it alteration finds’?). We employ the scenarios to interpret the significance and meaning of events, relations, interactions. They train our attention on certain aspects of a particular event or person or relation, suggest connections among different elements, produce an overarching framework. As de Sousa puts it, ‘when a paradigm scenario suggests itself as an interpretation of a current situation, it arranges or rearranges our perceptual, cognitive, and inferential dispositions’.<sup>35</sup>

The higher-level moral values deployed in the evaluation of practices, relations, actions have the same structure as such paradigm scenarios. We understand moral concepts like ‘justice’, ‘injustice’, ‘cruelty’, ‘reciprocity’, ‘solidarity’, and so on, not in terms of some particular instance of each concept, but in terms of their general properties, which will often involve temporally extended scenarios of just the kind at stake in our understanding of emotions. Such general properties are not, however, simply constructed by inductive generalization from particulars, or by seeking a set of necessary and sufficient conditions. They are better understood as constituted by characteristics shared by *paradigms* (hence the appropriateness of the name): the emotion or genre or value is best represented as a situation type united by a series of family resemblances.<sup>36</sup> The important point for our purposes is that the general interpretation of the paradigm scenario is not derived directly from the particular object we are interested in (though it may be subsequently shaped by it). In the moral and political case, we say that the higher-level moral value cannot be internal to the practice. Indeed, it must not be.

---

<sup>34</sup> Velleman makes interesting use of this idea. See J. David Velleman, ‘Narrative Explanation’, *The Philosophical Review* 112 (2003): 1-25.

<sup>35</sup> De Sousa, *The Rationality of Emotion*, p. 186.

<sup>36</sup> Cf. debates in cognitive science about how categorical judgments underlying the use of concepts are made. See, e.g., E. Rosch, ‘Principles of Categorization’, in *Concepts: Core Readings*, eds. E. Margolis and S. Laurence (Cambridge, MA: MIT Press, 1999), pp. 189-206.

For *Ondine I* to be a representation of ‘longing’, it must be the case that our idea of longing has been shaped by situations, experiences, relations, texts, and so on, other than those relating to *Ondine I*. The same thing goes with a value like reciprocity, or justice, and its realization in a specific practice. That does not exclude that our general paradigm scenario of ‘longing’ may be decisively transformed from seeing *Ondine I*, just like our paradigm scenario of reciprocity may be transformed by reflection on, say, its role in a state. It also does not exclude that our experience of *Ondine I* could be decisively transformed by our characterization of longing, just as our reflection on the state could be decisively transformed by our characterization of reciprocity. But that possibility, of course, presupposes the point made here: we must have had a pre-existing understanding of ‘longing’ or ‘reciprocity’ for the transformation to make any sense.

This confirms the EPT from within an account of interpretation: there is no way to provide a defense or characterization of a principle or value *internal* to a practice without assuming its basis is *external* to the practice. It also gives us more structure in explaining how we ‘get from’ some higher-level value to its application within a practice, how, that is, we generate claims like (12) above. So far, we know that the first step in interpreting an object (or, in our case, evaluating a practice) is to specify some interpretive aim and some corresponding paradigm scenario to make sense of it. In the esthetic example, we first look at the painting and form a conjecture about the emotion (or set of emotions) that might best explain its appeal and meaning (given our aim, namely internal reflection). We say things like ‘This is clearly a painting about longing not sorrow’. In the moral and political case, we first look at, say, the state (or the EU or marriage), and form a conjecture about the moral values that might best help us to evaluate it (given the aim, for example, of evaluating its justice). This is equivalent to the step in which we decide what we want an interpretation of a painting for. We then need to

characterize justice in some way that will be relevant and useful for our evaluation. We might give a higher-level interpretation of justice, for example, as a kind of reciprocity, or respect, or fairness, or as set of duties that protects some further value, like autonomy. Already here, just as in the case of the painting, we will be guided by the general properties of the object we are concerned with to help us select among the very different aspects of justice as a value (e.g., it is not justice as a virtue we are concerned with; the state is an institution that comprehensively regulates social life; it is coercive, etc.). ‘Reciprocity’, ‘respect’, ‘autonomy’ and so on, are, like ‘longing’, our paradigm scenarios. We have some sense of what each of them are ‘like’ that we draw from our past, from our experience of each of them in different contexts, and from paradigmatic instances.

It is at this point that we must make a crucial decision, with important consequences for how we will proceed further. The crucial decision is whether our higher-level characterization of justice ought to *remain* external to the practice, and so treat the practice as an instrument for achieving a higher-level aim that can be specified independently of it, or whether we need to enter *into* the practice in order to specify what the principle requires. The former would lead us to a version of *instrumental*; the latter would require *mediated deduction*. For example, we might be committed already at the first stage to a nonrelational conception of global distributive egalitarianism. From this perspective, the state (as apart of a state *system*) can only be justified to the extent that it promotes a globally egalitarian distributive scheme. The internal, contingent properties of states are otherwise irrelevant. Things are very different if we are committed to a relational conception of distributive justice that requires further specification. While there are relational principles that would not require such specification—such as, for example, the

principle ‘distribute equally iff there is mutual assurance of general compliance’<sup>37</sup>—most relational principles do. Take our running example of reciprocity. According to the method we are reconstructing, we would now need to move *internal* to the practice in order to see what justice, now understood as embodying a kind of reciprocity, requires given the kinds of practice-mediated relations created by states.

How do we ‘enter’ a practice in the relevant sense? We do so by providing an interpretation of it from the perspective of our paradigm scenario, which, as I mentioned above, helps us in isolating the relevant features of the practices we are evaluating, and assemble them into a plausible overall view of, say, *longing-as-it-is-depicted-in-Ondine-I or reciprocity-for-long-term-residents-of-a-state*. We note, to continue our analogy, the vertical tower of water, the angle of woman’s body as she tumbles into the water, the anguish of her expression as she swallows her hand; we see each of these elements as organized by a particular conception of longing, that unites and connects them into a comprehensible whole. When we do so, we think of the painting as bearing properties that are *intentionally* connected into some kind of unity. Whether or not we seek some grounding in what *Gauguin* intended—there is no requirement to do so from within the aim given by ‘internal reflection’—we still think of the painting as an intentional object, as having a *point and purpose*.<sup>38</sup> It is not, we say, simply a random collection of strokes on a canvas, as if it had been composed by a fortuitous wind blowing paint on an abandoned canvas. Its point and purpose, furthermore, need not be single; it could be complex (as surely it is). But its complexity will still contribute to its meaning *as a whole*. It is much the same with the state. In the case of reciprocity, our interpretation will be

---

<sup>37</sup> To apply this principle, all one needs to know is whether the practice in question satisfies the condition or not. I cannot argue this here, but I believe that all such principles are backed by higher-level relational principles that *do* require *mediated deduction*. In the example mentioned, there will be some higher-level principle that specifies why distributive justice *generally* requires mutual assurance, and some explanation for what it requires *specifically* in the context of the mutual assurance secured by states (but not clubs, etc.), which, in this example, is egalitarianism.

<sup>38</sup> On this point, see Dworkin, *Justice for Hedgehogs*; Ronald Dworkin, *Law’s Empire* (Cambridge: Harvard University Press, 1986).

directed towards the kinds of collective goods produced by the mutual production of a state's residents and citizens. Our interpretation will then give an account not only of *which* collective goods are produced by states but also *how* they contribute to the state's point and purpose. This interpretation need not be constrained by what residents' and citizens' joint intentions are (taken somehow as the actual intentions of a collective agent) any more than we need to be constrained by what Gauguin intended. The process of interpretation, as we have already discussed, is guided instead by our over-arching values, aims, and concerns. Let us say we interpret the state, when seen from the perspective of reciprocity, as providing a central class of collective goods whose main point and purpose is to set up and maintain the essential conditions for human flourishing. We can then go on to provide an interpretation of what a fair return requires among people who secure one another such essential conditions, just as we specified how longing is represented in *Ondine I*. And so we move from (10) to (12) in our schema.

With our contextually fixed principles in hand, we can now become *instrumentalists*. We now know what justice as reciprocity demands for citizens and residents. Armed with this knowledge, we can turn to any token-state, say, Britain, and wonder whether its institutions realize the demands that they create. If they do not, we can say that British law or policy or society is unjust, and seek to change it. British institutions are conceived of as instruments or tools for realizing the internal demands they themselves generate. Notice further that this is where the analogy to esthetic interpretation breaks down. The esthetic interpretation we engaged in with respect to *Ondine I* is not aimed to change the world (or the painting). While we were driven by some sense of the value of the painting (otherwise, it wouldn't be worth reflecting on its meaning), we were not narrowly trying to form principles for promoting or otherwise 'realizing' its goodness. Moral interpretation, of course, is concerned from the very

start with critique and criticism. But, as I hope to have shown, the special normativity of moral values represents no bar to taking their context-sensitivity seriously.

In concluding this section, I want to return to an objection. The objection is a familiar adaptation of the first argument we encountered, and wonders: What if a practice *in toto* must be dismantled or rejected or overthrown? Wouldn't the context-sensitive interpretation I have advocated make such rejection impossible? No. There are two ways such rejection can happen. First, we might decide that only a purely external perspective is possible with respect to some practices (e.g., torture).<sup>39</sup> There is nothing in the method or account of *mediated deduction* I have championed that precludes this possibility. Whether or not we ought to take such a purely external perspective depends on the substantive issues at stake, not on questions of methodology.<sup>40</sup> Second, we might engage in what we briefly referred to above as ideology or genealogical critique. We might take an internal value or principle at the center of the self-understanding of participants in some practice, and then aim to show that the value or principle, properly understood, cannot be realized by the practice in this or any other form; in order to realize the value, one must abandon the practice. Alternatively, and more radically, one might show that the practice actually *undermines* any basis for the value's realization. An example might be the idea of *wage slavery*. Here the ideology critic shows that the practice of capitalism undermines the very possibility of the freedom that it pretends to champion. Genealogical critique, similarly, serves to unmask the pretense that a practice contributes to the realization of some value internal to it by exploring the history by which the value became central to the practice. A paradigmatic example of such genealogy is Nietzsche's genealogy of

---

<sup>39</sup> It is important to remember, however, that such an external perspective need not be nonrelational. It could be that there is a wider practice (e.g., the state) such that an internal relational principle (e.g., (12)-type reciprocity) precludes the narrower practice (e.g., torture).

<sup>40</sup> On this point, see Saladin Meckled-Garcia, "The Practice-Dependence Red Herring and Better Reasons for Restricting the Scope of Justice" (2013).

morality, which demonstrates how the practices central to the modern moral system actually work to sustain astounding levels of cruelty and promote the deadening of man's creative potential.<sup>41</sup> The key point for our purposes is that ideology critique and genealogy require the same forms interpretation that I have elucidated above. The interpretation of the practice is a crucial step in showing how it renders the realization of the external value impossible. Ideology critique and genealogy are therefore best understood, as I mentioned above, as negative instances of *mediated deduction*.

A schematic summary of the interpretive method outlined here may be useful:

- (15) Identify a target social relation or practice or object. (E.g.: the state, or *Ondine I*)
- (16) Why take up the interpretive perspective with respect to that social relation or practice or object? (E.g.: ... to evaluate the state in terms of distributive justice; ... to find a meaning that can provide depth and richness to inner reflection)
- (17) What paradigm scenarios best characterize the interpretive aim? (E.g.: distributive justice in the state is best understood in relational terms as a species of reciprocity; the inner-reflection meaning of *Ondine I* is best understood in terms of longing)
- (18) At this stage, one provides a general conception of justice as reciprocity, defending it, for example, against rival fully relational and nonrelational views, or a general conception of longing, defending it against other candidate paradigm scenarios for inner reflection.
- (19) Turn back to the target social relation or practice or object. What is the best understanding of the social relation or practice or object in terms of the paradigm scenario(s) identified in (3)?
- (20) At this stage, one uses the general idea of justice as a form of reciprocity to make sense of the relation between citizens and residents of a state as mutual producers of collective goods, or the paradigm scenario of longing to make sense of the specific elements of *Ondine I*—the vertical tower of water, the angle of her body as she tumbles into the water, the anguish of her expression as she swallows her hand. Here one works out the specific form reciprocity takes in the state or the form longing takes in the painting. The result will be a conception of reciprocity or longing much more detailed and textured than the general conception with which one started; but it will also be, for that very reason, not (or not readily) generalizable. *Reciprocity-in-the-state* (or *longing-in-Ondine I*) is very different than *reciprocity-in-friendship* (or *longing-in-Courbet's-Woman-in-the-Waves*), and very different again from the paradigm scenario of reciprocity or longing.

---

<sup>41</sup> On Nietzsche and genealogy, see, e.g., David Owen, 'Nietzsche, Re - Evaluation and the Turn to Genealogy', *European Journal of Philosophy* 11 (2003): 249-7; Leiter, *Nietzsche on Morality*; Raymond Geuss, *Public Goods, Private Goods*, 2nd edn. (Princeton, NJ: Princeton University Press, 2003), Introduction.

## V. Conclusion

In this paper, I have argued that claim (1) (the EPT) is true and in any case innocuous. Claim (2) is merely definitional, and can easily be granted. Claim (3) (the FPT) is trivial. Claim (4) is false and claim (5), as a result, severely misleading.

In the process, I have identified four different ways in which fundamental principles can depend on practices. First, they can constrain the *domain* of a principle (e.g., relativism). I left relativism to the side in this paper. Second, they can have applicability conditions, and so the *scope* of a principle can depend on a practice or set of social relations. Third, and most controversially, the *grounds* of a principle can depend (in part) on the contingent properties of some set of practices. Here I claimed that, while it is true that ultimately all principles must bottom out in a fundamental and external principle, there are many fundamental, external principles (especially relational ones) that require context-sensitive interpretation in order to deliver determinate guidance. Fourth, I claimed that it is often our internal, practice-based commitments to values and principles that justify an appeal to the higher level principles that best explain them, rather than the other way around.

Once we had all this in view, we were in a position to see that the justification and formulation of fundamental principles in fact plays a small and largely insignificant role in what I called *substantive* moral and political philosophy. In the special sense of ‘fundamental’ that we have used here (adopted from Cohen), such philosophy spends most of the time *applying* fundamental principles to more specific contexts. The different ways in which one might apply higher-level principles deserves, I then argued, much more attention than it is has thus far received. Such diversity in modes of application has been overlooked largely as a result of claims like (4) and (5). I then went on to distinguish two patterns of application, namely *instrumental* and *mediated deduction*, and contended that it is rare to see (with the possible exception of Singer’s utilitarianism) theories which begin with a higher-level principle and then

proceed via an *instrumental* application of them to particulars. Most substantive philosophy proceeds rather, by way of *mediated deduction*, in which practices play an essential (if unavowed) role in moving from very abstract, higher-level principles to more determinate, lower-level ones. The final section of the paper went on to defend an interpretive method for generating, by *mediated deduction*, lower-level principles from higher-level principles in conjunction with specific practices. My aim there was to demonstrate that, if I am right, then interpretation must play an essential role in any plausible substantive moral and political philosophy. Making this role explicit, I hoped, would help to reframe current debates on both methodology and the role of facts in substantive moral and political philosophy.