

traditional issues of ethical, moral, and legal theory. My overall thesis is unpopular now—the fox has ruled the roost in academic and literary philosophy for many decades, particularly in the Anglo-American tradition.<sup>2</sup> Hedgehogs seem naïve or charlatans, perhaps even dangerous. I shall try to identify the roots of that popular attitude, the assumptions that account for these suspicions. In this introductory chapter I offer a road map of the argument to come that shows what I take those roots to be.

My advance summary could start in any chapter, fanning out from there, tracing the implications of that chapter for the rest. But I think it best to start at the end of the book, with political morality and justice, so that readers particularly interested in politics will have an advance understanding of why I think that the more abstract philosophical discussions of the book are required steps to what concerns them most. I hope that starting the summary there will also encourage other readers whose greater interest lies in more mainstream issues of philosophy—meta-ethics, metaphysics, and meaning—to find practical importance in what they might believe to be abstruse philosophical issues.

### Justice

*Equality.* No government is legitimate unless it subscribes to two reigning principles. First, it must show equal concern for the fate of every person over whom it claims dominion. Second, it must respect fully the responsibility and right of each person to decide for himself how to make something valuable of his life. These guiding principles place boundaries around acceptable theories of distributive justice—theories that stipulate the resources and opportunities a government should make available to people it governs. I put the matter that way, in terms of what governments should do, because any distribution is the consequence of official law and policy: there is no politically neutral distribution. Given any combination of personal qualities of talent, personality, and luck, what a person will have by way of resource and opportunity will depend on the laws in place where he is governed. So every distribution must be justified by showing how what government has done respects these two fundamental principles of equal concern for fate and full respect for responsibility.

A laissez-faire political economy leaves unchanged the consequences of a free market in which people buy and sell their product and labor as they wish and can. That does not show equal concern for everyone. Anyone impover-

ished through that system is entitled to ask: “There are other, more regulatory and redistributive, sets of laws that would put me in a better position. How can government claim that this system shows equal concern for me?” It is no answer that people must take responsibility for their own fate. People are not responsible for much of what determines their place in such an economy. They are not responsible for their genetic endowment and innate talent. They are not responsible for the good and bad luck they have throughout their lives. There is nothing in the second principle, about personal responsibility, that would entitle government to adopt such a posture.

Suppose government makes the extreme opposite choice, however: to make wealth equal no matter what choices people made for themselves. Every few years, as would be possible in a Monopoly game, government calls in everyone’s wealth and redistributes it in equal shares. That would fail to respect people’s responsibility to make something of their own lives, because what people chose to do—their choices about work or recreation and about saving or investment—would then have no personal consequences. People are not responsible unless they make choices with an eye to the costs to others of the choices that they make. If I spend my life at leisure, or work at a job that does not produce as much as I could of what other people need or want, then I should take responsibility for the cost this choice imposes: I should have less in consequence.

The question of distributive justice therefore calls for a solution to simultaneous equations. We must try to find a solution that respects both the reigning principles of equal concern and personal responsibility, and we must try to do this in a way that compromises neither principle but rather finds attractive conceptions of each that fully satisfy both. That is the goal of the final part of this book. Here is a fanciful illustration of a solution. Imagine an initial auction of all available resources in which everyone starts with the same number of bidding chips. The auction lasts a very long time, and will be repeated as long as anyone wishes. It must end in a situation in which nobody envies anybody else’s bundle of resources; for that reason the distribution of resources that results treats everyone with equal concern. Then imagine a further auction in which these people design and choose comprehensive insurance policies, paying the premium the market establishes for the coverage each chooses. That auction does not eliminate the consequences of good or bad luck, but it makes people responsible for their own risk management.

We can use that fanciful model to defend real-life distributive structures. We can design tax systems to model these imaginary markets: we can set tax

rates, for instance, to mimic the premiums it seems reasonable to assume people would pay in the hypothetical insurance market. The rates of tax designed in that way would be fairly steeply progressive; more so than our tax rates at present. We can design a health care system mimicking the coverage it seems reasonable to assume people would seek: this would require universal health care. But it would not justify spending, as Medicare now does, enormous sums keeping people alive in the last few months of their lives, because it would make no sense for people to give up funds useful for the rest of their lives to pay the very high premiums required by that sort of coverage.

*Liberty.* Justice requires a theory of liberty as well as a theory of resource equality, and we must be aware, in constructing that theory, of the danger that liberty and equality will conflict. It was Isaiah Berlin's claim that such conflict is inevitable. I argue, in Chapter 17, for a theory of liberty that eliminates that danger. I distinguish your freedom, which is simply your ability to do anything you might want to do without government constraint, from your liberty, which is that part of your freedom that government would do wrong to constrain. I do not endorse any general right to freedom. I argue, instead, for rights to liberty that rest on different bases. People have a right to ethical independence that follows from the principle of personal responsibility. They have rights, including rights to free speech, that are required by their more general right to govern themselves, which right also flows from personal responsibility. They have rights, including rights to due process of law and freedom of property, that follow from their right to equal concern.

This scheme for liberty rules out genuine conflict with the conception of equality just described because the two conceptions are thoroughly integrated: each depends on the same solution to the simultaneous equation problem. You cannot determine what liberty requires without also deciding what distribution of property and opportunity shows equal concern for all. The popular view that taxation invades liberty is false on this account provided that what government takes from you can be justified on moral grounds so that it does not take from you what you are entitled to retain. A theory of liberty is in that way embedded in a much more general political morality and draws from the other parts of that theory. The alleged conflict between liberty and equality disappears.

*Democracy.* But there is another supposed conflict among our political values. This is the conflict between equality and liberty, on the one hand, and

the right to participate as an equal in one's own governance, on the other. Political theorists sometimes call the latter a right to positive liberty and suppose that that right may conflict with negative liberty—the rights to freedom from government I just described—and also with the right to a just distribution of resources. The conflict is realized, on this view, when a majority votes for an unjust tax scheme or a denial of important liberties. I respond to that claim of conflict by distinguishing various conceptions of democracy. I distinguish a majoritarian or statistical conception from what I call the partnership conception. The latter holds that in a genuinely democratic community each citizen participates as an equal partner, which means more than just that he has an equal vote. It means that he has an equal voice and an equal stake in the result. On that conception, which I defend, democracy itself requires the protection of just those individual rights to justice and liberty that democracy is sometimes said to threaten.

*Law.* Political philosophers insist on yet another conflict among political values: the conflict between justice and law. Nothing guarantees that our laws will be just; when they are unjust, officials and citizens may be required, by the rule of law, to compromise what justice requires. In Chapter 19 I speak to that conflict: I describe a conception of law that takes it to be not a rival system of rules that might conflict with morality but as itself a branch of morality. It is necessary, to make that suggestion plausible, to emphasize what might be called procedural justice, the morality of fair governance as well as just outcome. It is also necessary to understand morality in general as having a tree structure: law is a branch of political morality, which is itself a branch of a more general personal morality, which is in turn a branch of a yet more general theory of what it is to live well.

You will by now have formed a suspicion. Poseidon had a son, Procrustes, who had a bed; he suited his guests to his bed by stretching or lopping them until they fit. You might well think me Procrustes, stretching and lopping conceptions of the great political virtues so that they neatly fit one another. I would then be achieving unity on the cheap: a meaningless victory. But I mean to submit each of the political conceptions I describe to the test of conviction. I will not rely on any assumption that a theory is sound just because it fits with other theories we also find agreeable. I hope to develop integrated conceptions that all seem right in themselves, at least after reflection. I do make an independent and very powerful claim, however. I argue throughout the book that in political morality integration is a necessary condition of truth.

But that is the point. A community that respects personal ethical responsibility must concentrate on a fair distribution of means when it fixes its political settlement. It must leave the choice of ends to its citizens one by one.<sup>8</sup>

## Equality of Resources

### *The Envy Test*

What political settlement, seeking what distribution of resources, fits our two principles taken together? I have proposed a fantasy answer.<sup>9</sup> Imagine people shipwrecked on an empty island with diverse natural resources. They are each given an equal number of clamshells as bidding tokens, and they compete in an auction for individual ownership of the island's resources. When the auction finally ends, and everyone is satisfied that he has used his clamshells most efficiently, the following "envy" test will necessarily be satisfied. No one will want to trade his bundle of resources for anyone else's bundle, because he could have had that other bundle in place of his own if he had so wanted. Because the result is an envy-free distribution in that sense, the strategy treats everyone with equal concern. Each person understands that his situation reflects that equal concern: his wealth is a function of what others want as well as what he wants. The strategy also respects the personal responsibility of each bidder for his own values. He uses his clamshells to acquire the resources that he deems best suited to the life he deems best. He is limited, in designing that life, by what he discovers are the choices of others, and therefore of what he can have available for whatever life he designs. His choices are not limited by any collective judgments about what is important in life, but only by the true opportunity costs to others of what he chooses. (I discuss the nature of true opportunity costs and their role in establishing a theory of justice, together with Samuel Freeman's comments on that role, in a note.<sup>10</sup>)

The fantasy distribution respects both our principles: it provides attractive conceptions of both equal concern and full respect. But you and I are not shipwrecked passengers on a newly discovered and abundantly stocked island. How far and in what way can we be guided by the fantasy in the very different situation of modern economies? The story has an immediate negative lesson. A command or socialist economy in which prices, wages, and production are set collectively by officials would be a very imperfect realization of our values. The decisions of a command economy are collective: they

reflect a collective decision about what ambitions, and hence which resources, are best suited to a good life. A free market is not equality's enemy, as is often supposed, but indispensable to genuine equality. An egalitarian economy is a basically capitalist economy.

That bald claim must, however, quickly be qualified in two crucial ways. First, it is essential to the justice of the island auction that the price someone pays for what he acquires reflects the true opportunity cost to others of his acquiring it, but actual markets in capitalist economies are often corrupted in ways that defeat that condition. Regulation is therefore often needed to perfect the freedom or efficiency of a market: to protect it against distortions of monopoly or externality. These distortions include (as we have recently come to learn) exaggerated risk in search of exaggerated profit when the risk falls largely on those who took no part in the decision and would have little share in any gain. Climate impact is another important example of distortion: because the market cannot easily be structured to reflect the opportunity costs of energy consumption now to future generations, extra-market regulation seems necessary. These adjustments to a free market do not contradict the spirit of this understanding of equal concern; on the contrary, they enforce that understanding by better matching people's resources to the true opportunity cost of what they do or consume.

The second qualification is very different and must occupy us at some length. The fantasy auction scheme shows equal concern, I said, because the result satisfies the envy test I described. What each islander has is fixed by his own choices, given the choices others make from an equal base. When the auction is finally over, however, and the islanders begin their economic lives, the envy test soon fails. They plant, manufacture, and consume using the resources they acquired in the auction, and they enter into transactions with one another, each trading to improve his situation. Some of the differences this activity generates reflect their choices—to consume rather than save, to rest rather than work, or to produce poetry that others do not much want rather than corn, which is popular. The envy test is still met in spite of these differences if we apply that test over time: people's resources continue to be sensitive to their choices. But other differences do corrupt the envy test. Some islanders do not have much talent to produce what the markets value, or they fall sick, or they make responsible investments that nevertheless fail. They then have fewer resources with which to build a life, not as a consequence of, but in spite of, the choices they have made. Now the envy test fails

because their resources do not depend, after all, only on their choices. The market is no longer egalitarian.

*Ex Ante or Ex Post?*

How should we respond? Runners in a fair race are equally placed, all at the starting line, before the race begins. They are *ex ante* equal. But they are not equally placed after the race has been run: *ex post* one has beaten the others. Which is the right temporal focus for justice? Does equal respect require trying to satisfy the envy test, so far as we can, *ex ante*, before the impact of transactions and luck? Or *ex post*, after those events have run their course? A government committed to *ex post* equality undertakes, so far as this is possible, to bring citizens who lack market skills to the same economic level as those with more skills and to restore those who have fallen ill or suffered handicaps to the position they would otherwise have occupied. A government that aims at *ex ante* equality, on the other hand, responds differently. It aims that its citizens face these contingencies in an equal position; in particular, that they have an opportunity to buy appropriate insurance against low productive talent or bad luck on equal terms.

At first blush *ex post* compensation might seem the more appropriate goal. People who are unemployed or who are badly injured or crippled and who receive only what an insurance policy might pay by way of compensation remain in a much worse position than others. Insurance payments typically do not compensate fully, and for some instances of bad luck—terrible physical disability—they fall sickeningly short of restoring people to their prior position. So long as the community can improve the situation of someone who has been a victim of bad luck, then equal concern might seem to require that it do that. In fact, however, the *ex post* approach, even so far as it is possible, is a very poor understanding of equal concern. The *ex ante* approach is better.

Investment luck, very broadly understood, is an important reason why people's income and wealth differ. You and I study financial charts with equal care and make equally intelligent though different choices. Your stocks thrive and mine wither; you are rich and I am poor and this is only because your luck has been better than mine. But if our political community undertook to erase this consequence of luck, it would undermine the responsibility each of us exercised; if it made our investment choices pointless in that way, we would cease to invest. Many of the most important decisions we make are

also investment decisions whose consequences turn on luck: any educational or training decision might be undermined by unforeseeable technological shifts that make our particular training useless, for instance. If the community aimed to insure that our fate in no way depended on how any such investment gambles fare—if it guaranteed that we are equal in wealth, whether or not our choice of career turned out to be suited to our tastes or talents or market conditions—it would end by crippling our own responsibility for our choices. So any plausible version of an *ex post* approach would have to draw a distinction between investment and other forms of luck and rule out the former as a ground for redistribution.

That distinction would be difficult to draw. But *ex post* compensation would be not a reasonable goal, even if restricted to noninvestment luck. Any community that undertook to spend all it could to improve the position of its blind or crippled members, for example, until further expenditure would not even marginally benefit them, would have nothing left to spend on anything else, and the lives of all other citizens would be miserable in consequence.<sup>11</sup> That policy would reflect no one's actual priorities, including the antecedent priorities of the victims of terrible accidents. If the choice had been up to them before they were injured, they would not have spent everything they had to buy the best possible accident insurance policy, because they would not have thought, given the odds, that it made sense to compromise their lives in every other respect to secure the most expensive possible insurance. The *ex post* compensation approach to bad luck is irrational.

It remains wrong even if we apply it to erase the consequences, not of bad luck as ordinarily understood, but only of the bad genetic luck of not having talents prized in the contemporary market. If the community restores people to a condition of equal wealth, no matter what choices they make about work and consumption, then, as I said, it destroys rather than respects this dimension of responsibility. But there is no way fully to erase the consequences of differences in talent without adopting that foolish remedy. It is impossible in principle, not simply practically impossible, to distinguish the consequences of choice and capacity across the range of economic decision, because preference and capacity interact in both directions. Our preferences both shape the talents we are disposed to develop and are shaped by the talents we believe we have. So we cannot separate choice from genetic luck in what might seem the most direct way: by making sure, *ex post*, that people's wealth reflects only the former and no tinge of the latter.

Equal concern does indeed require that a community compensate in some way for bad luck. But we need an understanding of compensation that is compatible with the right respect for individual responsibility, and we must therefore seek an *ex ante* approach. This aims, as I said, to situate people equally as they face both economic decisions and the contingencies that hedge those decisions. An economic market for investment, wage, and consumption is a crucial step toward that equality, because it allows people's decisions to carry costs or gain rewards that are measured by the impact of those decisions on other people. But we need a further step: we need to place people in the position they would have occupied if they had been, at a point before the decisions and events that shaped their lives began, equally able to protect themselves against these different dimensions of bad luck through appropriate insurance. That step unfortunately requires the kind of fantasy speculation that I referred to earlier. For of course it is impossible that people could ever be equally able to insure in any real insurance market; certainly impossible before their genetic luck begins, because before that point they do not even exist.

#### *Hypothetical Insurance*

We must return to our island. Now we notice that insurance is among the resources auctioned. Some islanders undertake to insure others, in competition with other insurers, at market clamshell rates. When the auction ends, *ex ante* equality has been preserved and future transactions maintain it. How does this expanded story help us? It teaches us the importance of the following hypothetical question. What level of insurance against low income and bad luck would people in our own actual community buy if the community's actual wealth was equally divided among them, if no information was available that would lead anyone or any insurer to judge that he was more or less at risk than others, and if everyone otherwise had state-of-the-art information about the incidence of different kinds of bad luck and the availability, cost, and value of medical or other remedies for the consequences of that bad luck?

We can sensibly speculate about answers to that question from information readily available about what kinds of insurance insurers actually do offer and people actually do buy. Of course, there must be a large range of uncertainty in any answer we give. We cannot specify any particular coverage level that we can be confident any specific number of people would buy under the fanciful

counterfactual conditions we imagine. But that need not be our aim. We can try to identify a top coverage level at which we can sensibly assume that most people in our community would have chosen to insure, given what we know about their needs and preferences, and given the premium structure that that coverage would require. We cannot answer even that question with any pretense to exactitude. But we can dismiss some answers as plainly too low. We can identify a coverage level such that it would be foolish for most people, given their preferences as we can ascertain them, not to buy coverage at that level.

We can then insist that our officials use at least that coverage level as a guide to redistributive programs of different kinds. We might aim to collect from the community, through taxes, an amount equal to the aggregate premium that would have been paid for universal coverage at that level and then distribute, to those who need it, services, goods, or funds that match what that coverage would have provided them in virtue of their bad luck. We would fund unemployment and low wage insurance, medical care insurance, and social security for people in retirement in that way. It is important to notice that by hypothesis any community can afford the programs that this insurance scheme describes: those programs would not be irrational in the way those mandated by a goal of *ex post* compensation would be. On the contrary, because the programs the scheme identifies reflect reasonable assumptions about the overall preferences of the community over risk and insurance, a government that did not provide them would fail in its economic responsibilities.

#### *Paternalism?*

Our overall ambition, remember, is to provide a scheme of distributive justice that satisfies both principles of dignity. It might now be objected that the hypothetical insurance scheme I just summarized offends the second principle because it is, in effect, mandatory. (Arthur Ripstein offers this objection and another concern.<sup>12</sup>) The scheme assumes that most citizens would have purchased insurance at least at the coverage levels and for the premiums it stipulates. But perhaps some would not, and taxing those citizens under the scheme (or indeed, awarding them benefits under it) is therefore, according to this objection, a paternalistic imposition of a supposedly reasonable choice upon them.

The point calls for further explanation, but the objection is not yet well framed. Paternalism means imposing a decision on someone supposedly for

his own good but contrary to his own sense of what that is. The hypothetical insurance scheme makes assumptions, on the contrary, about what citizens' preferences would have been in circumstances very different from those anyone has actually encountered. It is no more paternalistic to assume, for any individual, that he would have chosen to buy the insurance at what we judge to be a level at which most people would have insured than it would be to suppose that he would not have bought that insurance and to treat him accordingly.

So the scheme is not paternalistic. But it is probabilistic. No one can sensibly think or argue that he would not have made the decision we assume most people would have made. The counterfactuals are too deep for any such individualized judgment: the scheme's claims can only be statistical. But he can rightly say that he might not have made it. That fact presents an issue not of paternalism but of fairness. We can treat individual citizens on either of two assumptions, and it seems fair to treat them, lacking any information to the contrary, as if each would have done what we judge most would have done.

This is our justification. We aim to charge people the true opportunity costs of their choices. Though we must rely on actual markets in production and wage, we must supplement and correct those markets in a variety of ways. In particular we must try to eliminate the effects of bad luck and other misfortunes by judging what a more comprehensive and fairer market would have revealed as the opportunity costs of provision against those misfortunes. We must make probabilistic counterfactual assumptions in that exercise, to be sure. But that seems fairer than the alternatives, which are either to leave the misfortunes uncorrected or to choose some level of redistributive transfer payments through politics guided only by raw reactions of fairness that have no ground in theory and are likely to be stingy in practice. We choose the hypothetical insurance device, even though it requires rough judgments of probability, as more faithful to the overall opportunity costs conception of fairness. That is the best we can do to show equal concern and the right respect for individual responsibility. Our overall interpretive project endorses a redistributive scheme modeled on hypothetical insurance assumptions for that reason. (Amartya Sen offers a number of further objections to the hypothetical insurance scheme.<sup>13</sup>)

### *Laissez-faire Again*

That completes my summary sketch of a design for a political settlement that merges equal concern by government and personal responsibility for citizens.

(I have elsewhere described in much more detail the tax structure that this exercise would generate and the social programs it would justify.<sup>14</sup>) But we must take care not to confuse our *ex ante* approach, which features *ex ante* compensation, with a different *ex ante* approach—misleadingly called equality of opportunity—that is popular among political conservatives. This holds that we show equal concern by letting the chips fall where they may: it allows no redistribution of market rewards and insists those who have bad luck must bear it themselves. This is just a form of the *laissez-faire* doctrine I mentioned at the outset of this discussion. Proponents say that *laissez-faire* rewards individual responsibility. But people with little market talent or bad luck can reply that it does not show equal concern, because a different economic arrangement is available that also satisfies the requirements of individual responsibility and that shows more appropriate concern for them.

Equality of resources, understood as I have described it, may reward qualities of productive intelligence, industry, dedication, shrewdness, or contribution to the wealth of others. But that is not its aim. It does not even suppose that these are virtues; it certainly does not suppose that a life earning more money is a better or more successful life. It presumes only that we treat people with equal concern when we allow each to design his own life, aware that his choices will have, among other consequences, an impact on his own wealth. However, it is crucial to this understanding that the character and degree of that impact reflect the effect his choices have on the fortunes of others: the cost to others, in lost opportunities for themselves, of the various decisions he has made.