

persons are thought to be just insofar as they have, as one of the permanent elements of their character, a steady and effective desire to act justly. Aristotle's definition clearly presupposes, however, an account of what properly belongs to a person and of what is due to him. Now such entitlements are, I believe, very often derived from social institutions and the legitimate expectations to which they give rise. There is no reason to think that Aristotle would disagree with this, and certainly he has a conception of social justice to account for these claims. The definition I adopt is designed to apply directly to the most important case, the justice of the basic structure. There is no conflict with the traditional notion.

3. THE MAIN IDEA OF THE THEORY OF JUSTICE

My aim is to present a conception of justice which generalizes and carries to a higher level of abstraction the familiar theory of the social contract as found, say, in Locke, Rousseau, and Kant.⁴ In order to do this we are not to think of the original contract as one to enter a particular society or to set up a particular form of government. Rather, the guiding idea is that the principles of justice for the basic structure of society are the object of the original agreement. They are the principles that free and rational persons concerned to further their own interests would accept in an initial position of equality as defining the fundamental terms of their association. These principles are to regulate all further agreements; they specify the kinds of social cooperation that can be entered into and the forms of government that can be established. This way of regarding the principles of justice I shall call justice as fairness.

Thus we are to imagine that those who engage in social cooperation choose together, in one joint act, the principles which are to assign basic rights and duties and to determine the division of social benefits. Men are to decide in advance how they are to regulate their claims against one another and what is to be the foundation charter of their society. Just as each person must decide by rational reflection what constitutes his good,

4. As the text suggests, I shall regard Locke's *Second Treatise of Government*, Rousseau's *The Social Contract*, and Kant's ethical works beginning with *The Foundations of the Metaphysics of Morals* as definitive of the contract tradition. For all of its greatness, Hobbes's *Leviathan* raises special problems. A general historical survey is provided by J. W. Gough, *The Social Contract*, 2nd ed. (Oxford, The Clarendon Press, 1957), and Otto Gierke, *Natural Law and the Theory of Society*, trans. with an introduction by Ernest Barker (Cambridge, The University Press, 1934). A presentation of the contract view as primarily an ethical theory is to be found in G. R. Grice, *The Grounds of Moral Judgment* (Cambridge, The University Press, 1967). See also §19, note 30.

that is, the system of ends which it is rational for him to pursue, so a group of persons must decide once and for all what is to count among them as just and unjust. The choice which rational men would make in this hypothetical situation of equal liberty, assuming for the present that this choice problem has a solution, determines the principles of justice.

In justice as fairness the original position of equality corresponds to the state of nature in the traditional theory of the social contract. This original position is not, of course, thought of as an actual historical state of affairs, much less as a primitive condition of culture. It is understood as a purely hypothetical situation characterized so as to lead to a certain conception of justice.⁵ Among the essential features of this situation is that no one knows his place in society, his class position or social status, nor does any one know his fortune in the distribution of natural assets and abilities, his intelligence, strength, and the like. I shall even assume that the parties do not know their conceptions of the good or their special psychological propensities. The principles of justice are chosen behind a veil of ignorance. This ensures that no one is advantaged or disadvantaged in the choice of principles by the outcome of natural chance or the contingency of social circumstances. Since all are similarly situated and no one is able to design principles to favor his particular condition, the principles of justice are the result of a fair agreement or bargain. For given the circumstances of the original position, the symmetry of everyone's relations to each other, this initial situation is fair between individuals as moral persons, that is, as rational beings with their own ends and capable, I shall assume, of a sense of justice. The original position is, one might say, the appropriate initial status quo, and thus the fundamental agreements reached in it are fair. This explains the propriety of the name "justice as fairness": it conveys the idea that the principles of justice are agreed to in an initial situation that is fair. The name does not mean that the concepts of justice and fairness are the same, any more than the phrase "poetry as metaphor" means that the concepts of poetry and metaphor are the same.

Justice as fairness begins, as I have said, with one of the most general of all choices which persons might make together, namely, with the

5. Kant is clear that the original agreement is hypothetical. See *The Metaphysics of Morals*, pt. I (*Rechtslehre*), especially §§47, 52; and pt. II of the essay "Concerning the Common Saying: This May Be True in Theory but It Does Not Apply in Practice," in *Kant's Political Writings*, ed. Hans Reiss and trans. by H. B. Nisbet (Cambridge, The University Press, 1970), pp. 73-87. See Georges Vlachos, *La Pensée politique de Kant* (Paris, Presses Universitaires de France, 1962), pp. 326-335; and J. G. Murphy, *Kant: The Philosophy of Right* (London, Macmillan, 1970), pp. 109-112, 133-136, for a further discussion.

choice or the first principles of a conception of justice which is to regulate all subsequent criticism and reform of institutions. Then, having chosen a conception of justice, we can suppose that they are to choose a constitution and a legislature to enact laws, and so on, all in accordance with the principles of justice initially agreed upon. Our social situation is just if it is such that by this sequence of hypothetical agreements we would have contracted into the general system of rules which defines it. Moreover, assuming that the original position does determine a set of principles (that is, that a particular conception of justice would be chosen), it will then be true that whenever social institutions satisfy these principles those engaged in them can say to one another that they are cooperating on terms to which they would agree if they were free and equal persons whose relations with respect to one another were fair. They could all view their arrangements as meeting the stipulations which they would acknowledge in an initial situation that embodies widely accepted and reasonable constraints on the choice of principles. The general recognition of this fact would provide the basis for a public acceptance of the corresponding principles of justice. No society can, of course, be a scheme of cooperation which men enter voluntarily in a literal sense; each person finds himself placed at birth in some particular position in some particular society, and the nature of this position materially affects his life prospects. Yet a society satisfying the principles of justice as fairness comes as close as a society can to being a voluntary scheme, for it meets the principles which free and equal persons would assent to under circumstances that are fair. In this sense its members are autonomous and the obligations they recognize self-imposed.

One feature of justice as fairness is to think of the parties in the initial situation as rational and mutually disinterested. This does not mean that the parties are egoists, that is, individuals with only certain kinds of interests, say in wealth, prestige, and domination. But they are conceived as not taking an interest in one another's interests. They are to presume that even their spiritual aims may be opposed, in the way that the aims of those of different religions may be opposed. Moreover, the concept of rationality must be interpreted as far as possible in the narrow sense, standard in economic theory, of taking the most effective means to given ends. I shall modify this concept to some extent, as explained later (§25), but one must try to avoid introducing into it any controversial ethical elements. The initial situation must be characterized by stipulations that are widely accepted.

In working out the conception of justice as fairness one main task

clearly is to determine which principles of justice would be chosen in the original position. To do this we must describe this situation in some detail and formulate with care the problem of choice which it presents. These matters I shall take up in the immediately succeeding chapters. It may be observed, however, that once the principles of justice are thought of as arising from an original agreement in a situation of equality, it is an open question whether the principle of utility would be acknowledged. Off-hand it hardly seems likely that persons who view themselves as equals, entitled to press their claims upon one another, would agree to a principle which may require lesser life prospects for some simply for the sake of a greater sum of advantages enjoyed by others. Since each desires to protect his interests, his capacity to advance his conception of the good, no one has a reason to acquiesce in an enduring loss for himself in order to bring about a greater net balance of satisfaction. In the absence of strong and lasting benevolent impulses, a rational man would not accept a basic structure merely because it maximized the algebraic sum of advantages irrespective of its permanent effects on his own basic rights and interests. Thus it seems that the principle of utility is incompatible with the conception of social cooperation among equals for mutual advantage. It appears to be inconsistent with the idea of reciprocity implicit in the notion of a well-ordered society. Or, at any rate, so I shall argue.

I shall maintain instead that the persons in the initial situation would choose two rather different principles: the first requires equality in the assignment of basic rights and duties, while the second holds that social and economic inequalities, for example inequalities of wealth and authority, are just only if they result in compensating benefits for everyone, and in particular for the least advantaged members of society. These principles rule out justifying institutions on the grounds that the hardships of some are offset by a greater good in the aggregate. It may be expedient but it is not just that some should have less in order that others may prosper. But there is no injustice in the greater benefits earned by a few provided that the situation of persons not so fortunate is thereby improved. The intuitive idea is that since everyone's well-being depends upon a scheme of cooperation without which no one could have a satisfactory life, the division of advantages should be such as to draw forth the willing cooperation of everyone taking part in it, including those less well situated. The two principles mentioned seem to be a fair basis on which those better endowed, or more fortunate in their social position, neither of which we can be said to deserve, could expect the willing cooperation of others when some workable scheme is a necessary condition of the wel-

fare of all.⁶ Once we decide to look for a conception of justice that prevents the use of the accidents of natural endowment and the contingencies of social circumstance as counters in a quest for political and economic advantage, we are led to these principles. They express the result of leaving aside those aspects of the social world that seem arbitrary from a moral point of view.

The problem of the choice of principles, however, is extremely difficult. I do not expect the answer I shall suggest to be convincing to everyone. It is, therefore, worth noting from the outset that justice as fairness, like other contract views, consists of two parts: (1) an interpretation of the initial situation and of the problem of choice posed there, and (2) a set of principles which, it is argued, would be agreed to. One may accept the first part of the theory (or some variant thereof), but not the other, and conversely. The concept of the initial contractual situation may seem reasonable although the particular principles proposed are rejected. To be sure, I want to maintain that the most appropriate conception of this situation does lead to principles of justice contrary to utilitarianism and perfectionism, and therefore that the contract doctrine provides an alternative to these views. Still, one may dispute this contention even though one grants that the contractarian method is a useful way of studying ethical theories and of setting forth their underlying assumptions.

Justice as fairness is an example of what I have called a contract theory. Now there may be an objection to the term "contract" and related expressions, but I think it will serve reasonably well. Many words have misleading connotations which at first are likely to confuse. The terms "utility" and "utilitarianism" are surely no exception. They too have unfortunate suggestions which hostile critics have been willing to exploit; yet they are clear enough for those prepared to study utilitarian doctrine. The same should be true of the term "contract" applied to moral theories. As I have mentioned, to understand it one has to keep in mind that it implies a certain level of abstraction. In particular, the content of the relevant agreement is not to enter a given society or to adopt a given form of government, but to accept certain moral principles. Moreover, the undertakings referred to are purely hypothetical: a contract view holds that certain principles would be accepted in a well-defined initial situation.

The merit of the contract terminology is that it conveys the idea that principles of justice may be conceived as principles that would be chosen by rational persons, and that in this way conceptions of justice may be

explained and justified. The theory of justice is a part, perhaps the most significant part, of the theory of rational choice. Furthermore, principles of justice deal with conflicting claims upon the advantages won by social cooperation; they apply to the relations among several persons or groups. The word "contract" suggests this plurality as well as the condition that the appropriate division of advantages must be in accordance with principles acceptable to all parties. The condition of publicity for principles of justice is also connoted by the contract phraseology. Thus, if these principles are the outcome of an agreement, citizens have a knowledge of the principles that others follow. It is characteristic of contract theories to stress the public nature of political principles. Finally there is the long tradition of the contract doctrine. Expressing the tie with this line of thought helps to define ideas and accords with natural piety. There are then several advantages in the use of the term "contract." With due precautions taken, it should not be misleading.

A final remark. Justice as fairness is not a complete contract theory. For it is clear that the contractarian idea can be extended to the choice of more or less an entire ethical system, that is, to a system including principles for all the virtues and not only for justice. Now for the most part I shall consider only principles of justice and others closely related to them; I make no attempt to discuss the virtues in a systematic way. Obviously if justice as fairness succeeds reasonably well, a next step would be to study the more general view suggested by the name "rightness as fairness." But even this wider theory fails to embrace all moral relationships, since it would seem to include only our relations with other persons and to leave out of account how we are to conduct ourselves toward animals and the rest of nature. I do not contend that the contract notion offers a way to approach these questions which are certainly of the first importance; and I shall have to put them aside. We must recognize the limited scope of justice as fairness and of the general type of view that it exemplifies. How far its conclusions must be revised once these other matters are understood cannot be decided in advance.

4. THE ORIGINAL POSITION AND JUSTIFICATION

I have said that the original position is the appropriate initial status quo which insures that the fundamental agreements reached in it are fair. This fact yields the name "justice as fairness." It is clear, then, that I want to say that one conception of justice is more reasonable than another, or

6. For the formulation of this intuitive idea I am indebted to Allan Gibbard.

PREFACE

GROUND	317
COMMUNITY AND NATION	320
COMMUNITIES WHICH CHANGE	323
TOTAL COMMUNITIES	325
UTOPIAN MEANS AND ENDS	326
HOW UTOPIA WORKS OUT	331
UTOPIA AND THE MINIMAL STATE	333

Notes

Bibliography

Index

335
355
361

INDIVIDUALS have rights, and there are things no person or group may do to them (without violating their rights). So strong and far-reaching are these rights that they raise the question of what, if anything, the state and its officials may do. How much room do individual rights leave for the state? The nature of the state, its legitimate functions and its justifications, if any, is the central concern of this book; a wide and diverse variety of topics intertwine in the course of our investigation.

Our main conclusions about the state are that a minimal state, limited to the narrow functions of protection against force, theft, fraud, enforcement of contracts, and so on, is justified; that any more extensive state will violate persons' rights not to be forced to do certain things, and is unjustified; and that the minimal state is inspiring as well as right. Two noteworthy implications are that the state may not use its coercive apparatus for the purpose of getting some citizens to aid others, or in order to prohibit activities to people for their *own* good or protection.

Despite the fact that it is only coercive routes toward these goals that are excluded, while voluntary ones remain, many persons will reject our conclusions instantly, knowing they don't *want* to believe anything so apparently callous toward the needs and suffering of others. I know that reaction; it was mine when I first began to consider such views. With reluctance, I found myself becoming convinced of (as they are now often called) libertarian views, due to various considerations and arguments. This book contains little evidence of my earlier reluctance. Instead, it contains many of the considerations and arguments, which I present as forcefully as I can. Thereby, I run the risk of offending doubly: for the position

II THE PRINCIPLES REALIZED ON THE CAMPING TRIP

Two principles are realized on the camping trip, an egalitarian principle, and a principle of community. The community principle constrains the operation of the egalitarian principle by forbidding certain inequalities that the egalitarian principle permits. (The egalitarian principle in question is, as I shall explain, one of radical equality of opportunity: it is therefore consistent with certain inequalities of outcome.)

There are, in fact, a number of potentially competing egalitarian principles with which the camping trip, as I have

described it, complies, because the simple circumstances of the trip, unlike more complex ones, do not force a choice among them. But the only egalitarian principle realized on the trip that I shall bring into focus is the one that I regard as the correct egalitarian principle, the egalitarian principle that *justice* endorses, and that is a radical principle of equality of opportunity, which I shall call "socialist equality of opportunity."

Now, equality of opportunity, whether moderate or radical, removes obstacles to opportunity from which some people suffer and others don't, obstacles that are sometimes due to the enhanced opportunities that the more privileged people enjoy. Importantly, the removal of blocks to the opportunity of some people does

not always leave the opportunities of the initially better placed intact: sometimes it reduces the opportunities of those who benefit from inequality of opportunity. I underline this point because it means that promoting equality of opportunity is not only an *equalizing*, but also a *redistributing*, policy. Promoting equality of opportunity, in all of its forms, is not merely giving to some what others had and continue to enjoy.

We can distinguish three forms of equality of opportunity and three corresponding obstacles to opportunity: the first form removes one obstacle, the second form removes that one and a second, and the third form removes all three.

First, there is what might be called *bourgeois* equality of opportunity, by

which I mean the equality of opportunity that characterizes (at least in aspiration) the liberal age. Bourgeois equality of opportunity removes socially constructed status restrictions, both formal and informal, on life chances. An example of a formal status restriction is that under which a serf labors in a feudal society; an example of an informal status restriction is that from which a person whose skin is the wrong color may suffer in a society that is free of racist law but that nevertheless possesses a racist consciousness that generates racial disadvantage. This first form of equality of opportunity widens people's opportunities by removing constraints on opportunity caused by rights assignments and by bigoted and other prejudicial social perceptions.

Left-liberal equality of opportunity goes beyond bourgeois equality of opportunity. For it also sets itself against the constraining effect of social circumstances by which bourgeois equality of opportunity is undisturbed, the constraining effect, that is, of those circumstances of birth and upbringing that constrain not by assigning an inferior *status* to their victims, but by nevertheless causing them to labor and live under substantial disadvantages. The disadvantage targeted by left-liberal equality of opportunity derives immediately from a person's circumstances and does not depend for its constraining power on social perceptions or on assignments of superior and inferior rights. Policies promoting left-liberal equality of opportunity include head-start education

for children from deprived backgrounds. When left-liberal equality of opportunity is fully achieved, people's fates are determined by their native talent and their choices, and, therefore, not at all by their social backgrounds.

Left-liberal equality of opportunity corrects for *social* disadvantage, but not for native, or *inborn*, disadvantage. What I would call *socialist* equality of opportunity treats the inequality that arises out of native differences as a further source of injustice, beyond that imposed by unchosen social backgrounds, since native differences are equally unchosen. (Hence the similarity of the campers' attitudes to Sylvia's good luck and Leslie's, in scenarios b. and c. on pp. 8–9 above.) Socialist equality of opportunity seeks to correct

for *all* unchosen disadvantages, disadvantages, that is, for which the agent cannot herself reasonably be held responsible, whether they be disadvantages that reflect social misfortune or disadvantages that reflect natural misfortune. When socialist equality of opportunity prevails, differences of outcome reflect nothing but difference of taste and choice, not differences in natural and social capacities and powers.

So, for example, under socialist equality of opportunity income differences obtain when they reflect nothing but different individual preferences, including income/leisure preferences. People differ in their tastes, not only across consumer items, but also between working only a few hours and consuming rather little on the one hand, and working long hours

and consuming rather more on the other. Preferences across income and leisure are not in principle different from preferences across apples and oranges, and there can be no objection to differences in people's benefits and burdens that reflect nothing but different preferences, *when (which is not always) their satisfaction leads to a comparable aggregate enjoyment of life*. Such differences in benefits and burdens do not constitute inequalities of benefits and burdens.

Let me spell out the analogy at which I have just gestured. A table is before us, laden with apples and oranges. Each of us is entitled to take six pieces of fruit, with apples and oranges appearing in any combination to make up that six. Suppose, now, I complain that Sheila has five

apples whereas I have only three. Then it *should* extinguish my sense of grievance, a sense of grievance that is here totally inappropriate, when you point out that Sheila has only one orange whereas I have three, and that I could have had a bundle just like Sheila's had I forgone a couple of oranges. So, similarly, under a system where each gets the same income per hour, but can choose how many hours she works, it is not an intelligible complaint that some people have more take-home pay than others. The income/leisure trade-off is relevantly like the apples/oranges trade-off: that I have more income than you do no more shows, just as such, that we are unequally placed than my having four

apples from the table when you have two represents, just as such, an objectionable inequality. (Of course, some people love working, and some hate it, and that could be thought to [and I think it does] induce an injustice in the contemplated scheme, since those who love work will, *ceteris paribus*, relish their lives more than those who hate work do. But the same goes for some people enjoying *each* of apples and oranges more than others do; yet, even so, the apple/orange regime is a giant step toward equality [compared, for example, to a capitalist market society], and so, too, is equal pay for every hour worked, with each choosing the number of hours that she works. Accordingly, and for that reason, and also

because this is a short book, I ignore, henceforth, the complication that is exposed within these parentheses. I assume the fulfillment of the condition that was italicized at p. 19 above.)

Now, you might think that I have misused the term “socialist” in the phrase “socialist equality of opportunity,” for the simple reason that it is a familiar socialist policy to insist on equality of both income and hours of work: haven’t kibbutzim, those paradigms of socialism, worked that way?

I would draw attention, in reply, to the distinction between socialist principles and socialist modes of organization, the first, of course, being the putative justifications of the second. What I call “socialist equality of opportunity” is, as ex-

pounded here, a principle, one, so I say, that is satisfied on the camping trip, but I have not said what modes of organization would, and would not, satisfy it in general. And, although the suggested strictly equal work/wage regime would indeed contradict it, I acknowledge that socialists have advocated such regimes, and I have no wish, or need, to deny that those regimes can be called *socialist* work/wage regimes. What I do need to insist is that such systems contradict the fundamental principles animating socialists, when those principles are fully thought through. No defensible fundamental principle of equality or, indeed, of community, taken by itself, warrants such a system, which may nevertheless be amply justified as an appropriate

"second best" in light of the constraints of a particular place and time, and also in light of general constraints on the gathering and use of information about the tastes and powers of individuals, and, indeed, on their capacity to supply that information. Although justice might favor sensitivity to how happy or otherwise people are in similar circumstances—see the parenthetical remarks two paragraphs above—it is in general not possible, or even desirable, to fine-tune equality in that way.

What I have called *socialist* equality of opportunity is consistent with three forms of inequality, the second and third forms being subtypes of one type. Accordingly, I designate the three forms of inequality

as (i), (ii-a), and (ii-b). The first form of inequality is unproblematic, the second form is a bit problematic, and the third is very problematic.

(i) The first type, or form, of inequality is unproblematic because it does not constitute an inequality, all things considered. Variety of preference and choice across lifestyle options means that some people will have more goods of a certain sort than others do, but that is no inequality in the relevant sense, where comparable aggregate enjoyment obtains. That was the lesson of the apples/oranges example, and of its application to income/leisure choices.

(ii) The second type of inequality is problematic, since it does involve an inequality in aggregate benefit. For socialist

equality of opportunity tolerates inequality of benefit, where the inequality reflects the genuine choices of parties who are initially equally placed and who may therefore reasonably be held responsible for the consequences of those choices.

And this type of inequality takes two forms: inequality that is due to *regrettable choice* and inequality that is due to differences in *option luck*.

(ii-a) To illustrate the first of these forms, imagine that one apple/orange chooser (but not the other) carelessly waits so long that, by the time he picks up the fruit to which he's entitled, it has lost its full savor: the resulting inequality of benefit raises no grievance. And the same goes for someone in a work/pay re-

game whose ultimate fortune is inferior because she did not bother to examine her job opportunities properly.

These inequalities of aggregate benefit are justified by differential exercises of effort and/or care by people, who are, initially, absolutely equally placed, and who are equal even in their *capacities* to expend effort and care. A stock example is that of the grasshopper and the ant, who must be presumed to be equal in those capacities for the story to carry its customary moral. A mark of the fact that grasshopper/ant presents an inequality of benefit is that the homeless grasshopper retrospectively *regrets* his choice. He does so because he knows that, had he chosen as the ant did, his situation would now

be comparable to the ant's, instead of inferior in total benefit (including the benefits of idleness) to that of the ant.

Now you may be a skeptic about the grasshopper/ant story. You may believe (against the grain, I wager, of your reactions to people in ordinary life: see the paragraph that follows) that there is no such thing, ultimately, as being "truly responsible," or "choosing to be idle"; you may believe that greater negligence, for example, can reflect nothing but a smaller capacity for self-application, in the given circumstances, than others have, which should not be penalized. And if that is what you believe, then you will not countenance this second form of inequality. But even if, like me, you are not firmly disposed to disallow it, the question re-

mains, how *large* is this inequality likely to be? Well, that is a very difficult question, and my own view, or hope, is that, under an intelligent institutionalization of the relevant principle, it wouldn't be very large, *on its own*: it can, however, contribute to very high degrees of inequality when it's in synergy with the third and truly problematic form of inequality, that is, ii-b, which is consistent with socialist equality of opportunity.

I said that believing that no inequality could truly reflect real freedom of choice would contradict your reactions to people in day-to-day life, and that I lack that belief. I lack that belief because I am not convinced that it is true *both* that all choices are causally determined *and* that causal determination obliterates responsi-

bility. If you are indeed so convinced, then do not *blame* me for thinking otherwise, do not *blame* right-wing politicians for reducing welfare support (since, in your view, they can't help doing so), do not, indeed, blame, or praise, anyone for choosing to do anything, and therefore live your life, henceforth, differently from the way that we both know that you have lived it up to now.

(ii-b) The truly problematic inequality in overall benefit, the substantial inequality that is consistent with socialist equality of opportunity, is inequality that reflects differences in what philosophers call *option luck*. The paradigm case of option luck is a deliberate gamble. We start out equally placed, with \$100 each, and we are relevantly identical in all

respects, in character, in talents, and in circumstances. One of the features that we share is a penchant for gambling, so we flip a coin on the understanding that I give you \$50 if it comes up heads, and you give me \$50 if it comes up tails. I end up with \$150 and you end up with \$50, and with no extra anything to offset that monetary shortfall.

This inequality is consistent with socialist equality of opportunity; you and I simply used our radically similar opportunities, and, moreover, in exactly the same way. And unlike the grasshopper, the losing gambler, while of course regretting his loss, need not regret his decision to gamble in the way the grasshopper regrets his decision to be idle. The losing gambler can say: "Faced afresh

with the same options, I would have made the same choice; it was a reasonable gamble."

Now this form of inequality does not occur only as a result of gambling narrowly so called. There is also an element of option luck in the generation of market inequalities, which partly reflect gambles about where to put one's money or one's labor. Accordingly, some market-generated inequalities are partly compatible with, and, indeed, partly congruent with, socialist equality of opportunity. But one must not exaggerate the extent to which market inequalities have their genesis in pure market luck: market gambling differs significantly from standard gambling. Standardly, one has a choice whether or not to gamble: gambling is

avoidable. But the market is hardly avoidable in a market society: even the means of exiting any particular market society consist of resources that are accessible only on terms set by that market society. The market, one might say, is a casino from which it is difficult to escape, and the inequalities that it produces are tainted with injustice for that reason.

Whatever else is true, it is certainly safe to say that the yawning gulf between rich and poor in capitalist countries is not largely due to luck and the lack of it in optional gambling, but is rather a result of unavoidable gambling and straightforward brute luck, where no kind of gambling is involved. To be sure, avoidable option luck may figure in the explanation of cases where one entrepreneur prospers

threat to these notions that goes under the heading of the problem of free will. This problem takes a number of different forms, depending on what the threat is taken to be and on which judgments are seen as being threatened. The threat is often taken to be determinism, that is to say, roughly, the thesis that there are laws of nature from which it follows, given a complete description of the world at any one time, exactly what state the world must be in at any later time. But our moral ideas also seem to be threatened by a weaker claim, which I will call the Causal Thesis. This is the thesis that all of our actions have antecedent causes to which they are linked by causal laws of the kind that govern other events in the universe, whether these laws are deterministic or merely probabilistic. I will have something to say about both of these threats, but I will be particularly concerned with the compatibility of moral responsibility and the Causal Thesis because it seems to me that the space opened up by the falsity of determinism would be relevant to morality only if it were filled by something other than the cumulative effects of indeterministic physical processes. If the actions we perform result from the fact that we have a certain physical constitution and have been subject to certain causal influences, then an apparent threat to morality remains even if the laws linking these causes and their effects are not deterministic.

The idea that there is such a threat is sometimes supported by thought experiments such as the following. Suppose you were to learn that someone's present state of mind, intentions, and actions were produced in him or her a few minutes ago by "outside forces" such as electrical stimulation of his nervous system. You would not think it appropriate to blame that person for what he does under such conditions. But if the Causal Thesis is true, then all our actions are like this. The only differences are in the forms of outside intervention and the span of time over which it occurs, and surely these are not essential to the freedom of the agent in the sense relevant to moral responsibility.

Arguments of this kind are most commonly presented as a challenge to the applicability of moral praise and blame—that is, to "moral responsibility" in the sense I have called responsibility as attributability. But the significance of a person's choices for judgments of substantive responsibility is also threatened. We would not think, for example, that a promise was binding if the act of making the promise was brought about by outside causes of one of the kinds just listed. So some response to the apparent challenge of the Causal Thesis is required to defend judgments of responsibility of both of the kinds I have listed.

One form of response would be to argue that there are mistakes in the loose idea of causality to which this challenge appeals or in the assumption it makes about the relation between mental and physical events. No doubt there is much to be said on these topics, but my response will follow a different line. It is generally assumed that when an action is caused by factors outside the agent this makes moral appraisal inappropriate, but it is seldom explained why this should be so. I will argue that once we see clearly how the freedom of a choice or the voluntariness of an action is relevant to judgments of responsibility of the two kinds I have mentioned, it becomes apparent that neither form of responsibility is threatened by the general truth of determinism or of the Causal Thesis (although it can, of course, be undermined by particular kinds of causal interventions).

2. The Value of Choice

The task of this section and the next will be to explain the role that considerations of choice and responsibility play in reasons for rejecting principles. Two things need to be explained. The first is why principles that no one could reasonably reject often must be ones that make normative outcomes sensitive to individuals' choices, or at least to their having had the opportunity to choose. The second is how considerations of responsibility can diminish a person's reasonable grounds for rejecting a principle. How can the fact that a person could have avoided a certain burden by choosing appropriately make it the case that he cannot reasonably reject a principle that makes him bear that burden? My strategy will be to derive an answer to the second question from an answer to the first. Once we understand the positive reasons that people have for wanting opportunities to make choices that will affect what happens to them, what they owe to others, and what others owe to them, we can see also how their having had such opportunities can play a crucial role in determining what they can reasonably object to.¹

I will start, then, by examining some of the generic reasons that people have for wanting to have what happens depend on the way that they respond when presented with alternatives under the right conditions. When I go to a restaurant, for example, it is generally a good thing from my point of view to have what appears on my plate depend on the responses I give when presented with the menu. The most obvious reason for this is simply instrumental: I want what appears on

my plate to be something I will enjoy eating, and I believe that this is more likely to happen if what appears depends on my responses to the menu. This reason for valuing choice is both conditional and relative. It is conditional because the value of my response as a predictor of my future satisfaction depends on the nature of the question, my capacities of discernment, and the conditions under which my response is elicited. It is relative because it also depends on the reliability of the available alternative means for selecting the outcomes in question. In the restaurant case, for example, this value depends on how much I know about the cuisine in question and on my condition when the menu arrives: on whether I am drunk, for example, or too eager to impress my companions with my knowledge of French or my ability to swallow highly seasoned food. Thus the same interest that sometimes makes choice valuable—the interest in future satisfaction—can at other times provide reasons for wanting outcomes to be determined in some other way.²

What I have described so far might be called the “predictive” or “instrumental” value of choice. In the example I have given, choice is instrumental to my own future enjoyment, but the class of states that one might seek to advance by having outcomes depend on one’s choices is much broader. Aside from this range of instrumental reasons, however, there are other grounds for wanting to have what happens depend on my responses. One class of such reasons concerns the “representative” value of choice, which is illustrated by the example of choosing a gift.³ On our anniversary, I want not only to have a present for my wife, but also to have chosen that present myself. This is not because I think that I am more likely to come up with a present she will like (as far as that goes it would be better to have her choose the present herself). The reason is rather that the gift has a different meaning if I choose it myself—both the fact that I chose it and the choice that I make reflect my thoughts about her and about the occasion. In other cases, for reasons similar in character but opposite in sign, I may prefer that what happens should *not* depend on my choices. For example, I would prefer to have the question of who will get a certain job (whether it will be my friend or some well-qualified stranger) not depend on how I respond when presented with the choice: I want it to be clear that the outcome does not reflect my judgment of their respective merits or my balancing of the competing claims of merit and loyalty.

The reasons we have for wanting to see features of ourselves manifested in actions and their results are of course not limited to cases in

which our feelings for other people are at issue. I want to choose the furniture for my own living room, pick out the pictures for the walls, and write my own lectures despite the fact that these things might be done better by someone else. For better or worse, I want them to result from and hence to reflect my own taste, imagination, and powers of discrimination and analysis. We feel the same way, perhaps even more strongly, about important decisions affecting our lives in larger terms, such as what career to follow, where to work, and whom to marry.

These last examples, however, involve not only instrumental and representative but also what I will call symbolic value. In a situation in which people are normally expected to make choices of a certain sort for themselves, individuals have reason to value the opportunity to make these choices because not having or not exercising this opportunity would be seen as reflecting a judgment (their own or someone else’s) that they are not competent or do not have the standing normally accorded an adult member of the society. For example, in societies in which arranged marriages are not the norm, people have reason to want to choose their own mates rather than have their parents do it for them, not only because they think this will lead to a more satisfactory choice (instrumental value), or because they want their choice to be an expression of their own taste and affections (representative value), but also because having their parents make the choice would be “demeaning”—that is to say, would suggest that they are not competent, independent adults (symbolic value). In a different society, in which arranged marriage was common, reasons of this last sort would be much weaker.

I am not claiming that these three categories of reasons for valuing choice are mutually exclusive (representative and symbolic value may be difficult to distinguish in some cases, for example), or that taken together they encompass all the reasons one can have for valuing choice. My aim in listing and distinguishing them is simply to illustrate the variety of ways in which choice can be important and to make clear that its value is not always merely instrumental: the reasons people have for wanting outcomes to be dependent on their choices often have to do with the significance that this dependence itself has for them, not merely with its efficacy in promoting outcomes that are desirable on other grounds. I have presented them so far simply as reasons we commonly have for wanting to have certain powers and opportunities in our lives. But I take it as obvious that these are significant classes of

generic reasons that can figure as reasonable grounds for rejecting proposed moral principles.

Reasons of all three kinds would figure, for example, in determining the shape of an acceptable principle governing paternalistic interference in a person's life—that is to say, interference with a person's choices that is justified on the ground that it is “for his own good.” Possible grounds for rejecting principles permitting such interference include claims that the interference they permit (a) would deprive people of the opportunity to make choices with significant instrumental value, (b) would interfere with choices that have important representative value for people as ways of shaping their lives and expressing their values, or (c) would stigmatize those who are interfered with by labeling them as immature or incompetent. Where these three values are significant (and there are no sufficiently strong countervailing ones), the principles that no one could reasonably reject will forbid paternalistic intervention, thus making it a person's own (substantive) responsibility whether to risk the harms in question.

The pejorative ring of “paternalism” and the particular bitterness that it conveys stem from cases in which it is a matter of controversy whether the loss in question is a serious one or whether choices that may lead to it are foolish. Those who are inclined to make these choices may not see them as mistaken and, partly because of the controversy, may attach significant symbolic as well as representative value to being self-determining in this way and taking these risks. Consequently, they will resent paternalistic intervention which brands them as foolish when, as they see it, they merely differ from the majority in the things they value. Such sentiments are mistaken in some cases, although in others they can be the ground of reasonable objections. But this kind of resentment has no place at all in some other cases that are commonly called paternalistic, such as wage and hour laws or compulsory contributions to Social Security. These laws can both diminish and augment workers' ability to shape their lives to suit their preferences (diminish by ruling out choosing to work longer hours in order to make more money; augment by protecting people against being forced to make this choice). But they need not involve the kind of interference with representative values that gives paternalism its bad name, nor do they involve treating the workers they protect as foolish or incompetent. So it is useful, in weighing the force of objections to policies on the ground that they are “paternalistic,” to identify and distinguish between these various kinds of reasons rather than

simply appealing, in a general way, to the value of making one's own choices. I believe this is true in general of arguments about principles involving responsibility and choice.

As controversies about paternalism indicate, people disagree sharply about the value of various choices and opportunities to choose. Passionate disagreements about laws regarding seat belts and motorcycle helmets, for example, reflect disagreements about the instrumental, representative, and symbolic values of “unconstrained” choices about whether to use these devices, and hence about the permissibility of various forms of reminders and constraints. In addition to mere disagreement, however, different people have good reason to attach different value to various choices and opportunities. The same form of choice can have different instrumental value for people with different levels of knowledge, self-control, and discrimination, different representative value for people with different aims and attachments, and different symbolic value for people in different groups or societies. As I pointed out in Chapter 5, insofar as argument about the justifiability of moral principles has to proceed at a high level of generality, hence in ignorance of the distinctive situations of particular individuals, it must be based on “generic reasons”—that is to say, general conclusions about the reasons that individuals in a situation of a certain kind typically have. As I remarked earlier, there is always pressure to take account of individual differences by making principles more fine-grained and counterpressure to keep these principles simpler, hence more predictable and easier to apply. As we will see below, this tension is particularly clear in cases in which the opportunity to choose under a given set of conditions has very different value for different individuals.

The reasons I have listed for preferring principles that make what happens to us depend on the ways we respond when presented with alternatives are not undermined if it turns out that these responses have causes outside us. As long as these causes affect our responses only by affecting what we are like, it will remain true that these responses can be good predictors of what will bring us enjoyment or advance our aims. Similarly, in the case of representative value, it is quite plausible to suppose that many of the tastes and capacities for discernment that we want our choices to express have a basis in our causal makeup, but this fact does not make them less a part of us and hence does not diminish the value of choices that express them. In many cases what we want our choices to reflect is not simply our instinctive reactions but

rather our judgments of the relative force of competing reasons. Here again, as long as the causal processes in question do not deprive us of the ability to make such judgments at all they do not undermine the value of making choices that reflect these judgments; nor do they affect the fact that being allowed to make these judgments for ourselves is an important form of recognition as competent independent agents. So the instrumental, representative, and symbolic values of choice are not threatened by the Causal Thesis, whether the laws linking causes and effects are deterministic or merely probabilistic.

3. Responsibility and Choice

Having thus identified some of the positive reasons we have for wanting what happens to us to depend on our choices and other forms of response to alternatives, I want now to examine more fully the role that the value of choice plays in determining which principles it is reasonable to reject. It at least seems that when a person could have avoided a certain result by choosing appropriately, this fact weakens her grounds for rejecting a principle that would make her bear the burden of that result. Consider, for example, principles of the kind discussed in Chapter 5 that permit people to undertake useful projects that involve risk of harm to workers or others. It seems that the force of objections to such principles from the point of view of those who suffer this harm is significantly diminished if the principle requires that those who are exposed to this risk must be warned, in order to give them the opportunity of avoiding harm. The question to be addressed in this section is how this is to be explained.

To take a specific example, suppose that the officials of a city need to remove and dispose of some hazardous waste that has been found near a residential area. Apparently it has lain there for years, and they want to move it to a safer spot some distance away. Digging it up and moving it will inevitably release some hazardous chemicals into the air, but this is much less dangerous than leaving it in its present location, where it will in the long run seep into the water supply. Obviously they must take precautions to reduce the risks involved in this operation. They need to find a safe disposal site, away from where people normally have to go. They should build a fence around the new site, and another around the old one where the excavation is to be done, both of them with large signs warning people to keep away. They should also be sure to have the material wetted down and transported in closed

trucks to minimize the amount of hazardous material released into the air. Inevitably, however, some of it will be released—enough to cause lung damage to those who are directly exposed if, because of past exposure or genetic predisposition, they are particularly susceptible, but not enough to pose a serious threat to anyone who stays indoors and away from the excavation site. Given that this is so, the officials should be careful to warn people, especially those who know they are at risk, to stay indoors and away from the relevant areas while the work is being done.

Suppose that all these precautions were taken, but that some people were nonetheless exposed and suffered lung damage as a result. Let me stipulate that the officials did all that they could be expected to do to warn and protect these people (I am assuming that evacuation of everyone in the affected area was not an reasonable option), and that a principle permitting what they have done is thus one that no one could reasonably reject. In particular, this principle could not reasonably be rejected from the standpoint represented by the people who were in fact injured. The question at issue is what role the fact that they were warned, and thus given the choice of avoiding exposure, plays in making it the case that they cannot complain of the outcome. The Value of Choice account that I am proposing explains the role of choice in the justification of moral principles by appealing to the reasons (of the kinds described in the preceding section) we have for wanting outcomes to depend on the way we respond when presented with alternatives. In the present case these reasons are purely instrumental. No one has reason to place a positive value on having the opportunity to be exposed to hazardous chemicals. So according to the Value of Choice account "giving people the choice" is, like the fences, the careful removal techniques, and the remote location of the new site, just another means through which the likelihood of injury is reduced. But this may seem not to account for the full moral significance of the fact that those who were injured "knew what they were getting into." Consider the following two cases.

Suppose that one person was exposed because, despite the newspaper stories, mailings, posted signs, sound trucks, and radio and television announcements, he failed to hear about the danger. So he went for his usual walk without realizing what was going on. A second person heard the warnings but did not take the danger seriously. Curious to see what was being done, she went to the removal site and climbed over the fence to get a closer look.

There seems to be a clear difference between these two cases. In the first case, the officials have "done enough" to protect the person simply because they have gone to such lengths to reduce the likelihood that anyone will be injured (in particular, to reduce the likelihood that anyone will fail to know of the danger). After all, there is a limit to how much one has to do to protect people. But in the case of the second person, we can say more. Since she had been warned of the danger, and chose to go to the site anyway, we are inclined to say that she is (substantively) responsible for her own injury; and it is this fact, rather than the amount that has been done to protect her or the cost to others of doing more, that makes it the case that she cannot blame anyone for what happened. By choosing, in the face of warnings, to go to the excavation site, she laid down her right to complain of the harm she suffered as a result.

This familiar and intuitively appealing idea about the moral significance of choice, which I will call the Forfeiture View, is clearly distinct from the Value of Choice account that I have presented above.⁴ According to the Value of Choice account what matters is the value of the opportunity to choose that the person is presented with. If a person has been placed in a sufficiently good position, this can make it the case that he or she has no valid complaint about what results, whether or not it is produced by his or her active choice. On the Forfeiture View, on the other hand, it matters crucially whether an outcome actually resulted from a conscious decision in which the agent intentionally passed up specific alternatives. This is why that view accounts so well for our reaction to the person in the second example: not only does she have no one else to blame for her fate; she has *herself* to blame, since she chose to go to the site and climb over the fence.

This way of putting the matter can be misleading, however, insofar as it suggests that the person is responsible for her fate because she acted foolishly. What matters on the Forfeiture View is the fact of choice, not the faultiness of that choice. To see this, imagine a third person who, unlike the imprudently curious woman in my second example, did not run the risk of contamination foolishly or thoughtlessly. Suppose this person found that the day on which the excavation was to take place offered unusually good conditions for working outdoors on a scientific project to which she attached great value. Aware of the risk, she considered the matter carefully and decided that, taking into account her age and her likelihood of dying soon from an illness she already had, it was worth less to her to avoid this new

risk than to advance her project in the time she was likely to have remaining to her. Surely this person is just as fully responsible for her fate, according to the Forfeiture View, as the imprudent woman is. But this conclusion does not depend on our judging her decision to be foolish or mistaken. What lies behind the Forfeiture View is thus not a notion of desert, according to which people who behave wrongly or foolishly cannot complain about suffering as a result. The idea is rather that a person who could have chosen to avoid a certain outcome, but who knowingly passed up this choice, cannot complain of the result: *volenti non fit iniuria*.

The Forfeiture View is intuitively appealing, but this appeal should be resisted. Note, first, that the Value of Choice view can account for the apparent difference between the claims of the people in the first two examples described above. We may have taken all the steps we could be asked to take to protect the person in the first example, who nonetheless failed to learn what was being done. But because we did not succeed in making him aware of the danger, we did not make what happened to him depend on his response to this information. Given that this dependence is something we all would reasonably want to have under the circumstances, we did not succeed in making this person as well off as one would reasonably want to be. The woman in the second example, however, did have the benefit of being informed, even though this turned out to be worth less as a protection than it would have been to most other people. Since it is true in both cases that we did as much as we could be asked to do to protect these people against injury, neither can "complain"—neither could reasonably reject a principle permitting such a project to go forward with the safeguards it involved. But it remains true that only the woman in the second example was put in the position that these safeguards aim at providing for everyone.

What is different about the second example, then, has to do with the circumstances in which the woman was actually placed by our protective measures, not with the fact that she made a conscious decision to "take the risk." Consider, for example, a fourth person, who was informed of the risk of contamination but then simply forgot. As a result, when the trucks went by he was out in his yard exercising, breathing hard, with his Walkman turned up all the way. If, as in the other cases, enough was done to warn him, then this man is, like the woman in the second example, fully responsible for what happens to him, even though he made no conscious decision to take the risk.

From the fact that a person, under the right conditions, took a certain risk, we may conclude that he alone is responsible for what happens to him as a result. But this need not be seen as reflecting the special legitimating force of voluntary action, in the way that the Forfeiture View would suggest. The mere fact that he was placed in conditions in which he had the choice of avoiding the risk may be sufficient, as it was in the above examples. There are of course other cases in which it matters very much not only whether a person had the opportunity to make a conscious decision between certain alternatives but also whether he or she actually made such a decision. For example, a person is not bound by an agreement unless he or she actually consented to it. That is to say, a principle that no one could reasonably reject holds that if one has consented to an agreement under the right sort of conditions then one is bound by it; but it would be reasonable to reject any principle holding that one could be bound by contracts one had not consented to.⁵ The Value of Choice account can explain why this should be so. There are solid generic reasons to want the additional degree of control over one's obligations that a requirement of explicit consent provides, and these constitute reasonable grounds for rejecting principles that do not provide it. So contractualism can explain why various conclusions that look like instances of the Forfeiture View are in fact correct. It can explain why explicit consent or a conscious decision is required in some cases but not in others. Moreover, in explaining this, it does not need to appeal to the Forfeiture View to explain why various possible reasons for rejection have force or lack it. (This was the role for which the Forfeiture View was introduced in our discussion.) The explanations provided by the Value of Choice account are sufficient.

In cases of this last kind, such as voluntary agreements, the fact that a person chose to do something, or gave his or her consent to its being done, appears to have distinctive importance in part because it is the last justifying element to be put in place. If the conditions are right, then the person's choice or consent is sufficient (and in these cases generally necessary) to make the result morally legitimate. When we focus on this last step, the fact of choice appears to have the distinctive moral force that the Forfeiture View suggests. The conditions which must already have been in place in order for the choice to have this force recede into the background, or seem important only insofar as they affect the "voluntariness" of that choice. But this way of looking at the matter is misleading, for at least two reasons. First, as we have

seen, these "background" conditions are sometimes sufficient, by themselves, to determine the moral conclusion, without the occurrence of explicit choice or consent. Second, in cases in which explicit choice is important, it is unhelpful, and strains plausibility, to tie the significance of all the relevant background conditions to a notion of "voluntariness." For example, the difference between a person who entered into an unfavorable agreement because he completely ignored an alternative that was made known to him and a person who entered into such an agreement because that alternative was concealed from him does not lie in the fact that the former agreed voluntarily while the latter did not. One important virtue of the Value of Choice account is that it allows the various conditions under which a choice is (or could be) made to be taken into account separately from the fact of choice itself, and to be given the independent significance appropriate to them.

A defender of the Forfeiture View might raise a further objection at this point. According to the Value of Choice theory the moral significance of choice in cases like the hazardous waste removal example lies in its value as a protection. This is a special case of what I called, in the preceding section, the "predictive" or "instrumental" value of choice. "Having the choice"—having what happens depend on how one responds when presented with certain alternatives—is something we have reason to want because it decreases the likelihood of our suffering certain harms. But it might be objected that in the case of the imprudent woman who climbed over the fence being warned of the danger and given the choice of avoiding it turned out to have no positive value. The warning only aroused her impetuous curiosity, and she would have been better off if she had never been told at all. Yet we still feel that it is important that she was warned to stay home, and that the fact that she was warned is an important element supporting the conclusion that she cannot complain about what happened to her (and could not reasonably reject a principle permitting such a project to be undertaken). The objection concludes that since this moral significance cannot be accounted for by the positive value to her of being given this choice we need some other explanation. The natural alternative is something like the Forfeiture View.

Moreover, it may seem that if the Forfeiture View is to do the work required of it in such a case, it must rely on a more robust notion of freedom than that employed by the Value of Choice account. I claimed earlier that the Value of Choice account is not threatened by the

Causal Thesis. Choice can retain its value even if it is caused, ultimately, by factors outside us, I argued, as long as these factors operate "through us." Choice retains its instrumental and representative value, for example, if we are caused to make certain choices by the factors that cause us to be the kind of people who would so choose. But the case of the imprudent woman poses a challenge to this account. If the Causal Thesis is true, then she is caused to act as she does by factors operating "through her" in the relevant sense: she is an imprudent person. But if we accept the Forfeiture View as the right account of the moral significance of choice in her case, then we seem forced to say that she is responsible for what happens to her because she could have acted differently than she did—hence, if the Causal Thesis is accepted, differently than she was in fact caused to do. So if we accept the account of this case offered by the Forfeiture View we will need to come up with some way of understanding this "could have" and some way of explaining how it is compatible with the Causal Thesis and with the thesis of determinism.

I believe, however, that we should not accept this account. As was shown earlier by the comparisons with other examples (such as the cases of the man who did not get the word, and the one who did get it but then forgot), what is crucial in the case of the imprudent woman is not the fact of her choice but rather the circumstances in which she was placed. Moreover, an explanation that made this case turn on the fact that she "could have done otherwise" would lead to implausible results in other cases, since there are many conditions that undermine the legitimating force of choice despite the fact that a person choosing under such conditions still "could have done otherwise" in any sense that would apply in this case. It would, for example, be reasonable to reject a principle according to which a long-term contract is binding even when entered into by a fourteen-year-old without adult guidance. What is special about the case of fourteen-year-olds, however, is not that they *cannot* choose wisely (after all, many of them do), but rather that they are so likely not to do so. The same can be said, I believe, about the choices that adults make under various conditions, such as when they are drunk, or overcome with grief or fear. It is often difficult to decide whether these conditions are sufficient to render agreements invalid, but if "ability to do otherwise" were the test then it would be obvious that none of them are sufficient.

As I pointed out in the previous section, the instrumental value of choice is conditional and relative, and hence varies from one individual

to another. The case of the imprudent woman illustrates this variability: an opportunity that would have great instrumental value for most people as a way of avoiding risk turns out to do her no good, and perhaps even to make her worse off. The objection asks how, in light of this fact, having given her this choice can count as improving the conditions under which she was placed and thereby helping to make it the case that she cannot complain of the result. The answer has two parts. The first lies in the distinction between the generic reasons on which the justifiability of a moral principle must rest and the reasons that a specific individual may have, given all the facts about his or her situation. The reason why it is important that this woman was informed of the danger, and thus given the choice of avoiding it, is not that this was necessarily advantageous to her but rather that it is something that people in general have reason to value and hence to demand that an acceptable principle insist on. But (and this is the second point) this is not the only thing they have reason to insist on. Because some people are likely to choose unwisely, it is not enough merely to warn people even if they could all protect themselves by taking appropriate precautions. This is why, in the example given, it was necessary to put fences around the excavations and to wet the hazardous material down and transport it carefully in order to minimize the risk to those who failed to stay indoors. If the imprudent woman has no grounds for complaint it is because adequate precautions of this kind were taken. If she does have such grounds, it is because any principle that no one could reasonably reject would require more precautions of this kind in order for the project to go forward. The Forfeiture View, by concentrating on the fact of choice and the possibility of choosing otherwise as the morally significant features of the case, focuses our attention in the wrong place.

I conclude that the Value of Choice account provides the best explanation of the significance that our choices and other forms of response have in determining what we owe to each other. This account is particularly helpful insofar as it calls attention to the variety of ways in which these responses can be significant and to the importance of conditions other than the fact of choice itself. As a further illustration of this account, and as the basis for contrasting it with the account of moral responsibility that will be the subject of the next section, consider how it would apply to the justification of criminal punishment.

I will begin with a schematic comparison of the institution of punishment and the policy of waste removal that we have just been dis-

cussing. In each case we have the following elements. First, there is an important social goal: protecting the water supply in the one case, protecting ourselves and our possessions in the other. Second, there is a strategy for promoting this goal that involves the creation of another risk: the risk of contamination in the one case, the risk of punishment in the other. Third, the effect of this strategy is that there is (literally or metaphorically) a certain affected "area" that one can no longer enter without danger. In the one case this is the area in which the excavation, transport, and disposal of the waste are being carried out; in the other case it is the metaphorical "area" consisting of the range of activities that have been declared illegal. Fourth, even if safeguards are introduced to reduce people's exposure to the risk created, it remains overwhelmingly likely that many of those who enter the affected areas, and perhaps some others, will be harmed. Some of the safeguards that are introduced (such as requirements of due process, and careful methods of excavation and transport) serve to protect those who stay out of the affected area. Other safeguards enhance the value of choice as a protection by making it less likely that people will choose to enter this area. In the example of hazardous waste these include signs and other publicity informing people of the risk, as well as fences, guards, and the choice of an obscure disposal site where no one has reason to go. Analogous features in the case of punishment are education (including moral education), the dissemination of information about the law, and the maintenance of social and economic conditions that reduce the incentive to commit crime by offering the possibility of a satisfactory life within the law. Restrictions on entrapment by law enforcement officers also belong in this category of safeguards, as do provisions which excuse from punishment those who, because of mental illness or defect, are unable to regulate their conduct in accordance with the law. Without safeguards of these kinds, the value of choice as a protection would be unacceptably low.

In each case, in order to defend the practice in question we need to argue that the importance of the social goal justifies creating the risk and making the "affected area" unsafe, and that given the safeguards that have been put in place enough has been done to protect people against suffering harm from the threat that has been created.

There are some significant differences between the two cases. First, insofar as the activities that make up the "affected area" in the case of punishment are ones that it is morally wrong to engage in, being deprived of the opportunity of "entering this area" without risk does

not count as a morally significant loss. This eases the task of justification somewhat in that case. A second difference makes that task more difficult, however. It is not the aim of either policy that anyone should be harmed by the threat that is created. But in the case of punishment, when this harm does occur it is deliberately inflicted on particular people, as required by the institution itself. It is an essential part of the institution of punishment that those who violate the law should be punished, but it is no part of the waste-removal program that those who go to the excavation site should suffer contamination. Thus, since a policy of deliberately inflicting harm is more difficult to justify than a policy of creating a risk while trying (no doubt imperfectly) to protect people against it, the institution of punishment carries a heavier burden of justification than the program of waste removal.

When such an institution is justified, however, this justification entails a stronger form of "forfeiture" than we found in considering the example of hazardous waste removal. A person who knowingly and intentionally violates a justifiable law lays down his or her right not to suffer the prescribed punishment: that is to say, such a person has no legitimate complaint against having this penalty inflicted. This forfeiture is a consequence of the justification for the institution of punishment and for the particular law in question, however, not an element in that justification. More specifically, it is a consequence of the "heavier burden of justification" to which I just referred. Because the institution assigns punishment to those who fulfill certain conditions, justifying the institution involves justifying the infliction of these penalties. If the conditions for punishment include having made a certain kind of choice, then a justification for punishment justifies making that choice a necessary and, when other conditions are fulfilled, sufficient condition for punishment. No such assignment and no such forfeiture are involved in the justification of the program of hazardous waste removal. A person who recklessly chooses to enter the affected area does not lay down a right to further protection against contamination. She has, by assumption, already received all the prior protection she is entitled to, and she does not lay down her right to rescue or treatment unless this has been prescribed and the policy including this prescription is justified. Forfeiture is a creature of particular institutions and relatively specific principles such as those governing promising (provided, of course, that these are justifiable). It is not a moral feature of choice in general. As I argued above, what figures in the justification of such institutions and principles is not

forfeiture but, rather, the less sharp-edged notion of the value of choice.

A justification for an institution of punishment along the lines just sketched would not be a retributivist account, since it makes no appeal to the idea that it is morally desirable for wrongdoers to suffer punishment. But the moral status of the acts for which someone can be punished would remain important in this justification, in two ways that can be brought out by considering the possibility of "strict liability" offenses. As I said above, if the actions for which punishment is prescribed are ones that it would be morally wrong to perform, then a person cannot complain of a significant loss of liberty when this "affected area" is made one that one cannot enter without fear of punishment. Systems of punishment that cover a wider range of actions could be justified, however. If, for example, a legal penalty is attached to selling contaminated milk (not merely to doing so knowingly, recklessly, or negligently), then one "enters the affected area" just by going into the milk business, and if such a law is justified then doing this involves laying down one's (legal) right not to be penalized in the event that the milk one sells turns out to be impure even though one has not been negligent in handling it. But this enlargement of the affected area to include morally unobjectionable activities (such as conscientious engagement in the milk business) makes such laws more difficult to justify.

There is also a second source of difficulty. The criminal law is not just an organized system of threats, or a system of taxes designed to encourage some activities and discourage others. It provides rules and standards that good citizens are supposed to respect—that is, to regard as norms which they accept as reason-giving. Punishment is thus in part an expression of an authoritative judgment that the criminal has not done this—an expression of "legal blame."⁶ Insofar as this is so, it will seem inappropriate to punish, and hence to condemn, someone whose conduct is admitted to be blameless, such as a person who through no fault of his own has sold milk that turns out to be impure. This will seem particularly inappropriate when the punishment takes one of the forms (such as imprisonment) that has the clearest condemnatory force.

This inappropriateness remains even if the person punished had ample opportunity to avoid the sanction. After all, no one has to go into the milk business, and, we may suppose, everyone who does so knows what the penalties are for selling adulterated milk. Punishment

thus has two related aspects: the penalties that are inflicted, such as fines and loss of liberty, and the condemnation that these normally express. (These are related since the unpleasantness of being publicly condemned is itself a penalty.) The schematic justification described above dealt only with the first of these aspects. It argued that, like the injuries involved in moving the hazardous waste, these penalties can be justified when certain conditions are fulfilled, prominent among which is the requirement that those who are punished should have had the opportunity to avoid punishment by choosing appropriately. As we have just seen, however, the condemnatory aspect of punishment is subject to a further requirement: the condemnation must be appropriate. What triggers this requirement is not the unpleasantness of the condemnation, but the content of the judgment expressed. Insofar as punishment involves an assertion that the agent governed himself in a way that was faulty, it is appropriate only when this is true.

So an analysis of the kind presented in this section, based on the value of choice as a protection, provides only a partial account of the conditions under which punishment is justified—an account dealing only with the "penalty" aspect of punishment. Since the penalties involved in the criminal law are so severe, however, the justifiability of this aspect of punishment is by far the more important question. In the case of morality the reverse is true: moral blame is fundamentally a judgment of condemnation, not a penalty. Accordingly, as I will argue in the following section, the notion of responsibility that is a condition for moral blame is quite different from the one we have just been discussing.

4. Moral Appraisal

I turn now to the other sense of responsibility distinguished at the beginning of this chapter: responsibility as a precondition for moral appraisal, rather than as one of the ideas figuring in the principles on which this appraisal is based. The task of this section will be to spell out more fully the conception of moral appraisal that follows from the theory developed in earlier chapters. I will argue that this conception offers plausible accounts both of the range of things that can be objects of moral assessment and of the force and content of our notions of moral blame and guilt. This account will provide the basis for considering, in the following section, the conditions under which moral appraisal must be suspended—that is to say, the conditions that mark