

I. Taking Ourselves Seriously
II. Getting It Right

HARRY FRANKFURT

THE TANNER LECTURES ON HUMAN VALUES

Delivered at

Stanford University
April 14–16, 2004

HARRY G. FRANKFURT is Professor Emeritus of Philosophy at Princeton University. He was educated at Johns Hopkins University, where he received his Ph.D. He has held positions at Ohio State University, the State University of New York at Binghamton, Rockefeller University, and at Yale. He was also a visiting fellow at All Souls College, Oxford. He is a fellow of the American Academy of Arts and Sciences. He has written more than 50 scholarly articles, essays, and reviews and is the author of three books: *Demons, Dreamers and Madmen: The Defense of Reason in Descartes' Meditations* (1970); *The Importance of What We Care About* (1988); and *Necessity, Volition and Love* (1999); and the editor of *Leibniz: A Collection of Critical Essays* (1972).

I. TAKING OURSELVES SERIOUSLY

I

I suppose some of you must have noticed that human beings have a tendency to be heavily preoccupied with thinking about themselves. Blind rollicking spontaneity is not exactly the hallmark of our species. We put very considerable effort into trying to get clear about what we are really like, trying to figure out what we are actually up to, and trying to decide whether anything can be done about this. The strong likelihood is that no other animal worries about such matters. Indeed, we humans seem to be the only things around that are even *capable* of taking themselves seriously.

Two features of our nature are centrally implicated in this: our rationality and our ability to love. Reason and love play critical roles in determining what we think and how we are moved to conduct ourselves. They provide us with decisive motivations, and also with rigorous constraints, in our careers as self-conscious and active creatures. They have a great deal to do, then, with the way we live and with what we are.

We are proud of the human abilities to reason and to love. This makes us prone to rather egregious ceremonies and excursions of self-congratulation when we imagine that we are actually making use of those abilities. We often pretend that we are exercising one or the other—that we are following reason, or that we are acting out of love—when what is truly going on is something else entirely. In any case, each of the two is emblematic of our humanity, and each is generally acknowledged to merit a special deference and respect. Both are chronically problematic, and the relation between them is obscure.

Taking ourselves seriously means that we are not prepared to accept ourselves just as we come. We want our thoughts, our feelings, our choices, and our behavior to make sense. We are not satisfied to think that our ideas are formed haphazardly, or that our actions are driven by transient and opaque impulses or by mindless decisions. We need to direct ourselves—or at any rate to *believe* that we are directing ourselves—in thoughtful conformity to stable and appropriate norms. We want to get things right.

It is reason and love—the directives of our heads and of our hearts—that we expect to equip us most effectively to accomplish this. Our lives are naturally pervaded, therefore, by an anxious concern to recognize what they demand and to appreciate where they lead. Each has, in its own way, a penetrating and resonant bearing upon our basic condition—the condition of persons, attempting to negotiate the environments of their internal as well as of their external worlds.

Sometimes, to be sure, we energetically resist what reason or love dictates. Their commands strike us as too burdensome, or as being in some other way unwelcome. So we recoil from them. Perhaps, finally, we reject them altogether. Even then, however, we ordinarily allow that they do possess a genuine and compelling authority. We understand that what they tell us really does count. Indeed, we have no doubt that it counts a great deal—even if, in the end, we prefer not to listen.

Among my aims in these lectures is to explore the roles of reason and of love in our active lives, to consider the relation between them, and to clarify their unmistakable normative authority. In my judgment, as you will see, the authority of practical reason is less fundamental than that of love. In fact, I believe, its authority is grounded in and derives from the authority of love. Now love is constituted by desires, intentions, commitments, and the like. It is essentially—at least as I construe it—a *volitional* matter. In my view, then, the ultimate source of practical normative authority lies not in reason but in the will.

I hope that you will find my analyses and arguments at least more or less convincing. I also hope, of course, that they will be clear. In this connection, I must confess to being a bit unsettled by a rather mordant piece of advice that comes (I understand) from the quantum physicist Nils Bohr. He is said to have cautioned that one should never *speak* more clearly than one can *think*. That must be right; but it is rather daunting. In any event, here goes.

2

What is it about human beings that makes it possible for us to take ourselves seriously? At bottom it is something more primitive, more fundamental to our humanity, and more inconspicuous than either our capacity for reason or our capacity to love. It is our peculiar knack of separating from the immediate content and flow of our own consciousness and introducing a sort of division within our minds. This elementary maneuver establishes an inward-directed, monitoring oversight. It

puts in place an elementary reflexive structure, which enables us to focus our attention directly upon ourselves.

When we divide our consciousness in this way, we *objectify* to ourselves the ingredient items of our ongoing mental life. It is this self-objectification that is particularly distinctive of human mentality. We are unique (probably) in being able simultaneously to be engaged in whatever is going on in our conscious minds, to detach ourselves from it, and to observe it—as it were—from a distance. We are then in a position to form reflexive or higher-order responses to it. For instance, we may approve of what we notice ourselves feeling, or we may disapprove; we may want to remain the sort of person we observe ourselves to be, or we may want to be different. Our division of ourselves situates us to come up with a variety of supervisory desires, intentions, and interventions that pertain to the several constituents and aspects of our conscious life. This has implications of two radically opposed kinds.

On the one hand, it generates a profound threat to our well-being. The inner division that we introduce impairs our capacity for untroubled spontaneity. This is not merely a matter of spoiling our fun. It exposes us to psychological and spiritual disorders that are nearly impossible to avoid. These are not only painful; they can be seriously disabling. Facing ourselves, in the way that internal separation enables us to do, frequently leaves us chagrined and distressed by what we see, as well as bewildered and insecure concerning who we are. Self-objectification facilitates both an inhibiting uncertainty or ambivalence and a nagging general dissatisfaction with ourselves. Except in their most extreme forms, these disorders are too commonplace to be regarded as pathological. They are so integral to our fundamental experience of ourselves that they serve to define, at least in part, the inescapable human condition.

On the other hand, however, this very capacity to divide and to objectify ourselves provides the foundational structure for several particularly cherished features of our humanity. It accounts for the fact that we possess such a thing as practical reason; it equips us to enjoy a significant freedom in the exercise of our will; and it creates for us the possibility of going beyond simply wanting various things, and of coming instead to care about them, to regard them as important to ourselves, and to love them. The same structural configuration that makes us vulnerable to disturbing and potentially crippling disabilities also immeasurably enhances our lives by offering us—as I will try to explain—opportunities for practical rationality, for freedom of the will, and for love.

3

When we begin attending to our own feelings and desires, to our attitudes and motives, and to our dispositions to behave in certain ways, what we confront is an array of—so to speak—psychic raw material. If we are to amount to more than just biologically qualified members of a certain animal species, we cannot remain passively indifferent to these materials. Developing higher-order attitudes and responses to oneself is fundamental to achieving the status of a responsible person.

To remain wantonly unreflective is the way of nonhuman animals and of small children. They do whatever their impulses move them most insistently to do, without any self-regarding interest in what sort of creature that makes them. They are one-dimensional, without the inner depth and complexity that render higher-order responses to oneself possible. Higher-order responses need not be especially thoughtful, or even entirely overt. However, we become responsible persons—quite possibly on the run and without full awareness—only when we disrupt ourselves from an uncritical immersion in our current primary experience, take a look at what is going on in it, and arrive at some resolution concerning what we think about it or how it makes us feel.

Some philosophers have argued that a person becomes responsible for his own character insofar as he shapes it by voluntary choices and actions that cause him to develop habits of discipline or indulgence and hence that make his character what it is. According to Aristotle, no one can help acting as his virtuous or vicious character requires him to act; but in some measure a person's character is nonetheless voluntary, because "we are ourselves. . . part-causes of our state of character" (*Nicomachean Ethics* III.5, 1114.b22). In other words, we are responsible for what we are to the extent that we have caused ourselves—by our voluntary behavior—to become that way.

I think Aristotle is wrong about this. Becoming responsible for one's character is not essentially a matter of *producing that character* but of *taking responsibility for it*. This happens when a person selectively identifies with certain of his own attitudes and dispositions, whether or not it was he that caused himself to have them. In identifying with them, he incorporates those attitudes and dispositions into himself and makes them his own. What counts is our current effort to define and to manage ourselves, and not the story of how we came to be in the situation with which we are now attempting to cope.

Even if we *did* cause ourselves to have certain inclinations and tendencies, we can decisively rid ourselves of any responsibility for their

continuation by renouncing them and struggling conscientiously to prevent them from affecting our conduct. We will still be responsible, of course, for having brought them about. That cannot be changed. However, we will no longer be responsible for their *ongoing* presence in our psychic history, or for any conduct to which that may lead. After all, if they do persist, and if they succeed in moving us to act, it will now be only against our will.

4

When we consider the psychic raw materials with which nature and circumstance have provided us, we are sometimes more or less content. They may not exactly please us, or make us proud. Nevertheless, we are *willing* for them to represent us. We *accept* them as conveying what we really feel, what we truly desire, what we do indeed think, and so on. They do not arouse in us any determined effort to dissociate ourselves from them. Whether with a welcoming approval, or in weary resignation, we consent to having them and to being influenced by them.

This willing acceptance of attitudes, thoughts, and feelings transforms their status. They are no longer merely items that happen to appear in a certain psychic history. We have taken responsibility for them as authentic *expressions of ourselves*. We do not regard them as disconnected from us, or as alien intruders by which we are helplessly beset. The fact that we have adopted and sanctioned them makes them intentional and legitimate. Their force is now our force. When they move us, we are therefore not *passive*. We are *active*, since we are being moved just by ourselves.

Being identified with the contents of one's own mind is a very elementary arrangement. It is so ubiquitous, so intimately familiar, and so indispensable to our normal experience, that it is not easy to bring it into sharp focus. It is so natural to us, and as a rule it comes about so effortlessly, that we generally do not notice it at all. In very large measure, it is simply the default condition.

5

Of course, the default condition does not always prevail. Sometimes we do not participate actively in what goes on in us. It takes place, somehow, but we are just bystanders to it. There are obsessional thoughts, for instance, that disturb us but that we cannot get out of our heads; there are peculiar reckless impulses that make no sense to us, and upon which we could never think of acting; there are hot surges of anarchic emotion

that assault us from out of nowhere and that have no recognizable warrant from the circumstances in which they erupt.

These are psychic analogues of the seizures and spasmodic movements that occur at times in our bodies. The fact is that we are susceptible to mental tics, twitches, and convulsions, as well as to physical ones. These are things that *happen to us*. When they occur, we are not participating agents, who are expressing what we really think or want or feel. Just as various bodily movements occur without the body being moved by the person whose body it is, so various thoughts, desires, and feelings enter a person's mind without being what that person truly thinks or feels or wants.

Needless to say, however dystonic and disconnected from us these mental events may be, they do occur in *our* minds—just as the analogous physical events occur nowhere else but in our bodies. They are, at least in a gross literal sense, our thoughts, our feelings, and our desires. Moreover, they often provide important indications of what else is going on in our minds. Uncontrollably spasmodic movements of the limbs are likely to be symptomatic of some deeper and otherwise hidden physical condition. Similarly, the fact that I have an obsessional thought that the sun is about to explode, or a wild impulse to jump out the window, may reveal something very significant about what is going on in my unconscious. Still, that is not what I really think about the sun; nor does the impulse to jump express something that I really want to do.

6

What a person finds in himself may not just seem oddly disconnected from him. It may be dangerously antithetical to his intentions and to his conception of himself. Some of the psychic raw material that we confront may be so objectionable to us that we cannot permit it to determine our attitudes or our behavior. We cannot help having that dark side. However, we are resolved to keep it from producing any direct effect upon the design and conduct of our lives.

These unacceptable intruders arouse within us, then, an anxious disposition to resist. By a kind of psychic immune response—which may be mobilized without our even being aware of it—we push them away, and we introduce barriers of repression and inhibition between them and ourselves. That is, we dissociate ourselves from them and seek to prevent them from being at all effective. Instead of incorporating them, we *externalize* them.

This means that we deny them any entitlement to supply us with motives or with reasons. They are outlawed and disenfranchised. We refuse to recognize them as grounds for deciding what to think or what to do. Regardless of how insistent they may be, we assign their claims no place whatever in the order of preferences and priorities that we establish for our deliberate choices and acts. The fact that we continue to be powerfully moved by them gives them no rational claim. Even if an externalized desire turns out to be irresistible, its dominion is merely that of a tyrant. It has, for us, no legitimate authority.

Some philosophers maintain that, just by virtue of having a desire, a person *necessarily* has a reason for trying to satisfy it. The reason may not be a very strong one; there may be much better reasons to perform another action instead. Nevertheless, it counts for *something*. The very fact that a person wants to do something always means, on this view, that there is at least that much of a reason in favor of his doing it.

However, the mere fact that a person has a desire does not give him a reason. What it gives him is a *problem*. He has the problem of whether to identify with the desire and thus validate it as eligible for satisfaction, or whether to dissociate himself from it, treat it as categorically unacceptable, and try to suppress it or rid himself of it entirely. If he identifies with the desire, he acknowledges that satisfying it is to be assigned *some* position—however inferior—in the order of his preferences and priorities. If he externalizes the desire, he determines to give it no position in that order at all.

7

Reflexivity and identification have fundamental roles in the constitution of practical reason. Indeed, it is only by virtue of these elementary maneuvers that we *have* such a thing as practical reason. Without their intervention, we could not regard any fact as giving us a reason for performing any action.

When does a fact give us a reason for performing an action? It does so when it suggests that performing the action would help us reach one or another of our goals. For example, the fact that it is raining gives me a reason for carrying an umbrella insofar as doing that would be helpful as a means to my goal of keeping dry.

Having a goal is not the same, however, as simply being moved by a desire. Suppose I have a desire to kill someone, and that firing my pistol at him would be an effective way to accomplish this. Does that mean I

have a reason to fire my pistol at him? In fact, I have a reason for doing that *only* if killing the man is not just an outcome for which a desire happens to be occurring in me. The desire must be one that I accept and with which I identify. The outcome must be one that I really want.

Suppose that the man in question is my beloved son, that our relationship has always been a source of joy for me, and that my desire to kill him has no evident connection to anything that has been going on. The desire is wildly exogenous; it comes entirely out of the blue. No doubt it signifies God knows what unconscious fantasy, which is ordinarily safely repressed. In any case, it instantly arouses in me a massive and whole-hearted revulsion. I do whatever I can to distance myself from it, and to block any likelihood that it will lead me to act.

The murderous inclination is certainly real. I do have that lethal desire. However, it is not true that I want to kill my son. I don't really want to kill him. Therefore, I don't have any reason to fire my pistol at him. It would be preposterous to insist that I do have at least a weak reason to shoot him—a reason upon which I refrain from acting only because it is overridden by much stronger reasons for wanting him to remain alive. The fact that shooting him is likely to kill him gives me *no reason at all* to shoot him, even though it is true that I have a desire to kill him and that shooting him might do the trick. Since the desire is one with which I do not identify, my having it does not mean that killing my son is actually among my goals.

8

Practical reasoning is, in part, a procedure through which we determine what we have most reason to do in order to reach our goals. There could be no deliberative exercise of practical reason if we were related to our desires only in the one-dimensional way that animals of nonreflective species are related to whatever inner experience they have. Like them, we would be mutely immersed in whatever impulses happen at the moment to be moving us; and we would act upon whichever of those impulses happened to be most intense. We would be no more able than they are to decide what we have reason to do because, like them, we would be unable to construe anything as being for us an end or a goal.

In fact, without reflexivity we could not make decisions at all. To make a decision is to make up one's mind. This is an inherently reflexive act, which the mind performs upon itself. Subhuman animals cannot perform it because they cannot divide their consciousness. Since they cannot take themselves apart, they cannot put their minds back

together. If we lacked our distinctive reflexive and volitional capacities, making decisions would be impossible for us too.

That would not alter the fact that, like all animals in some degree, we would be capable of behaving *intelligently*. Being intelligent and being rational are not the same. When I attempt to swat an insect, the insect generally flies or scurries rapidly away to a place that is more difficult for me to reach. This behavior reduces the likelihood that it will die. The insect's self-preservative movements are not structured in detail by *instinct*. They are not inflexibly *modular* or *tropistic*. They are continuously adjusted to be effective in the particular, and often rapidly changing, circumstances at hand. In other words, the insect—although it does not deliberate or reason—behaves intelligently. Even if we too were unable to reason or to deliberate, we too would nevertheless often still be able—by appropriately adaptive adjustments in our behavior—to find our way intelligently to the satisfaction of our desires.

9

Let us suppose that a certain motive has been rejected as unacceptable. Our attempt to immunize ourselves against it may not work. The resistance we mobilize may be insufficient. The externalized impulse or desire may succeed, by its sheer power, in defeating us and forcing its way. In that case, the outlaw imposes itself upon us without authority, and against our will. This suggests a useful way of understanding what it is for a person's will to be free.

When we are doing exactly what we want to do, we are acting freely. A free act is one that a person performs simply because he wants to perform it. Enjoying freedom of action consists in maintaining this harmonious accord between what we do and what we want to do.

Now sometimes, similarly, the desire that motivates a person as he acts is precisely the desire by which he wants to be motivated. For instance, he wants to act from feelings of warmth and generosity; and in fact he *is* warm and generous in what he does. There is a straightforward parallel here between a free action and a free will. Just as we act freely when what we do is what we want to do, so we will freely when what we want is what we want to want—that is, when the will behind what we do is exactly the will by which we want our action to be moved. A person's will is free, on this account, when there is in him a certain volitional unanimity. The desire that governs him as he is acting is in agreement with a higher-order volition concerning what he wants to be his governing desire.

Of course, there are bound to be occasions when the desire that motivates us when we act is a desire by which we do *not* want to be motivated. Instead of being moved by the warm and generous feelings that he would prefer to express, a person's conduct may be driven by a harsh envy, of which he disapproves but which he has been unable to prevent from gaining control. On occasions like that, the will is not free.

But suppose that we are doing what we want to do, that our motivating first-order desire to perform the action is exactly the desire by which we want our action to be motivated, and that there is no conflict in us between this motive and any desire at any higher order. In other words, suppose we are thoroughly wholehearted both in what we are doing and in what we want. Then there is no respect in which we are being violated or defeated or coerced. Neither our desires nor the conduct to which they lead are imposed upon us without our consent or against our will. We are acting just as we want, and our motives are just what we want them to be. Then so far as I can see, we have on that occasion all the freedom for which finite creatures can reasonably hope. Indeed, I believe that we have as much freedom as it is possible for us even to conceive.

IO

Notice that this has nothing to do with whether our actions, our desires, or our choices are causally determined. The widespread conviction among thoughtful people that there is a radical opposition between free will and determinism is, on this account, a red herring. The possibility that everything is necessitated by antecedent causes does not threaten our freedom. What it threatens is our power. Insofar as we are governed by causal forces, we are not omnipotent. That has no bearing, however, upon whether we can be free.

As finite creatures, we are unavoidably subject to forces other than our own. What we do is, at least in part, the outcome of causes that stretch back indefinitely into the past. This means that we cannot design our lives from scratch, entirely unconstrained by any antecedent and external conditions. However, there is no reason why a sequence of causes, outside our control and indifferent to our interests and wishes, might not happen to lead to the harmonious volitional structure in which the free will of a person consists. That same structural unanimity might also conceivably be an outcome of equally blind chance. Whether causal determinism is true or whether it is false, then, the wills of at least some of us may at least sometimes be free. In fact, this freedom is clearly not at all uncommon.

11

In the Scholium to Proposition 52 in Part IV of his *Ethics*, Baruch Spinoza declares that “the highest good we can hope for” is what he refers to as “*acquiescentia in se ipso*.” Various translators render this Latin phrase into English as “self-contentment,” as “self-esteem,” or as “satisfaction with oneself.” These translations are a little misleading. The good to which Spinoza refers is certainly not to be confused with the contentment or pride or satisfaction that people sometimes award themselves because of what they think they have accomplished, or because of the talents or other personal gifts with which they believe they are endowed. It is not Spinoza’s view that the highest good for which we can hope has to do either with successful achievement or with vanity or pride.

There is something to be said for a bluntly literal construction of his Latin. That would have Spinoza mean that the highest good consists in *acquiescence to oneself*—that is, in acquiescence to being the person that one is, perhaps not enthusiastically but nonetheless with a willing acceptance of the motives and dispositions by which one is moved in what one does. This would amount to an inner harmony that comes to much the same thing as having a free will. It would bring with it the natural satisfaction—or the contentment or self-esteem—of being just the kind of person one wants to be.

Unquestionably, it is a very good thing to be in this sense contented with oneself. Spinoza does not say that it is the *best* thing one can hope for; he doesn’t say even that it is enough to make life good. After all, it may be accompanied by terrible suffering, disappointment, and failure. So why say, as he does, that it is the highest thing for which one can hope?

Perhaps because it resolves the *deepest* problem. In our transition beyond naive animality, we separate from ourselves and disrupt our original unreflective spontaneity. This puts us at risk for varieties of inner fragmentation, dissonance, and disorder. Accepting ourselves reestablishes the wholeness that was undermined by our elementary constitutive maneuvers of division and distancing. When we are acquiescent to ourselves, or willing freely, there is no conflict within the structure of our motivations and desires. We have successfully negotiated our distinctively human complexity. The unity of our self has been restored.

12

The volitional unity in which freedom of the will consists is purely structural. The fact that a person’s desire is freely willed implies nothing

as to *what* is desired or as to whether the person actually cares in the least about it. In an idle moment, we may have an idle inclination to flick away a crumb; and we may be quite willing to be moved by that desire. Nonetheless, we recognize that flicking the crumb would be an altogether inconsequential act. We want to perform it, but performing it is of no importance to us. We really don't care about it at all.

What this means is not that we assign it a very low priority. To regard it as truly of *no* importance to us is to be willing to give up having any interest in it whatever. We have no desire, in other words, to *continue* wanting to flick away the crumb. It would be all the same to us if we completely ceased wanting to do that. When we do care about something, we go beyond wanting it. We want to *go on* wanting it, at least until the goal has been reached. Thus, we feel it as a lapse on our part if we neglect the desire, and we are disposed to take steps to refresh the desire if it should tend to fade. The caring entails, in other words, a commitment to the desire.

Willing freely means that the self is at that time harmoniously integrated. There is, within it, a *synchronic* coherence. Caring about something implies a *diachronic* coherence, which integrates the self across time. Like free will, then, caring has an important structural bearing upon the character of our lives. By our caring, we maintain various thematic continuities in our volitions. We engage ourselves in guiding the course of our desires. If we cared about nothing, we would play no active role in designing the successive configurations of our will.

The fact that there are things that we do care about is plainly more basic to us—more constitutive of our essential nature—than what those things are. Nevertheless, *what* we care about—that is, what we consider important to ourselves—is obviously critical to the particular course and to the particular quality of our lives. This naturally leads people who take themselves seriously to wonder how to get it right. It leads them to confront fundamental issues of normativity. How are we to determine what, if anything, we *should* care about? What *makes* something genuinely important to us?

13

Some things are important to us only because we care about them. Who wins the American League batting title this year is important to me if I am the kind of baseball fan who cares about that sort of thing, but probably not otherwise. My close friends are especially important to me; but

if I did not actually care about those individuals, they would be no more important to me than anyone else.

Of course, many things are important to people even though they do not actually care about them. Vitamins were important to the ancient Greeks, who could not have cared about them since they had no idea that there were such things. Vitamins are, however, indispensable to health; and the Greeks did care about that. What people do not care about may nonetheless be quite important to them, obviously, because of its value as a means to something that they do in fact care about.

In my view, it is *only* in virtue of what we actually care about that anything is important to us.¹ The world is everywhere infused for us with importance; many things are important to us. That is because there are many things that we care about just for themselves, and many that stand in pertinent instrumental relationships to those things. If there were nothing that we cared about—if our response to the world were utterly and uniformly flat—there would be no reason for us to care about anything.

14

Does this mean that it is all simply up to us—that what is truly important to us depends just upon what goes on in our minds? Surely there are certain things that are *inherently* and *objectively* important and worth caring about, and other things that are not. Regardless of what our own desires or attitudes or other mental states may happen to be, surely there are some things that we *should* care about, and others that we certainly should *not* care about. Is it not unmistakably apparent that people should at least care about adhering to the requirements of morality, by which all of us are inescapably bound no matter what our individual inclinations or preferences may be?

Some philosophers believe that the authority of morality is as austere independent of personal contingencies as is the authority of logic. Indeed, their view is that moral principles are grounded in the same fundamental rationality as logically necessary truths. For instance, one advocate of this moral rationalism says: “Just as there are rational requirements on thought, there are rational requirements on action”; and since “the requirements of ethics are rational requirements. . . , the

1. I will not discuss whether this needs to be modified to refer also to what we *would* care about if we were properly acquainted with it. In any case, the modification could readily be absorbed into the voluntaristic account of practical normativity that I am developing.

motive for submitting to them must be one which it would be contrary to reason to ignore.”² On this account, failure to submit to the moral law is irrational. The authority of the moral law is the authority of reason itself.

The normative authority of reason, however, cannot be what accounts for the normative authority of morality. There must be some other explanation of why we should be moral. For one thing, our response to immoral conduct is very different from our response to errors in reasoning. Contradicting oneself or reasoning fallaciously is not, as such, a *moral* lapse. People who behave immorally incur a distinctive kind of opprobrium, which is quite unlike the normal attitude toward those who reason poorly. Our response to sinners is not the same as our response to fools.

Moreover, if it *were* possible for people to justify their conduct strictly by reason (i.e., with rigorous proofs demonstrating that acting otherwise would be irrational) that would provide no advantage to morality. In fact, it would render the claims of morality far less compelling, because it would take people off the hook. After all, being convinced by proofs does not implicate any of a person’s individual preferences or predilections. Reason necessitates assent and leaves no room for individual choice. It is entirely impersonal. It does not reveal character.

Construing the basis of morality rationalistically misses the whole point of moral norms. Morality is essentially designed to put people *on* the hook. Whether or not a person adheres to the moral law is not supposed to be independent of the kind of person he is. It is presumed to reveal something about him deeper and more intimate than his cognitive acuity. Moral principles cannot rest, therefore, simply upon rational requirements. There must be something behind the authority of the moral law besides reason.

15

Let us assume, then, that moral authority cannot be satisfactorily established by invoking just the bloodless support of strict rationality. Is there not a sufficient basis of some other kind for recognizing that moral requirements (and perhaps normative requirements of various other types as well) are genuinely important in themselves, regardless of any-

2. Thomas Nagel, *The Possibility of Altruism* (Oxford: Oxford University Press, 1970), p. 3.

one's beliefs or feelings or inclinations? In my judgment, there is not. There can be no rationally warranted criteria for establishing anything as *inherently* important.

Here is one way to see why. Nothing is important if everything would be exactly the same with it as without it. Things are important only if they make a difference. However, the fact that they do make a difference is not enough to show that they are important. Some differences are too trivial. A thing is important only if it makes an *important* difference. Thus, we cannot know whether something is important until we already know how to tell whether the difference it makes is important.

The unlimited regress to which this leads is clearly unacceptable. If it were possible for attributions of inherent importance to be rationally grounded, they would have to be grounded in something besides other attributions of inherent importance. The truth is, I believe, that it is possible to ground judgments of importance only in judgments concerning what people care about. Nothing is truly important to a person unless it makes a difference that he actually cares about. Importance is never inherent. It is always dependent upon the attitudes and dispositions of the individual. Unless a person knows what he *already* cares about, therefore, he cannot determine what he has reason to care about.

The most fundamental question for anyone to raise concerning importance cannot be the *normative* question of what he *should* care about. That question can be answered only on the basis of a prior answer to a question that is not normative at all, but straightforwardly *factual*: namely, the question of what he actually does care about.³ If he attempts to suspend all of his convictions, and to adopt a stance that is conscientiously neutral and uncommitted, he cannot even begin to inquire methodically into what it would be reasonable for him to care about. No one can pull himself up by his own bootstraps.

16

What we care about has to do with our particular interests and inclinations. If what we *should* care about depends upon what we *do* care about, any answer to the normative question must be derived from considerations that are manifestly subjective. This may make it appear that what we should care about is indeed up to us, and that it is therefore likely to vary from one person to another and to be unstable over time.

3. If the modification mentioned in note 1 above is adopted, the pertinent question (concerning what the person *would* care about) will still be straightforwardly factual.

Answers to the normative question are certainly up to us in the sense that they depend upon what we care about. However, what we care about is not always up to us. Our will is not invariably subject to our will. We cannot have, simply for the asking, whatever will we want. There are some things that we cannot *help* caring about. Our caring about them consists of desires and dispositions that are not under our immediate voluntary control. We are committed in ways that we cannot directly affect. Our volitional character does not change just because we wish it to change, or because we resolve that it do so. Insofar as answers to the normative question depend upon carings that we cannot alter at will, what we should care about is not up to us at all.

Among the things that we cannot help caring about are the things that we love. Love is not a voluntary matter. It may at times be possible to contrive arrangements that make love more likely or that make it less likely. Still, we cannot bring ourselves to love, or to stop loving, by an act of will alone—that is, merely by choosing to do so. And sometimes we cannot affect it by any means whatsoever.

The actual causes of love are various, and often difficult to trace. It is sometimes maintained that genuine love can be aroused only by the perceived value of the beloved object. The *value* of the beloved is what captivates the lover and moves him to love. If he were not responsive to its value, he would not love it. I do not deny that love may be aroused in this way. However, love does not *require* a response by the lover to any real or imagined value in what he loves. Parents do not ordinarily love their children so much, for example, because they perceive that their children possess exceptional value. In fact, it is the other way around: the children seem to the parents to be valuable, and they are valuable to the parents, only because the parents love them. Parents have been known to love—quite genuinely—children that they themselves recognize as lacking any particular inherent merit.

As I understand the nature of love, the lover does not depend for his loving upon reasons of any kind. Love is not a conclusion. It is not an outcome of reasoning, or a consequence of reasons. It *creates* reasons. What it *means* to love is, in part, to take the fact that a certain action would serve the good of the beloved as an especially compelling reason for performing that action.

We care about many things only for their instrumental value. They are intermediate goals for us, which we pursue as means to other things. Conceivably, a person's goals might all be intermediate: whatever he wants, he wants just for the sake of another thing; and he wants that other thing just in order to obtain something else; and so on. That sort of life could certainly keep a person busy. However, running endlessly from one thing to another, with no conclusive destinations, could not provide any full satisfaction because it would provide no sense of genuine achievement. We need *final* ends, whose value is not merely instrumental. I believe that our final ends are provided and legitimated by love.

Love is paradigmatically personal. What people love differs and may conflict. There is often, unfortunately, no way to adjudicate such conflicts. The account of normativity that I have been giving may therefore seem excessively skeptical. Many people are convinced that our final ends and values—most urgently our moral values—must be inescapably secured by reason, and must possess an inescapable authority that is altogether independent of anyone's personal desires and attitudes. What we should care about, they insist, must be determined by a reality entirely other than ourselves. My account is likely to strike them as radically neglectful of these requirements. They will have the idea that it is unacceptably noncognitive and relativistic. I think that idea is wrong, and I will try to correct it in my next lecture.

II. GETTING IT RIGHT

I

Suppose you are trying to figure out how to live. You want to know what goals to pursue and what limits to respect. You need to get clear about what counts as a good reason in deliberations concerning choice and action. It is important to you to understand what is important to you.

In that case, your most fundamental problem is not to understand how to identify what is valuable. Nor is it to discover what the principles of morality demand, forbid, and permit. You are concerned with how to make specific concrete decisions about what to aim at and how to behave. Neither judgments of value in general nor moral judgments in particular can settle this for you.

From the fact that we consider something to be valuable, it does not follow that we need to be concerned with it. There are many objects, activities, and states of affairs that we acknowledge to be valuable but in which we quite reasonably take no interest because they do not fit into our lives. Other things, perhaps even of lesser value, are more important to us. What we are actually to care about—what we are to regard as really important to us—cannot be based simply upon judgments concerning what has the most value.

In a similar way, morality too fails to get down to the bottom of things. The basic concern of morality is with how to conduct ourselves in our relations with other people. Now why should *that* be, always and in all circumstances, the most important thing in our lives? No doubt it is important; but, so far as I am aware, there is no convincing argument that it must invariably override everything else. Even if it were entirely clear what the moral law commands, it would remain an open question how important it is for us to obey those commands. We would still have to decide how much to care about morality. Morality itself cannot satisfy us about that.

What a person really needs to know, in order to know how to live, is what to care about and how to measure the relative importance to him of the various things about which he cares. These are the deepest as well as the most immediate normative concerns of our active lives. To the extent that we succeed in resolving them, we are able to identify and to order our goals. We possess an organized repertoire of final ends. That puts us in a position to determine, both in general and in particular instances, what we have reason to do. It is our understanding of what to care about, then, that is the ultimate touchstone and basis of our practical reasoning.

2

So, what *are* we to care about? This is not a matter that we can settle arbitrarily, or by deploying some shallow and unstable measure. In designing and committing our lives, we cannot rely upon casual impulse. Our deliberations and our actions must be guided by procedures and standards in which it is appropriate for us to have a mature confidence. The final ends by which we govern ourselves require authentication by some decisive rational warrant.

There is a famous passage in David Hume's *Treatise of Human Nature* that appears to rule out the possibility of providing any rational basis for deciding what we are to care about. Even the most grotesque pre-

ferences, Hume insists, are not irrational. He declares, for instance, that "'tis not *contrary* to *reason* to prefer the destruction of the whole world to the scratching of my finger."¹

Now it is true that this preference involves no purely *logical* mistake. So far as logic alone is concerned, it is unobjectionable. Someone who chooses to protect his finger from a trivial injury at the cost of unlimited destruction elsewhere is not thereby guilty of any contradiction or faulty inference. In this purely formal sense of rationality, his choice is not at all irrational.

But what would we say of someone who made that choice? We would say he must be *crazy*. In other words, despite the unassailability of his preference on logical grounds, we would consider both it and him to be wildly irrational. Caring more about a scratched finger than about "destruction of the whole world" is not just an unappealing personal quirk. It is *lunatic*. Anybody who has that preference is *inhuman*.

3

When we characterize the person in Hume's example as "crazy," or as "lunatic," or as "inhuman," these epithets do not function as mere vituperative rhetoric. They are literal denials that the person is a rational creature. There is a familiar mode of rationality, then, that is not exclusively defined by *a priori*, formal necessities. Hume's madman may be as competent as we are in designing valid chains of inference and in distinguishing between what is and what is not logically possible. His irrationality is not fundamentally a *cognitive* deficiency at all. He is *volitionally* irrational. He has a defect of the *will*, which bears upon how he is disposed to choose and to act.

Our basis for considering him to be volitionally irrational is not that his preferences happen to be merely very different from ours. It is that the relative importance to him of protecting his finger and of destroying the world is altogether *incommensurate* with how much we care about those things. He is moved to bring about unimaginable destruction for a reason that strikes us as so inconsequential as hardly to justify incurring any cost at all. An outcome from which we recoil in horror is, to him, positively attractive. The critical point has to do with possibilities: he is prepared to implement voluntarily a choice that we *could* not, under any circumstances, *bring ourselves* to make.

1. David Hume, *A Treatise of Human Nature*, edited by L. A. Selby-Bigge (Oxford: Oxford University Press, 1888), book II, part III, section III, p. 416 (emphasis added).

There are structural analogues between the requirements of volitional rationality and the strictly formal, *a priori* requirements of pure reason. Both modes of rationality limit what is possible, and each imposes a corresponding necessity. The boundaries of formal rationality are defined by the necessary truths of logic, to which no alternatives are conceivable. The boundaries of volitional rationality are defined by contingencies that effectively constrain the will. They limit what it is in fact possible for us to care about, what we can accept as reasons for action, and what we can actually bring ourselves to do. Violations of volitional rationality are not *inconceivable*. Rather, what stands in their way is that they are *unthinkable*.

Being volitionally rational is not just a matter of the choices that a person actually makes. It involves being *incapable*, under any circumstances, of making certain choices. If someone attempts to reach a cool and balanced judgment about whether it would be a good idea to destroy the entire world in order to avoid being scratched on his finger, that is not a demonstration of sturdy rationality. Even if he finally concludes that destroying the world to protect his finger is after all *not* such a good idea, the fact that he had to deliberate about this would make it clear that something is wrong with him.

Rationality does not *permit* us to be open-minded and judicious about everything. It requires that certain choices be utterly out of the question. Just as a person transgresses the boundaries of formal reason if he supposes of some self-contradictory state of affairs that it might really be possible, so he transgresses the boundaries of volitional rationality if he regards certain choices as genuine options.

A rational person cannot bring himself to do various things that, so far as his power and skill are concerned, he would otherwise be entirely capable of doing. He may think that a certain action is appropriate, or even mandated; but, when the chips are down, he finds that he just cannot go through with it. He cannot mobilize his will to implement his judgment. No reasons are good enough to move him actually to carry out the action. He cannot bring himself to destroy the world in order to avoid a scratch on his finger. In virtue of the necessities by which his will is constrained, making that choice is not among his genuine options. It is simply unthinkable.

What makes it unthinkable? Why are we unable to bring ourselves to do certain things? What accounts for our inability, or our inflexible refusal, to include among our alternatives various actions that we are otherwise quite capable of performing? What is the ground of the constraints upon our will that volitional rationality entails?

One view is that these volitional necessities are responses to an independent normative reality. On this account, certain things are *inherently* important. They therefore provide incontrovertible reasons for acting in certain ways. This is not a function of our attitudes or beliefs or desires, or of subjective factors of any kind. It does not depend in any way upon the condition of our will, or upon what we happen to regard as reasons for acting. In virtue of their unequivocal objectivity, moreover, these reasons possess an inescapable normative authority. It is the natural authority of the real, to which all rational thought and conduct must seek to conform.

In some way—just how is commonly left rather obscure—the independent reality of these reasons becomes apparent to us. We recognize, with a vivid clarity, that various things are inherently important. Then we cannot help accepting the authority of the reasons that they provide. It is impossible for us to deny, or to hold back from acknowledging, the importance that is—so to speak—right before our eyes. Seeing is believing. Thus, our will comes to be constrained by the forceful immediacy of reality.

This is the doctrine of “normative realism.” It holds that there are objective reasons for us to act in various ways, whether we know them, or care about them, or not. If we fail to appreciate and to accept those reasons, we are making a mistake. Some philosophers presume that normative realism is implicitly supported by the presumption that, as Robert Adams puts it, “keeping an eye out for possible corrections of our views is an important part of the seriousness of normative discourse.”² In their view, our concern to avoid mistakes—our belief that we need to get our normative judgments and attitudes *right*—“strongly favors” the supposition that the importance of reasons is inherent in them and that practical reason is therefore securely grounded in the independent reality of its governing norms.

2. Robert M. Adams, *Finite and Infinite Goods* (Oxford: Oxford University Press, 1999), p. 18.

6

My own view is different. I do not believe that anything is inherently important. In my judgment, normativity is not a feature of a reality that is independent of us. The standards of volitional rationality and of practical reason are grounded, so far as I can see, only in ourselves. More particularly, they are grounded only in what we cannot help caring about and cannot help considering important.

Our judgments concerning normative requirements can certainly get things wrong. There is indeed an objective normative reality, which is not up to us and to which we are bound to conform. However, this reality is not objective in the sense of being entirely outside of our minds. Its objectivity consists just in the fact that it is outside the scope of our voluntary control.

Normative truths require that we submit to them. What makes them inescapable, however, is not that they are grounded in an external and independent reality. They are inescapable because they are determined by volitional necessities that we cannot alter or elude. In matters concerning practical normativity, the demanding objective reality that requires us to keep an eye out for possible correction of our views is a reality that is within ourselves.

7

Let me begin to illustrate and to explain this by considering what I suppose everyone will agree is a clear paradigm of something that is genuinely important to us.

Except perhaps under very extraordinary conditions, the fact that an action would protect a person's life is universally acknowledged to be a reason for that person to perform the action. He may have a better reason for doing something else instead. There may even be entirely convincing reasons for him to prefer to die. However, self-defense is rarely (if ever) either thought to be a wholly *irrelevant* consideration in the evaluation of alternatives or thought to be in itself a reason *against* performing an action. Generally it is acknowledged without reserve to be at least a reason *in favor of* performing any action that contributes to it.

As a source of reasons for acting, our interest in staying alive has enormous scope and resonance. There is no area of human activity in which it does not generate reasons—sometimes weaker, sometimes stronger—for doing certain things or for doing things in a certain way. Self-preservation is perhaps the most commanding, the most protean,

and the least questioned of our final ends. Its importance is recognized by everyone, and it radiates everywhere. It infuses importance into innumerable objects and activities, and it helps to justify innumerable decisions. Practical reason could hardly get along without it.

8

How come? What accounts for the fact that we are always at least minimally attentive to the task of protecting our lives? What is it about survival that makes it at all important to us? What warrants our invariable acceptance of self-preservation as a reason that supports preferring one course of action over another?

Many people claim to believe that every human life is intrinsically valuable, regardless of how it is lived. Some individuals profess that what *they* are doing with *their* lives, or what they are likely to do with them, gives their lives a special importance. However, even when people have ideas like these about the value or importance of human life, that is ordinarily not the sole or even the primary explanation of why they are determined to go on living. It is not what really accounts for the fact that, in making decisions concerning what to do, they regard preserving their lives as a significant justifying consideration. Someone who acts in self-defense is universally conceded to have a pertinent reason for doing what he does, regardless of how he or others may evaluate his life.

Another view purports to identify reasons for living that do not require any assumption concerning the value of our lives. One of the best recent moral philosophers, the late Bernard Williams, suggests that it is a person's ambitions and plans—what he calls the person's "projects"—that provide "the motive force [that] propels [the person] into the future, and gives him a reason for living." These projects are "a condition of his having any interest in being around" in the world at all. Unless we have projects that we care about, Williams insists, "it is unclear why [we] should go on."³ In other words, we have a reason to do what it takes to go on living if we have projects that require our survival, but not otherwise.

That can't be right. It seems to me that what Williams says pertains just to people who are seriously depressed. The individuals he describes have no natural vitality. Their lives are inert, lacking any inherent momentum or flow. The movement from one moment to the next does

3. These quotations are from "Persons, Character, and Morality," in Bernard Williams, *Moral Luck* (Cambridge: Cambridge University Press, 1981), pp. 12–14.

not come to these people in the usual way—as a matter of course. They need a special push. They will move willingly into the future only if they are “propelled” into doing so. Unless they can supply themselves with an effectively propulsive fuel—“projects”—they will conclude that there is no reason for them to go on at all, and they will lose interest in being around.

Surely Williams has it backward. Our interest in living does *not* commonly depend upon our having projects that we desire to pursue. It’s the other way around: we are interested in having worthwhile projects because we do intend to go on living, and we would prefer not to be bored. When we learn that a person has acted to defend his own life, we do not need to inquire whether he had any projects in order to recognize that he had a reason for doing whatever it was that he did.

9

What ordinarily moves us to go on living, and also to accept our desire to continue living as a legitimate reason for acting, is not that we think we have reasons of *any* kind for wanting to survive. Our desire to live and our readiness to invoke this desire as generating reasons for performing actions that contribute to that end are not themselves based on reasons. They are based on love. They derive from and express the fact that, presumably as an outcome of natural selection, we love life. That is, we love living.

This does not mean that we especially enjoy it. Frequently we do not. Many people willingly put up with a great deal of suffering simply in order to stay alive. It is true, of course, that some people are so very miserable that they do really want to die. But this hardly shows that they do not love life. It only shows that they hate misery. What they would certainly prefer, if only they could arrange it, is not to end their lives but just to end the misery.

The desire to go on living is not only universal. It is irreducible. It is only if our prerational urge to preserve our lives has somehow become drastically attenuated that we demand reasons for preserving them. Otherwise, we do not require reasons at all. Our interest in self-preservation is a lavishly fecund source of reasons for choice and for action. However, it is not itself grounded in reasons. It is grounded in love.

In addition to their interest in staying alive, people generally have various other similarly primitive and protean concerns as well, which also provide them with reasons for acting. For instance, we cannot help caring about avoiding crippling injury and illness, about maintaining at least some minimal contact with other human beings, and about being free from chronic suffering and endlessly stupefying boredom. We love being intact and healthy, being satisfied, and being in touch. We cannot bring ourselves to be wholly indifferent to these things, much less categorically opposed to them. To a considerable degree, moreover, it is our concerns for them that give rise to the more detailed interests and ambitions that we develop in response to the specific content and course of our experience.

These fundamental necessities of the will are not transient creatures of social prescription or of cultural habit. Nor are they constituted by peculiarities of individual taste or judgment. They are solidly entrenched in our human nature from the start. Indeed, they are elementary constituents of volitional reason itself. It is conceivable, of course, that someone might actually not care a bit about these presumptively universal final ends. There is no logical barrier to rejecting them altogether or to being devoted to their opposites. Loving death, or incapacity, or isolation, or continuously vacant or distressing experience involves no contradiction. If a person did love those things, however, we would be unable to make sense of his life.

It is not terribly difficult to understand that a sensible person might regard certain states of affairs as giving him sufficient reason to commit suicide, or to incur crippling injuries, or to seek radical and permanent isolation, or to accept endless boredom or misery. What *would* be unintelligible is someone pursuing those things for their own sakes, rather than just to attain other goals that he cared about more. We could not empathize with, or expect ourselves to be understood by, someone who loves death or disability or unhappiness. We would be unable to grasp how he could possibly be drawn to what we cannot help being so naturally driven to avoid. His preferences, his deliberations, and his actions are guided by final ends that to us would be flatly incomprehensible. It makes no sense to us that anyone could love them.

I I

What is at stake here is not a matter of avoiding mistakes and getting things right. The volitionally irrational lover of death or disability or suffering has not overlooked something, or misunderstood something, or miscalculated, or committed *any* sort of error. From our point of view, his will is not so much in error as it is deformed. His attitudes do not depend upon beliefs that might be demonstrated by cogent evidence or argument to be false. It is impossible to reason with him meaningfully concerning his ends, any more than we could reason with someone who refuses to accept any proposition unless it is self-contradictory.

Many philosophers believe that an act is right only if it can be justified to other rational beings. For this to be plausible, it is not enough that the rationality of the others be merely of the formal variety. Those whom we seek to convince must be volitionally rational as well. If they are not, then their practical reasoning—however formally correct it may be—builds upon a foundation that is in radical opposition to ours. What justifies something to us will, to them, serve only to condemn it. We can therefore do no more with them than to express the bewilderment and revulsion that are inspired in us by the grotesque ends and ideals that they love.

I 2

So what is love? My conception does not aim at encompassing every feature of the hopelessly inchoate set of conditions that people think of as instances of love. The phenomenon that I have in mind includes only what is, for my purposes, philosophically indispensable. Most especially, it is not to be confused with romantic love, infatuation, dependency, obsession, lust, or similar varieties of psychic turbulence.

As I construe it, love is a particular mode of caring. It is an involuntary, nonutilitarian, rigidly focused, and—as is any mode of caring—self-affirming concern for the existence and the good of what is loved. The object of love can be almost anything—a life, a quality of experience, a person, a group, a moral ideal, a nonmoral ideal, a tradition, whatever. The lover's concern is *rigidly focused* in that there can be no equivalent substitute for its object, which he loves in its sheer particularity and not as an exemplar of some general type. His concern is *non-utilitarian* in that he cares about his beloved for its own sake, rather than only as a means to some other goal.

It is in the nature of the lover's concern that he is *invested* in his beloved. That is, he benefits when his beloved flourishes; and he suffers when it is harmed. Another way of putting it is that the lover *identifies himself* with what he loves. This consists in the lover accepting the interests of his beloved as his own. Love does not necessarily include a desire for union of any other kind. It does not entail any interest in reciprocity or symmetry in the relationship between lover and beloved. Moreover, since the beloved may be entirely unaware of the love, and may be entirely unaffected by it, loving entails no special obligation to the beloved.

Loving is risky. Linking oneself to the interests of another, and exposing oneself to another's vicissitudes, warrants a certain prudence. We can sometimes take steps that inhibit us from loving, or steps that stimulate us to love; more or less effective precautions and therapies may be available, by means of which a person can influence whether love develops or whether it lasts. Love is nonetheless *involuntary*, in that it is not under the immediate control of the will. We cannot love—or stop loving—merely by deciding to do so.

The causes of love are multifarious, and frequently obscure. In any event, love is not essentially a matter of judgment or of reasoned choice. People often think of what causes them to love something as giving them reasons to love it. However, loving is not the rationally determined outcome of even an implicit deliberative or evaluative process. Parmenides said that love is “the first-born offspring of necessity.”⁴ We come to love because we cannot help loving. Love *requires* no reasons, and it can have *anything* as its cause.

However, love is a powerful *source* of reasons. When a lover believes that an action will benefit his beloved, he does not need to wonder whether there is a reason for him to perform it. Believing that the action will have that effect means that he already has a reason. Insofar as a person loves something, he *necessarily* counts its interests as giving him reasons to serve those interests. The fact that his beloved needs his help is in itself a reason for him to provide that help—a reason that takes precedence, other things being equal, over reasons for being comparably helpful to something that he does not love. That is part of what it means to love. Loving thus *creates* the reasons by which the lover's acts of devotion to his beloved are dictated and inspired.

4. J. Burnet, *Early Greek Philosophy* (London: A. & C. Black, 1948), p. 177 (fragment 13).

Love entails two closely related volitional necessities. First, a person cannot help loving what he loves; and second, he therefore cannot help taking the expectation that an action would benefit his beloved as a powerful and often decisively preemptive reason for performing that action. Through loving, then, we acquire final ends to which we cannot help being bound; and by virtue of having those ends, we acquire reasons for acting that we cannot help but regard as particularly compelling.

It is not essential to love that it be accompanied by any particular feelings or thoughts. The heart of the matter is not affective or cognitive, but strictly volitional. The necessities of love, which drive our conduct and which circumscribe our options, are necessities of the will. Their grip means that there are certain considerations by which we cannot help being moved to act, and which we cannot help counting as reasons for action. What is essential to love is just these constrained dispositions to reason and to act out of concern for the beloved.

To be sure, the necessities that configure the lover's will are often associated with extravagant passion, and also with representations of the beloved as exceptionally worthy or attractive. It is not difficult to understand why. Love commits us to significant requirements and limitations. These are boundaries that delineate the substance and the structure of our wills. That is, they define what—as active beings—we most intimately and essentially are. Accordingly, love is not only risky. It profoundly shapes our personal identities and the ways in which we experience our lives.

Therefore, it is only natural that loving tends to arouse strong feelings in us. It is also only natural that we may hold ourselves away from loving until we are satisfied that it will be worth the anxieties, distractions, and other costs that it is likely to bring. Thus, love is often accompanied both by vivid enthusiasms and by reassuring characterizations of the beloved. These may be very closely related to loving, but the relationship is only contingent. They are not conceptually indispensable elements of love.

It is important to appreciate the difference between the necessities of love and various other deeply entrenched constraints upon the will, which are due to unwelcome and more or less pathological conditions such as compulsions, obsessions, addictions, and the like. These condi-

tions do not involve what I understand by the term “volitional necessity.” The necessities that they do involve may be even more urgent and more relentless than those of love; and their influence upon our lives may be no less pervasive and profound. However, they differ fundamentally from the volitional necessities of love in that we only submit to them unwillingly—that is, because they force us to do so. They are generated and sustained from outside the will itself. Their power over us is external, and merely coercive. The power of love, in contrast, is not like that.

Unfortunately, in attempting to explain the difference, it is easy to get lost in a thicket of complexities and qualifications. The trouble is that people are maddeningly nuanced and equivocal. It is impossible to grasp them accurately in their full depth and detail. They are too subtle, too fluid, and too mixed up, for sharp and decisive analysis. So far as love is concerned, people tend to be so endlessly ambivalent and conflicted that it generally cannot be asserted entirely without caveat either that they do love something or that they don't. Frequently, the best that can be said is that part of them loves it and part of them does not.

In order to keep my discussion here fairly simple, I therefore propose just to *stipulate* that a lover is never troubled by conflict, or by ambivalence, or by any other sort of instability or confusion. Lovers do not waver or hold back. Their love, I shall assume, is always robustly wholehearted, uninhibited, and clear.

Now the necessities of wholehearted love may be irresistible, but they are not coercive. They do not prevail upon the lover against his will. On the contrary, they are constituted and confirmed by the fact that he *cannot help* being wholeheartedly behind them. The lover does not passively submit to the grip of love. He is fully identified with and responsible for its necessities. There is no distance or discrepancy between what a lover is constrained to will and what he cannot help wanting to will. The necessities of love are imposed upon him, then, by himself. It is by his own will that he does what they require. That is why love is not coercive. The lover may be unable to resist the power it exerts, but it is his own power.

Moreover, the wholehearted lover cannot help being wholehearted. His wholeheartedness is no more subject to the immediate control of his will than is his loving itself. There may be steps that would cause his love to falter and to fade; but someone whose love is genuinely wholehearted cannot bring himself to take those steps. He cannot deliberately try to stop himself from loving. His wholeheartedness means, by definition,

that he has no reservations or conflicts that would move him to initiate or to support such an attempt. There is nothing within him that tends to undermine his love, or that gives him any interest in freeing himself from it. If the situation were otherwise, that would show either that his love had already somehow been undermined or that it had never been truly wholehearted to begin with.

15

The volitional necessities that I have been considering are absolute and unconditional. No rational person ever aims at death or disability or misery purely for their own sakes. In no possible circumstances could a rational person choose those things as final ends, or consider the likelihood that an action would achieve them as being in itself a reason for performing the action. Those judgments and choices are out of the question no matter what. They are precluded by volitional constraints that cannot be eluded and that never change.

Are these constraints “objective”? Well, in one sense they are obviously *not* objective. They derive from our attitudes; they are grounded nowhere but in the character of our own will. That evidently means that they are subjective. On the other hand, we cannot help having the dispositions that control the actions, choices, and reasons at issue. The character of our will could conceivably be different from what it is. However, its actual contingent necessities are rigorous and stable; and they are outside our direct voluntary control. This warrants regarding them as objective, despite their origin within us.

It seems to me that what the principles of morality essentially accomplish is that they elaborate and elucidate universal and categorical necessities that constrain the human will. They develop a vision that inspires our love. Our moral ideals define certain qualities and conditions of life to which we are lovingly devoted. The point of the moral law is to codify how personal and social relationships must reasonably be ordered by people who cannot help caring about the final ends that are most fundamental in the lives of all fully rational beings.

It is sensible to insist that moral truths are, and must be, stringently objective. After all, it would hardly do to suppose that the requirements of morality depend upon what we happen to want them to be, or upon what we happen to think they are. So far as I can see, all the objectivity required by the moral law is provided by the real necessities of our volitionally rational will. There is no need to look elsewhere to explain how

moral judgments can be objective. In any case, there is really nowhere else to look.

The truths of morality do not appear to be merely contingent. The appearance that they are necessary truths is, I believe, a reflection or a projection of the volitional necessities from which morality derives. We are aware that we have no choice, and we locate this inescapability in the object instead of in its actual source, which is within ourselves. If we suppose that the moral law is timeless and unalterable, that is because we suppose—rightly or wrongly—that the most fundamental volitional features of human nature are not susceptible to change.

The particular mode of opprobrium that is characteristic of our response to immorality is easy to account for when we recognize that our moral beliefs promote a vision of ideal personal and social relationships that has inspired our love. Attributing moral blame is distinctively a way of being angry at the wrongdoer. The anger is itself a kind of punishment. This is perhaps most transparent when a person directs his anger inward and suffers the lacerations of self-imposed feelings of guilt. What makes moral anger understandable and appropriate is that the transgression of an immoral agent consists in his willfully rejecting and impeding the realization of our moral ideal. In other words, he deliberately injures something that we love. That is enough to make anyone angry.

16

Needless to say, many of our volitional necessities and final ends are far from universal. The fact that I care about various specific individuals, groups, and ways of doing things is not a function simply of generic human nature. It arises from my particular makeup and experience. Some of the things that I happen to love are also loved by others; but some of my loves are shared only by, at most, a small number of people. The very fact that these more personal volitional necessities are not universal implies that they depend upon variable conditions. Naturally, we cannot change them at will; but they can be changed. Even within the life of a single individual, love comes and it goes.

This certainly does *not* mean that loving one thing is as good as loving another. It is true that nothing is *inherently* either worthy or unworthy of being loved, independently of what we are and what we care about. The ground of normativity is relative in part to the common nature of human beings and in part to individual experience and character. Still,

despite this relativity, there are plenty of ways that our loving can go absolutely wrong. There is plenty of room for demonstrating the seriousness of our normative discourse, in the way that counts so much for Adams and other normative realists, by “keeping an eye out for possible corrections of our views.”

We may need to correct our views concerning what is important to us because our love for one thing conflicts with our love for another. Perhaps we care about worldly success and also about peace of mind, and then it comes to our attention that pursuing the one tends to interfere with attaining the other. Determining which of the two we love more is likely to be facilitated by increasing our understanding of them. As we learn more about what each is and what it entails, it will often become clear that one arouses in us a more substantial interest and concern than the other.

Even when we are not aware of any conflict among our goals, it is only reasonable for us to be alert to the possibility that we do not understand the people and the ideals and the other things that we love well enough. Getting to know them better may reveal conflicts that previously were unnoticed. Our loving may turn out to have been misguided because its objects are not what we thought they were, or because the requirements and consequences of loving them differ from what we had supposed. In love, no less than in other matters, it is helpful to be clear about what we are getting into and what that lets us in for.

In addition to the fact that our understanding of the things we love may require correction, there is also the fact that we often do not understand ourselves very well. It is not so easy for people to know what they really care about or what they truly love. Our motives and our dispositions are notoriously uncertain and opaque, and we often get ourselves wrong. It is hard to be sure what we can bring ourselves to do, or how we will behave when the chips are down. The will is a thing as real as any reality outside us. The truth about it does not depend upon what we think it is, or upon what we wish it were.

Once we have learned as much as possible about the natural characteristics of the things we care about, and as much as possible about ourselves, there are no further substantive corrections that can be made. There is really nothing else to look for, so far as the normativity of final ends is concerned. There is nothing else to get right.

The legitimacy and the worthiness of our final ends are not susceptible to being demonstrated by impersonal considerations that all rational agents would accept as appropriately controlling. Sometimes normative disagreements cannot be rationally resolved. It may even be true that other people are required by what they care about to harm or to destroy what we love. Our love may be inspired by an endearing vision of how relationships between individuals might ideally be arranged; but other people may be driven by what they care about to struggle against arranging things in that way. There may be no convincing basis for regarding either them or ourselves as rationally defective or as having made some sort of mistake.

So far as reason goes, the conflict between us may be irreducible. There may be no way to deal with it, in the end, other than to separate or to slug it out. This is a discouraging outcome, but it does not imply a deficiency in my theory. It is just a fact of life.⁵

18

Wholehearted love definitively settles, for each of us, issues concerning what we are to care about. It expresses what we, as active individuals, cannot help being. We have no recourse other than to accept its dictates. Moreover, wholehearted love expresses—beyond that—what we cannot help wholeheartedly *wanting* to be. This means that we accept its authority not merely as inescapable, but as legitimate too. It is the only legitimate authority upon which, for each of us, our normative attitudes and convictions can properly and finally rely.

Even after we have recognized what it is that we love, and acquiesced to it as establishing the defining necessities of our volitional nature, problems do of course remain. We can fail what we love, through ignorance or ineptitude; and we can betray what we love, and thereby betray ourselves as well, through a shallow indulgence that leads us to neglect its interests and hence also to neglect our own. These problems have to do with competence and character.

On the other hand, for normative guidance in understanding what we should want or what we should do, there can be no authority superior to the welcome necessities of our own nature. As in the realm of politics,

5. There may be similarly irreducible conflicts within a single person, for whom there will then be no alternatives but to separate one part of himself radically from the other or to endure tumultuous inner conflict.

the legitimacy of authority here can derive only from the will of the governed. A rational acquiescence to this authority requires a clear self-understanding and a wholehearted acceptance of the essential requirements and boundaries of our will. This amounts to finding a mature confidence, which is not vulnerable to destruction of the self's integrity by familiar varieties of hyper-rationalistic skepticism.

This confidence, in which the authority of our norms of conduct is grounded, is a confidence in what we cannot help being. That provides us with the deepest and most secure foundation for practical reason. Without it, we could not even know where the exercise of practical rationality ought to begin. Without this confidence, in fact, there is no point in trying to become confident about anything else at all.⁶

6. It is worth noticing that Descartes found it impossible to rely confidently upon theoretical reason without first acquiring—through his argument that God could not have made him so defective as to be misled by the clear and distinct perceptions that he could not help accepting—a firm confidence in the necessities of his own cognitive nature. My argument about the ground of practical normativity is, I believe, significantly analogous to his argument about the ground of theoretical reason.