

LAWYERS AND JUSTICE

AN ETHICAL STUDY

DAVID LUBAN

PRINCETON UNIVERSITY PRESS
PRINCETON, NEW JERSEY

SPONSORED BY THE INSTITUTE FOR PHILOSOPHY AND PUBLIC POLICY

THE STRUCTURE OF ROLE MORALITY

The parameters of our problem are now clear: we want to give considerations of role morality their proper weight without permitting them simply to override common morality. In this chapter I shall sketch a model of institutional excuses that is able to do that. I shall do so by exhibiting inadequacies in "my station and its duties"—the formal pattern of institutional excuses expounded earlier, which placed policies over acts and erected a structure of exclusionary reasons based on these policies. This pattern is inadequate in several respects, I believe, and in the course of demonstrating this, better patterns will emerge.

The duties associated with a professional role are of three sorts. First, there are those that are essential to the proper functioning of the role. Thus, for example, it is often argued that a lawyer must keep client confidences because otherwise clients would not tell lawyers all the facts needed for their lawyers to represent them. This argument is a functional justification of the duty of confidentiality.

Second, there are "side constraints"—rules that are not essential to the basic activities of the role, but which prevent abuses or further other worthy ends. Examples of side constraints in the Code of Professional Responsibility include rules requiring lawyers to divulge adverse legal authority to the courts or to report ethical violations by other lawyers.¹ Neither of these is functionally essential to the role of an adversary advocate, but both are worthwhile side constraints to impose upon that role.

Finally, there are the customary or accepted practices of the role, those things that every lawyer does only because they are the things that every lawyer does. These may change over time, but at any given time they are part of the role-related activity.²

For the purposes of the present argument, I will confine the notion of role morality to the first of these categories—the duties de-

1. ABA Code, DR 7-106(B)(1) and DR 1-103(A), respectively.

2. These may in turn become duties. For an analysis of the process by which custom becomes legal duty, see FINNIS, pp. 238-45.

rived from a functional analysis of the means necessary to a role's ultimate ends. That is because the pattern of institutional excuses sketched earlier under the heading of "my station and its duties" consists in denying the agent's personal moral responsibility for role acts by appealing to the fact that the role itself is good. It is hard to see what relevance the goodness of the role has for the justification of particular duties unless we are talking about duties that contribute to the fulfillment of the role's ends. Side constraints and customs should enter into the justification of the role as it is currently understood, rather than being themselves justified by the role. Their justification is exogenous to the functioning of the role, and when we ask whether a role is justified, we ask, "given that it operates under these side constraints and with these customs, is the role a good one?"

An institutional excuse, we said, consists of the two-step argument that (1) the role is morally desirable; and (2) in a morally desirable role, the moral responsibility for actions falls on the role (or on the institution that creates it) and not on the role agent.³

Evidently, the idea here is that the institution creating the role passes along its moral cachet to the requirements of the role (just as, in rule-utilitarianism, utility-maximizing rules pass along their moral cachet to the acts they require us to perform). I shall argue, however, that the weaker the justification of the institution, the slighter the moral significance of special institutional duties.⁴

THE FOURFOLD ROOT OF SUFFICIENT REASONING

Let us look at the structure of an institutional excuse when the institution is strongly justified, when it is a positive moral good. We do this to see how justification passes from institution to agent.

Consider, as an example, a philanthropic organization such as Ox-

3. I have changed the wording slightly from the original formulation to account for the objection that an immoral role does not excuse its occupant.

4. The same is true, I believe, in the case of utilitarianism. Rule-utilitarianism tells us that if a rule is justified (no matter how marginally) we must perform the acts it requires. On my view, this cannot be right: the question of whether the rule is strongly or weakly justified must affect its ability to require acts that considered in isolation are undesirable on utilitarian grounds. If a rule is warranted only by what I have called a "pragmatic justification," that is, only because it is not significantly worse than its (almost equally unattractive) alternatives, it cannot be right for utilitarians to perform acts with very low utility just because the rule requires them. If, on the other hand, the rule has a very powerful utilitarian justification, it may well justify such acts.

fam, whose sole function is to distribute food to famine-stricken people in impoverished areas of the world. Let this be the institution. The division of labor within it creates different jobs or institutional tasks: the roles. Each of these has specified duties or role obligations. These might be quite general: the logistics officer, for example, might have as her role obligation procuring the means of transporting food. (The role obligations may also be quite specific.) To carry out her role obligations, she must perform various actions; we'll call them role acts.

Let us suppose that to get food to a remote village in an underdeveloped country stricken by famine, the logistics officer must obtain several trucks from a powerful local boss named P. P, as it happens, is involved in a number of unsavory activities, including a plan to murder a local man because P wants to sleep with his wife. Imagine further that the logistics officer overhears P dispatching a murderer to kill the man that very night; that P discovers that the logistics officer has overheard him; and that P tells her that if the man is warned and escapes, P will withhold the trucks.

The officer is in a moral dilemma. Other things being equal, she is under a moral obligation to warn the victim or the police. Let us, at any rate, suppose that this is so. But here, if anywhere, we may wish to permit an institutional excuse. Suppose the officer complies with P's demand. Asked to justify this, she says, "My job is more important."

The structure of this institutional excuse may be spelled out as follows: she points out that the role act of complying with P is required by her role obligation, which in turn is necessary to perform the role's institutional task, which (finally) is justified by the positive moral good of the institution: the saving of many innocent lives. Taken together, these justifications for the role act outweigh the obligation to warn P's unfortunate victim.

It may be objected that this formulation misrepresents the logistics officer's dilemma. Even if she did not occupy a role, she would face the dilemma of choosing between saving the life of P's victim and saving the lives of the famine victims. Her excuse, then, is not "My job is more important than saving one life," but "Saving many lives is more important than saving one life." The moral dilemma now appears to be a straightforward consequentialist problem of whether to save one life or many: it has nothing to do with the appeal to a professional role. So goes the objection.

In reality, however, it is the objection that misrepresents the officer's dilemma. Strictly speaking, the logistics officer has no choice

between saving one life or saving many, because, strictly speaking, *she* cannot save many. If she had superhuman powers, perhaps she could: she could singlehandedly raise money, buy food, carry it to another continent, make arrangements with the local authorities, transport it, distribute it, and so forth. But, because her powers are merely mortal, only many people working together can save the lives of the starving populace.

That, of course, is why the institution and her role within it exist. The objection treats her role as an accidental or contingent feature occasioning the officer's dilemma of saving one and saving many; but this is incorrect, for without the role (and without superhuman powers) she could not conceivably be in this dilemma. Far from being a contingent feature of a more general problem, the fact that she is in the role is necessary for the problem even to be stated, to be intelligible. Correctly speaking, the dilemma is not "to save one life or many?" but "to save one life or to do my job (which is a necessary but not a sufficient condition for saving many lives)?" The dilemma is precisely a choice between a role obligation and a common moral obligation.⁵

The general problem, which creates the dilemma, is that the positions, "The institution is a morally good one" and "The institution imposes role obligations on its officers, some of which may mandate morally bad role acts" can both be true. In fact, the problem is more complicated than that because the structure of arguments I have just described contains two terms between institutions and role acts (namely role and role obligations), and the general problem that moral good can authorize moral wrong is replicated at all the stages. Good institutions can require unsavory roles; good roles can have unwelcome obligations attached to them; and honorable obligations can require awful actions.

In any such case, the institutional excuse, fully spelled out, will take the form I have indicated: the agent (1) justifies the institution by demonstrating its moral goodness; (2) justifies the role by appealing to the structure of the institution; (3) justifies the role obligations by showing that they are essential to the role; and (4) justifies the role act by showing that the obligations require it.

Already we have modified our original pattern of institutional excuses: we need four separate arguments, rather than just one that shows that the institution is morally good. Let us call this pattern—

5. This argument was added as a result of a discussion with Mary Gibson and Judith Lichtenberg.

with apologies to Arthur Schopenhauer—the “Fourfold Root of Sufficient Reasoning.”

We may represent the fourfold root argument schematically as a chain consisting of four links. Link one consists of the institution and its justification. Link two is a role and its derivation from the institution’s requirements. Link three is a role obligation and its derivation from the role’s requirements. And link four is a role act and the demonstration that it is an instance of the role obligation. The Fourfold Root of Sufficient Reasoning is an argument that begins with link one and proceeds through the chain to link four.

Here is a schematic diagram of the fourfold root argument, as applied to our Oxfam example, as well as to the lawyer’s duty of zeal:

TABLE 1.

THE FOURFOLD ROOT OF SUFFICIENT REASONING	
Link 1:	<i>Institutions</i> (OXFAM; The Adversary System)
Link 2:	<i>Roles</i> (Logistics Officer; Advocate)
Link 3:	<i>Role Obligations</i> (Procuring Trucks; The Duty of Zeal)
Link 4:	<i>Role Acts</i> (Obeying P.; Various Role Acts of Lawyering)
Each link is justified by appeal to the one preceding it; Link 1 is justified by direct argument	

I am not, let me emphasize, claiming that an institutional excuse is inevitably appropriate when the fourfold root is satisfied. You and I may differ in our assessment even of the Oxfam example. I am only claiming that in such a case, a difficult moral dilemma exists, from which an institutional excuse is one possible way out. Whether the institutional excuse succeeds will be determined by balancing its soundness and significance against the common moral obligation from which it is excusing the agent.

In any event, the fourfold root argument may be undermined at any of its links. The role may not be essential to the institution (the logistics officer is only a back-up: someone else has already procured trucks); the obligation may not be an essential part of the role (rather than a logistics officer, the visitor is a dietician merely passing through the village on his way to the famine area, and thus needs no favors from P); or the act may not be required to fulfill the obligation (the logistics officer can warn the endangered man or the police without P knowing).

Most important for our purposes, however, is to see that the first

step may be undermined—that even though the institution may be a justifiable one, its justification is too weak to support the fourfold root’s conclusion.

Consider this example. Rather than an Oxfam logistics officer, the visitor who overhears P dispatch the murderer is an anthropology graduate student, a member of a team doing field research on village culture. Suppose as well that the research will be ruined if the anthropologists interfere with social relations in the “subject population”—participant observers are strictly ruled out, and deviation invalidates the results. This is a reason for the student not to warn the murder victim; in addition, of course, she does not want P to interfere with her research.

Here the institution in question is the scientific research team. We may suppose that it is a worthy one: the attempt to understand other cultures and (in general) to seek knowledge is morally admirable. I suppose as well that its restrictions on interfering with village affairs are at least arguably necessary for the job to be done (you can read the arguments in the team’s grant proposal).

Yet here I think we would agree that all this is not enough reason to refrain from warning the target of P’s designs. Once again, of course, we can question any of the four links in the institutional excuse. Do her role obligations require her to permit the murder (i.e., is there any way she can prevent it and still do the research)? How strong is the argument that only this methodology will yield meaningful scientific results? How crucial are the results in this village to the study as a whole? And, finally, how important is the study in the scheme of things?

In this example, the last is the main sticking point. Pure knowledge is great stuff, there’s no denying it. But in general we strike the balance differently here—we expect Dr. Frankenstein to call the police rather than quietly watch a murder outside his house, even if doing so will cost science a badly needed esophagus. The sides of the anthropologist’s dilemma—the role act as justified by the fourfold foot versus the common moral obligation—do not have equal weight and the institutional excuse collapses. The question, then, is not just whether the institution is morally justified, but how strongly it is justified.

TWO PATTERNS OF INSTITUTIONAL EXCUSE

This question already raises serious problems for “my station and its duties.” It will be easiest to see these by asking another question:

what combined effects do weak links in the fourfold root have on the strength of the whole argument? Suppose in particular that each link in the fourfold root is justified, but only weakly; does this diminish the strength of the fourth and final link?

At each link *X* of the fourfold root (except the final one) we may ask both (a) What justifies *X*? (Is *X* justified?) and (b) What does *X* justify? The exclusionary-reasons approach of "my station and its duties" assumes that the answers to these questions are logically independent of each other, so that after we have an answer to (a) showing that *X* is justified, we can simply take *X* as a given, as part of the "perimeter," and proceed to answer (b). This, in effect, is a *minimum-threshold test* for each link: once *X* is over the threshold, it no longer matters whether its justification was strong or weak—*X* is justified, pure and simple.

This is the key to an exclusionary-reasons analysis of institutional excuses: once the institution is justified, further inquiry into how it was justified is excluded. We must regard the justification as a given and ask only what it in turn justifies. The arguments pertaining to the justification of each link "forget" those pertaining to previous links.

Our discussion of the two examples suggests, however, that this is wrong: a weak link weakens the entire chain and several weak links together weaken it more than any single link does on its own. If I may employ a mathematical metaphor, we should use a *cumulative-weight test* rather than a minimum-threshold test. That is, the total force of the fourfold root must depend on the *product* of the forces of the several steps: each step of the argument establishing an institutional excuse weights the ensuing steps. If the institution is justified, but only weakly, that fact weights the justification that roles within it can provide. If, in turn, a role is justified, but only weakly—if, that is, it is useful but not essential for the functioning of the institution—that fact weights the justification that its role obligations can provide. The two weak justifications together seriously undermine the moral force of appeals to these role obligations; and so on, down the line.

The cumulative-weight approach, like the minimum-threshold approach, distinguishes question (a), "What justifies *X*?" from question (b), "What does *X* justify?" But it no longer regards them as independent of each other: the answer to (b) will depend on the kind of answer given to (a) and not just on whether (a) had a positive answer. Only in this way can we account for the rather obvious idea that an act absolutely essential for an institution of merely marginal importance itself has only marginal importance.

It is nevertheless true that a weak justification is a justification. If an institution is over the threshold of justification, even by a millimeter, it has reason to exist; if a role within it is justified, weakly or not, it too has reason to exist. That is the element of truth in a minimum-threshold test. It is only when we get to the final question, the question of whether the agent should perform a role act that is contrary to common morality, that the cumulative weight of justification becomes relevant. Thus, the fourfold root justifies roles, role obligations, and role acts by a minimum-threshold test—asking the justificatory question for each link is conditional on the previous link having passed a minimum-threshold test—but at the same time, we keep a "running total" of justificatory strength given by a cumulative-weight test, which comes into play when we finally wonder what to do.

The quasi-mathematical language here is of course only a metaphor for a process of deliberation that is not quantitative, because the factors involved do not admit of measurement. Even talk of "weighing" or "balancing" factors, talk to which lawyers no less than decision theorists are addicted, is little more than pickle-smoke, suggesting precision where none is to be had.

We talk about "weighing" factors when we deliberate, despite the fact that we cannot assign numerical weights to them. We even apply the metaphor of weights to factors that are incommensurable. The metaphor is a placeholder or stand-in for processes of valuation, deliberation, and judgment that are difficult even to name, let alone to describe, and that is why we use the metaphor. The problem with its use lies in its suggestion that these hard-to-think-about deliberative processes amount to a kind of computation or cost-benefit accounting.

It is this suggestion that is absurdly misleading. A question like those we are exploring here, e.g., "Which is more important, gaining anthropological knowledge or saving one life?" may actually be nonsensical, just as the question "Which is better, no-fault auto insurance or Beethoven's Second Symphony?" is nonsensical. It will certainly seem so to someone trying to answer it in practice, particularly if one tries to answer it by taking the metaphor of "weighing the value of the terms" literally. (What on earth would it mean to take it literally?)

We nevertheless do make choices and trade-offs between apples and oranges, often with a fair degree of confidence that they are the right choices. It is this simple fact that makes talk of "weighing" and "balancing" more than nonsense. And the decision theorist's model of deliberation—assigning numerical values to the factors in a deci-

sion and then performing calculations—may therefore not be as silly as it appears, provided that we never forget that it is only a metaphor. Deliberating-by-number bears precisely the same relation to genuine practical reasoning that painting-by-number bears to art.

Bearing these warnings in mind, I would like to stay with the mathematical metaphor a bit longer, in order to indicate more clearly what a cumulative-weight test is. Demonstrating, in the tidy (if fictitious) world of arithmetic, that institutional excuses can be derived on a very different pattern from that provided by the exclusionary-reasons analysis will help us see that the derivation must be different in real life.

Suppose, then, that the various pro and con factors that bear on the justification of a link in the fourfold root could be given numerical values. Suppose also that the strength of justification rises and falls as the difference between benefits and costs, or credits and debits, rises and falls.

To form the cumulative weight of justification, the credits at each link must be discounted by the weakness of the previous link, but the debits must not. A debit is a reason for believing that a role is not essential to an institution, or a role obligation to a role, or a role act to an obligation. Such a reason is always an objection to performing the role act, regardless of whether the institution, role, or role obligation are strongly or weakly justified; if anything, the objection is more strenuous when the justification of other links is weaker.

This point, to put it in other words, means that if an institution is weakly justified, that fact may weaken arguments that one's role in it is crucial, but it cannot weaken arguments that the role is useless, on the institution's own terms. Thus, the cumulative-weight test may be stated as follows: A cumulative-weight test weights the credits, but not the debits, in each link by the justificatory strength of the preceding link.⁶

Dear reader, let us now leave the arithmetic to those who are made happy by it and return to the topic of institutional excuses. So far we have modified "my station and its duties" by substituting the Fourfold Root of Sufficient Reasoning and the cumulative-weight test for the original two-step form of institutional excuse and the minimum-threshold test.

6. It is a simple matter to state this symbolically, but the exercise does not really have much point; there is such a thing as being too literal-minded about a mathematical metaphor, and surely elaborating a formula that no one will ever use is a prime instance of this.

These are improvements over "my station and its duties"; but the model we have been developing is still insufficient.

Notice that although general assessments of institutions, roles, and role obligations are crucial to the fourfold root inquiry, ultimately they matter only as policies underlying role acts. This is clearly true in "my station and its duties," because the justifications of institutions, roles, and role obligations are "forgotten," and only the bare fact that they were successful is remembered.

But it is also true on the approach I have defended. For here, too, the arguments themselves are forgotten and only their strength is remembered. Moreover, even the strengths of the arguments at each link are forgotten—the cumulative weight of all of them, not the weights of each, determine their significance for the role act. All of the components of our deliberation are collapsed into one weighting that bears equally on each role act required by the role obligation or rule.

What this leaves out is the all-important fact that even if two role acts are both required by a rule, the policies underlying that rule may bear on those role acts with different relevances.

We touched on this point in chapter 3. The simple example, please recall, was a rule barring those under sixteen years old from driving a car. The concern underlying the rule is obvious: to keep drivers off the road until they are old enough to handle a car prudently. But of course this concern is more applicable to some under-sixteen-year-olds than others. The rule applies to all under-sixteen-year-olds, but its justification does not.

This is equally so when the rule in question is a role obligation and the application is a role act: the rule might apply to two cases, even though its justification does not. To take this into account, we must ask not only about the weight of the various justifications in the fourfold root, but also about their relevance. When we do that, we break decisively with the mathematical metaphor of cumulatively weighing arguments; we now must stop calculating and start thinking. And this is the final modification in the theory of institutional excuses that we have been developing.

HOW OUR ANALYSIS DIFFERS FROM "MY STATION AND ITS DUTIES"

What has become of the distinction between acts and policies employed in our last chapter's discussion of institutional excuses? A partial answer is that it is still present. The level of greatest particularity

in the fourfold root structure is that of role acts, and entities at the other three levels—institutions, roles, and role obligations—are all evaluated by arguments that generalize beyond this particular case, arguments that might well be called “policy arguments.”

The primary novelty in this approach is that it rejects the question “Which is the logical subject of evaluation: the act or the rule (the policy)?” as posing a false dichotomy. Unlike a policies over acts approach—rule-utilitarianism, for example—the fourfold root does not eschew the moral assessment of individual acts because it concludes by weighing the arguments favoring the performance of this particular role act against the objections to this particular role act arising from common morality. Unlike an acts over policies approach—act-utilitarianism, for example—the fourfold root takes the moral assessment of general rules, policies, and institutions very seriously, for all these must be brought to bear on the question of whether a morally dissonant role act might nevertheless be required of us.

Another way to look at the relationship between acts and policies is to notice that at least two perspectives are possible. The role agent takes the institution, the role, and the role obligations as givens. A lawyer may think that the legal system, the practice of law, and her own professional ethics are morally corrupt, but she cannot very easily change them. She can, of course, stop being a lawyer, but then she is no longer a role agent. Thus for her, the only possible subject of moral evaluation is the particular act, and policy arguments come into play only as moral arguments about what act she should perform. The extent to which she finds her role corrupt will manifest itself in the range of cases in which she is willing to break role, and in the size of her liquor bill.

But we can also adopt the point of view of the role designer or legislator. For the legislator, policy arguments about a certain institution, role, or role obligation bear directly on whether that institution, role, or obligation should be reformed, rather than on how a role occupant should weigh the requirements of the role against the demand to acknowledge other people morally. The latter question is considered by the legislator only in the thought that a role that routinely requires morally disquieting behavior of its agents should be reformed.

The agent can’t change the role; she can only decide whether or not to perform it.⁷ The legislator can change the role, but stands on

the sidelines while it is being performed. This suggests that the question of the “logical subject of moral evaluation” must be answered conditionally: whichever entity you have the power to affect becomes the logical subject of (your) evaluation. Both the role agent and the legislator will consider the whole fourfold root structure; but they will draw different consequences from it: the role agent will decide what to do, while the legislator will decide how to change the conditions under which agents make such decisions.

And this proviso—that both legislator and role agent must consider all the arguments in the fourfold root, differing only in the consequences they draw from them—shows that we have dispensed with the exclusionary-reasons analysis altogether. For the role agent must now consider moral arguments beyond the perimeters set by the role, while the legislator must heed the effects roles have on agents’ capacity to acknowledge the moral status of other people in particular cases.

On an exclusionary-reasons approach, the fact that my role is justified gives me a reason not to reconsider its credits and debits whenever I deliberate. But the current approach insists that these reasons be considered, both as to their relevance and as to their importance. In fact, the current approach and an exclusionary-reasons analysis run in opposite directions. On the exclusionary-reasons approach, arguments justifying a role may be “cashed in” for authorizations to treat the role as a given, a presupposition, out of further deliberation. The current approach, on the other hand, treats the fact that a role is justified as a mere placeholder or abbreviation for a more complex claim that certain policy arguments are both relevant and sufficiently strong. In other words, it “cashes out” the role (and role obligations and institution) into a set of policy arguments.

IS IT TOO MUCH TO ASK?

Let us recapitulate the analysis of institutional excuses developed in the preceding pages and answer some objections.

The analysis begins, of course, with a moral dilemma: a dilemma between a common moral consideration and an act apparently required by one’s role that conflicts with it. Ultimately the dilemma will be resolved one way or another, of course: the agent’s final decision will be to pick which of the two conflicting moral demands is more

are in effect redesigning the role. This cannot amount to much unless many agents do so at once.

7. This is not to deny that agents, by routinely breaking role in certain circumstances,

important. All the rest is simply an explanation of how the agent assesses the moral demand that she perform the role act. And the process of assessment goes like this:

1. Identifying the institution, the role, the role obligation and the role act.
2. Assessing the institution, role, and role obligation in the light of the ends they are to serve.
3. Applying the minimum-threshold test: determining whether, at each link, the credits and debits indicate that the entity (institution, role, role obligation, role act) is justified.
4. Applying the cumulative-weight test: determining the total significance of the various policy arguments to the role act.
5. Assessing the relevance of the policy arguments to the case at hand.
6. Resolving the dilemma by weighing the justification of the role act against the moral offense of performing it.
7. Acting.

This seems like a lot to ask; it all looks like a veritable night in Gethsemane every time a lawyer wants to take what might be described in some circles as "a lousy four-bill landlord-tenant case."

Now there is one obvious way in which this objection simply misses the point. The Fourfold Root of Sufficient Reasoning is a theoretical account of moral justification, not a recipe for real-time deliberation. Even if an agent is unable to press the fourfold root into service when she is confronted with an actual moral problem, it still serves a valuable function. It tells us that institutional excuses are not inevitably appropriate even in relation to good institutions; it tells us that appealing to my station and its duties does not absolve me from moral criticism; and (I will argue in the next chapter) it tells us that the standard conception of the lawyer's role is untenable.

The fact that the fourfold root is fundamentally a theory of justification and only secondarily a theory of deliberation makes it an important tool for the analysis of rules of professional obligation. The justification for rules such as those of zeal and confidentiality will consist of a progression along the first three links of the fourfold root structure—institutions, roles, and role obligations in the program I just presented—and this inquiry can be undertaken without waiting for a concrete moral dilemma. In chapter 8 I shall follow this strategy to argue that the principles of partisanship and nonaccountability must be modified; in chapters 9 and 10, I will argue in this fashion that the present rules of confidentiality are inadequate.

Using the fourfold root structure in this way to assess the strength

of professional obligations abstracted from concrete moral dilemmas can also make real-time deliberation less strenuous, and thereby disarm the objection that the program demands too much. The point is that the most cerebral parts of the deliberation, the abstract evaluation of the institution, role, and role obligations contained in the first three links of the fourfold root, need not be done over and over again. The purposes and justifications of your occupational role and the institutions in terms of which it exists are familiar: you must come to grips with them some time in your career, but they do not have to be reinvented each time you confront a dilemma. They don't even have to be thought through by you personally, except to assess their relevance to the case at hand: you can read about what justifies confidentiality, or the adversary system, in a book like this one. The only deliberative acts that must be performed on the spot are those needed to bring these rather abstract assessments to bear on the dilemma at hand. And this may take mere seconds to think through—we are talking, after all, of deliberation, not a cost-benefit analysis by Coopers & Lybrand.

Although the fourfold root is basically a structure of justification and not of deliberation, it is not bad as a structure of deliberation either. What looks to be hopelessly complex is not so bad in real life. If you try to describe the mental operations involved in doing your week's grocery shopping ("Enumerate who will be at dinner and how many dinners you must cook; plan the menus; break them down into ingredients; add the amounts of each ingredient across menus; etc.") it seems astonishing that you eat at all; obviously, you do manage somehow.

Perhaps, however, the objection is not that this pattern of moral deliberation is too complicated, but that the demand to deliberate comes too frequently. Asking us to deliberate every single time that a role obligation comes into conflict with common morality—even if it's just a lousy four-bill landlord-tenant case—seems to screw the springs of moralism impossibly tight, and translates all of morality into the deadly realm of the supererogatory.⁸ To this I have three things to say.

First—and this is obviously the heart of the matter—your "lousy four-bill landlord-tenant case" may result in an innocent person being evicted on a technicality; it is facts like this that cause the moral dilemma to exist in the first place. To dismiss the moral problem is

8. These worries are developed in Wolff: see especially pp. 38–40. The phrase "the deadly realm of the supererogatory" comes from Judith Lichtenberg.

not merely wrong, it also treats the lawyer who faces it with unwarranted contempt, as though his or her occupation could not possibly raise an issue worth thinking about.

Second—and this pertains particularly to lawyers—the sort of thinking about institutions and general practices required by this pattern of deliberation is just what is done for professional reasons every time there is a complicated brief to write. It does not seem too much to ask that what one does in order to live with one's client might also be done in order to live with one's self, even in the hard cases where you have to think about it.

Third, it is important to remind ourselves of what the alternative is to raising moral questions when we are faced with the prospect of doing something that may be wrong. The alternative is not raising them.

I would like to consider a related and important objection to the kind of moral thinking I am recommending. All the talk of "weighing" various "factors" against each other can seem deceptively anti-septic. Those "factors" include commitments to the duties of a profession, to a career, or to major social institutions such as adversary justice. These can be, they frequently are, among the deepest loyalties and commitments in our lives; and it cannot be right to ask us to reconsider them, to trade them off, again and again. A person who was willing to do this, we may think, is morally frightening, not commendable.

Bernard Williams has argued in this way about our "ground" projects—those to which we are most profoundly committed. Williams's argument (too complex to be presented here) is roughly this: an agent who conceived of his or her projects as chronically open for reconsideration would have no integrity, and integrity is not merely a virtue; its possession is a precondition of human agency (integrity is not far from psychological identity). The fact that we have "ground" projects is tantamount to the fact that we do not seriously contemplate suicide—they are, so to speak, the autonomic functions of the moral life, and to include them as just some more factors in deliberation, which may be neglected or abandoned at times, is like trying to put the circulation of your blood under voluntary control.⁹

This argument can be read in two contradictory ways. On the strong reading, a person who regards even the things that are dearest to her, that make life worth living, as always open to question

9. WILLIAMS (1), pp. 100–18; WILLIAMS (4), pp. 11–14.

and in principle compromiseable cannot exist, or act, or be a person.¹⁰ On the weak reading, such a person can exist, but she is morally frightening. She has weakened her moral defenses against doing horrible things, because such defenses need a category of the "unthinkable"—one is reminded of David Rousset's aphorism, "Normal men do *not* know that everything is possible." A person willing to deliberate on anything is one who "knows" that everything is possible.¹¹

The first of these interpretations of Williams's argument says that a person cannot deliberate in the way I have suggested, while the second says that a person can, but should not. And each of these in turn has a strong and a weak version: on the strong version, a person cannot (should not) ever deliberate in this way, while on the weak version, a person cannot (or should not) routinely deliberate in this way.

The strong versions are false. A person can and sometimes should rethink and abandon even a basic commitment, when it is leading her into an untenable life—when, for example, it is forcing her to compromise other basic commitments and in particular a basic commitment to morality. Many Western Communists, for example, found themselves compelled to abandon their party when the Hitler-Stalin pact was signed. Such an experience is bound to be shattering; it will be experienced by the agent as a betrayal of herself by her project, and it may create a sharp break in her life, in effect rupturing her personal identity. (This may explain why many Communists did not simply abandon Moscow but became rabid anti-Communists.) Nevertheless, such rethinking is obviously possible, and few would deny that at times it is desirable and necessary.

More troubling is the weak version of the argument, that one can not (or should not) routinely engage in such crisis thinking. This version of the argument seems to call into question the possibility and worth of the model of deliberation I am advocating.

Ultimately, however, the model emerges unscathed because it does not in fact require the agent to, so to speak, put it all on the line whenever it is applied. It would require this if a negative conclusion to the fourfold root argument meant that the agent's role should henceforth cease to exist, or (less drastically) that the agent should abandon the role. But in fact, a negative conclusion requires only that the agent break role on this occasion, which must surely be possible if we are to abandon the totalitarian conception of role involved in

10. This is my interpretation of WILLIAMS (1), pp. 116–17, and WILLIAMS (4), p. 14.

11. And this comes from WILLIAMS (1), pp. 92–93.

"my station and its duties." One might even say that adopting this pattern of deliberation gives the agent a more livable way of engaging in the fundamental project embodied in the role. That, after all, is its purpose: because it gives the institutional values their proper weight, it is a pattern of thought for role agents who take these values dead seriously.

Perhaps, however, the worry is that because the agent is to entertain in a somewhat dispassionate way the limited importance of institutional commitments, the pattern is incompatible with seeing those commitments as "ground projects." But I do not believe this is so: there is no real-life incompatibility between a project's being fundamentally important to me and my appreciating that its absolute importance is limited. I may find music a thing that makes life worth living and commit my life, my fortune, and my sacred honor to hearing, playing, and studying music. This does not prevent me from appreciating that music-making is of strictly limited importance, and I will join the bucket-brigade to stop Rome from burning even if I must give up a once-in-a-lifetime opportunity to hear Peter Nero. I believe, therefore, that one may accept Williams's argument without deluding oneself about the worth of one's projects or enslaving oneself to the commitments they involve. The practice of law may be the most important thing in a lawyer's life (that this is so may account for the very high divorce rate in the legal profession) but that need not be accompanied by overvaluing the role of lawyer. If I am right, taking the role seriously means in fact taking seriously the moral limits of the role, and not giving in to the temptation toward rote application of its institutionally imposed duties.

THE DIVISION OF LABOR AND THE MORALITY OF ACKNOWLEDGMENT

Let me conclude by scrutinizing once again some of the differences between deliberation on the pattern I have recommended and "my station and its duties," which may give rise to further objections. There are two such points of comparison: the division of moral labor, and the morality of acknowledgment.

The division of moral labor is retained in the current approach in a weakened form. It is still agreed that reasons exist to parcel out different duties to different stations. These reasons, in fact, are precisely the reasons offered in the fourfold root to justify the institution, the role, and the role obligations. What has changed, however, is the absolute claim of the division of moral labor over the agent's undi-

vided moral judgment. The claims of the role may be overridden if they are irrelevant to the case or insufficiently strong.

It might be objected that on this approach the claims of the role are honored in name only—that, in other words, the tension between role morality and common morality is now resolved in favor of common morality in virtually all cases. For in any case in which common morality diverges from role morality, it will turn out that the moral arguments for the role fail in respect of relevance—otherwise, common morality would not diverge from the moral basis of the role. Thus, it may appear that the deck is now stacked in such a way that role morality can never prevail.

But this is not so. Sometimes the policies underlying a role can be both forceful and relevant, even though applying them works what by the lights of common morality is a substantive wrong. For example, there are powerful reasons for a lawyer defending a guilty criminal-accused to try her damndest for an acquittal: these are simply the reasons we have discussed for making sure that no accused person faces a hostile state and tribunal without legal counsel, and they create the role of the defense lawyer with its adamant obligation of zeal. Nevertheless, if the lawyer succeeds in winning an acquittal, an injustice is done.

The criminal defense lawyer is one of the clearest cases of a role occupant who will often find that the justifications of the role are so crucial that they override all but the most stringent demands of common morality. Military roles can also provide such examples; so can certain political roles (a topic to which I shall return in chapter 14). These, the most clear-cut examples, graphically illustrate that the current approach takes role morality seriously.

We next ask what all this has to do with the morality of acknowledgment. After all, we embarked on this desperate and tedious safari through the theory of practical reason only in order to find an alternative to "my station and its duties," an alternative that is necessary because "my station and its duties" blocks the acknowledgment of the other.

Posing the question in this way makes it, I hope, almost self-answering. The morality of acknowledgment as I described it earlier requires that we be able to respond to others, even in violation of the duties of our station. We must, that is, acknowledge that their moral status exercises claims on us that do not always coincide with the claims of our role. The pattern of deliberation I have been describing accords with this requirement because the moral force of our role is strictly limited; because the particularities of the case count; and

above all, because the status of other people *qua* people is registered in the common morality that figures into our deliberations. The need to break role, in a particular case, for the sake of a particular person, is accommodated by the pattern of deliberation I have been describing.

A moment ago we considered the objection that this pattern is in fact biased against role morality and toward common morality in situations where they clash. I argued that this is not so, because when forceful policies relevant to the case at hand back the role obligation, the agent may demur to it, even though a substantive immorality is worked. At this point, however, one may offer the opposite objection: if any cases exist in which the agent demurs to requirements of a social role even though doing so ill-uses another person, then the agent is not acknowledging the other. The objection, in other words, is that our model is inadequate to the morality of acknowledgment precisely because it is not biased in favor of common morality.

Back and forth we go. The objection, apparently, does not accept that a social role exerts any moral force at all when we are faced with the need to respond to another person. This ignores several facts, however. First of all, the role itself is packed with moral content in the form of the policies underlying it. This fact cannot simply be dismissed, for doing so confuses acknowledging the other with doing whatever the other requires, even when that itself violates common morality. Responding to the other's moral worth cannot mean treating the other as being of infinite worth. We acknowledge other people's worth as little by exaggerating it as by dismissing it. No man is an island, to be sure, but none of us is the world either, and to treat someone as the world is to fail in acknowledging everyone else.

Moreover, as I suggested earlier, roles also play a constructive part in acknowledgment. They give content to acknowledgment by creating social channels through which respect for persons is conveyed. If I am in a slow restaurant and about to miss my movie, the waiter acknowledges my predicament by quietly asking the chef to prepare my food quickly and then bringing everything promptly. In absolute terms, of course, that isn't much, but precisely because of the constraints on the waiter's role, it is a thoughtful way to act that acknowledges the importance of my evening's project. This is so even if the waiter fails and I miss my movie. By contrast, the waiter does not acknowledge my personhood by running to the theater and pleading with the projectionist to delay starting the movie, even if that is the only way I will have time to eat. For then the waiter is not treating me as a person. He is treating me as an ass.

Social roles, by constraining behavior and creating expectations, actually create new possibilities of human acknowledgment based on the exercise of tact, delicacy, discretion, and acts that become expressive of concern precisely because they are performed from within tight constraints. Breaking role, on the other hand, may at times lead one to violate respectful distance from the other—it may be condescending or overbearing. And, of course, in those emergencies where it would be obscene to act in role, where the other's situation is truly desperate or miserable, the sort of deliberation I have recommended allows the role agent to acknowledge that fact, because the strength of the justification will be overridden by the emergency.