

# Part 1

---

## *Social Justice*

The idea of distributive justice has been around for a very long time – the Greek philosopher Aristotle (384–322 BC) wrote about it. *Social* justice is different. That idea is relatively recent, creeping into use from about 1850 on, and not everybody likes it. It developed only as philosophers came to see society’s key social and economic institutions, which crucially determine the distribution of benefits and burdens, as a proper object for moral and political investigation. Some philosophers aren’t happy with it. People can act justly or unjustly, but what does it mean to say that *society* is just or unjust? Some politicians aren’t crazy about it either. For them, those who talk about social justice tend to hold the mistaken belief that it is the state’s job to bring about certain distributive outcomes, which means interfering with individual freedom and the efficient working of a market economy. (To get a common confusion out of the way, let’s be clear from the start that social and distributive justice are usually regarded as different from retributive justice. That is concerned with the justification of punishment, with making the punishment fit the crime. So we’re not going to be dealing with the kind of justice administered by the criminal justice system, the kind where we would talk about ‘miscarriages of justice’.)

Given that it is controversial, and relatively new, wouldn't it make more sense to begin with liberty, or community – ancient ideas that everybody values? I start with social justice for two reasons.

First, and most important, most political philosophers would say that it was the publication of a book on social justice – *A Theory of Justice* (1971) by the American philosopher John Rawls (b. 1921) – that transformed and revived their discipline. I would agree with them. For many years before Rawls, academic political philosophy was either the history of political thought or quasi-technical linguistic analysis of the meaning of political concepts. Since Rawls there has been systematic and substantive argument about what the societies we live in should actually be like. ('Substantive' means 'to do with substance or content, not just form'.) Much of what has been written since then can helpfully be understood as engaging with Rawls's theory – like it or not, those writing in his wake have to think about how their arguments relate to his – so it makes sense to lay out the basics of his position right at the beginning. His theory invokes and incorporates ideas of liberty, equality and community. These concepts are all closely interrelated, and thinking about his approach to justice provides the most convenient way in.

Second, one of Rawls's most famous claims is that 'justice is the first virtue of social institutions'. That is debatable, as we shall see: one might judge that other goals, goals that conflict with justice, are more important. But it is at least quite common for people to believe that other goals can only be pursued to the extent that that pursuit is compatible with the claims of justice. Think about the situation where one can make a lot of people very happy by killing an innocent man. (Suppose they mistakenly think he is guilty and that's why they would be happy.) Most people feel that to do that would be wrong, because the most important thing is not to treat people unjustly. Something similar underlies the thought that it is better to let the guilty go free than unjustly punish the innocent. On this kind of view, justice is a constraint on what we can do. It doesn't tell us everything – remember we are talking about the virtues of social institutions, not the virtues we might exemplify in our individual lives. But it

does tell us what must be our top priority when it comes to deciding the rules we are going to live under.

### Concept v. conceptions: the case of justice

Let's begin with an elementary but very useful analytical tool: the distinction between a concept and the various conceptions of that concept. Much confusion can be avoided by holding on to this distinction, which applies to many political concepts, not just those discussed in this book (e.g. 'democracy', 'power'). With this clearly in mind, it gets a lot easier to see what is going on in political debates where, typically, those on different sides use the same word to mean things that, when probed, turn out to be rather different. Understanding how they differ, and what underlies the disagreements, is the first step towards deciding which side is right.

The 'concept' is the general structure, or perhaps the grammar, of a term like justice, or liberty, or equality. A 'conception' is the particular specification of that 'concept', obtained by filling out some of the detail. What typically happens, in political argument, is that people agree on the general structure of the concept – the grammar, the way to use it – while having different conceptions of how that concept should be fleshed out. Take the case of justice. The basic concept of justice is that it is about giving people what is due to them, and not giving them what is not due to them. (This, at least, is how a lot of people think about it, though it is true that there might be disagreement even about this. I don't want to get on to that, more properly philosophical, terrain.) What is due to them. Not what it would be nice for them to have. Not what it would be polite to give them. Not even what it would be morally good to give them. (I'll explain this one in a minute.) What they have as their due.

This analysis, then, ties justice to duty – to what it is morally required that we, perhaps collectively through our political and social institutions, do to and for one another. Not just to what it would be morally *good* to do, but what we have a duty to do,

what morality *compels* us to do. And, of course, there are many different conceptions of this concept, because people who agree that this is what ‘justice’ means, as a concept, can still endorse different conceptions of justice, can (and do) disagree about what justice ‘means’ in terms of the content fleshing out the grammar of that term. This part of the book will say a bit more about the overarching concept of justice, and then lay out three influential conceptions – Rawls’s justice as fairness, Robert Nozick’s justice as entitlement, and the conception of justice as desert. Most people endorse bits of all three. Sometimes this is done in an informed self-reflective way that has worried about whether the overall package of beliefs about justice is consistent (for there are ways of combining elements of these – and other – conceptions into a coherent whole). More often, however, it happens unthinkingly, in a way that turns out, on inspection, to contain a deal of confusion.

Back to the concept of justice. There might be things it would be morally good to do that aren’t requirements of justice. Think of justice as a specific subset of morality. If Rawls is right that justice is the first virtue of social institutions, then that means that the most important set of moral considerations relevant to politics and the organization of society is that which concerns giving people their due. And what is due to people has a good deal, though not everything, to do with what they have a right to. That’s why justice and rights are so closely connected. Consider the contrast between justice and charity. One might think it was morally good to give charitably to those in distress without thinking that it was a requirement of justice. Indeed, if one thought of oneself as giving *charitably*, then one would precisely not be thinking of one’s act as a requirement of justice. (Of course you might give to particular needy individuals or organizations calling themselves ‘charities’ because you felt that their claims on you were indeed claims of justice, but then you would not be giving charitably.) It is quite common, I think, for people to regard their reasons for helping those who are starving in far-off countries as reasons of charity, or as deriving from a principle of humanity (say, a concern and respect for fellow human beings), but not as reasons of justice. We ought to help them in times of

need, it is morally praiseworthy to do so, and the reasons to do so are moral ones, but there is no duty to do so, for their claims on us are claims of common humanity, not claims of justice. The same kind of thinking is applied by some – such as the libertarian Nozick, whose views we’ll examine shortly – to our obligations to help needy members of our own society. It’s a morally good thing to do, but justice is about protecting legitimate property rights and it should be up to the individual to decide whether to help or not.

This brings us to the big reason why the distinction between justice and other kinds of moral claim is typically seen as so important. The state is justified in making sure that people carry out their duties to one another. It is justified in using its coercive power to force people to do what they might not do voluntarily. This is a big deal. As I said in the introduction, the state, as political philosophers think about it, is not something separate from and in charge of those who are subject to its laws. It is – or should be – the collective agent of the citizens, who decide what its laws should be. So to say that the state is justified in forcing people to comply with their duties is to say that citizens are justified in using the coercive apparatus of the state (laws, police, courts, prisons) to force one another to act in certain ways – including ways that some citizens might believe to be wrong. This, of course, raises big and difficult issues to do with the justification of state authority and whether, or in what circumstances, individuals are obliged to obey (and perhaps sometimes to disobey) laws they disagree with. Fortunately, this book is not about those big and difficult issues. What matters here is the significance of justice, given a common and plausible view of what the state can and cannot make people do. If you think that the state can justifiably force people to be charitable to one another, you are guilty of conceptual confusion. But thinking that the state can justifiably force people to carry out their duties to one another is, for many, part of the point or significance of the concept of duty. So justice is central to political morality, because of the widely held claim that once we know what our duties are to one another then we also know when we can justify using the machinery of the state to get

people to do things they might not otherwise do, and might even regard as wrong.

Clearly, if justice is about identifying the scope and content of coercively enforceable duties, or if we think that by definition the duties that arise are coercively enforceable, then it becomes particularly important correctly to identify the scope and limits of justice. And it's not surprising that there are big disagreements about that scope and those limits. Everybody will agree that it is legitimate for the state to (try to) enforce the law against murder. We all have a duty not to murder one another, and a duty to do what we can to prevent people performing the unjust act of murdering others. That some people might want to murder others, or might disagree that they have a duty not to, is neither here nor there. But claims about social or distributive justice go way beyond this kind of claim, in terms of the extent of the duties they imply. Do talented, productive people have a duty to forgo some of the money they earn to help those less fortunate than themselves, a duty, compliance with which we can – or even have a duty to – enforce upon them? Or is that properly a matter of charity – something beyond the realm of the state? The three conceptions of justice we will look at shortly give different answers to these questions.

Justice can be the first virtue without being the only one. This is an instance of a quite general point that it is always useful to keep in mind. Different morally valuable political concepts – justice, liberty, equality, democracy – need not coincide completely. This is a hard thing for politicians to accept, since they tend to be reluctant to acknowledge that their preferred policies or positions might involve anything other than the complete and harmonious realization of all good things. You don't often find a politician being honest enough to say something like: 'I believe in social justice of type  $x$ . I accept that this involves significant restrictions of individual freedom, that it does not provide anything I could honestly call equality of opportunity, and that its realization requires substantial limitations on the scope of democratic decision-making. Nonetheless, here are my reasons for believing in it.' Why not? Because their opponents would make a big fuss about the loss of freedom, the lack of equality of

opportunity and/or the restriction on democracy – each of which would doubtless be described in terms much more confused and vague than they intended. Compared to real politicians – who have to worry about how their statements will be interpreted, twisted, used and abused rhetorically, and spun – political philosophers have it easy. They can say precisely what they mean, with a reasonable degree of confidence that they will be taken as meaning precisely what they say.

This point about conflicts between political values should not be misunderstood. Of course, our aim is indeed to achieve the best reconciliation possible – in the sense of coming up with an overall position which does the best job of giving proper weight to these differing values. Of course there are different conceptions of the various concepts in question, and which conception we favour may in part reflect our other value commitments, which will in turn influence our preferred conception of another concept. We may well have an overall vision about how society should be that informs the way we think about all of them. But none of this means that we should start by simply assuming that, since equality and liberty or justice and democracy are good things, we must be looking for a way of thinking about these concepts which avoids the possibility of conflict between them. On the contrary, clarity is best achieved by keeping concepts as distinct as possible, resisting the temptation to let them melt into one another.

The most common example of confusion on this issue concerns the idea of democracy, a concept with such positive connotations that it is typically stretched in all sorts of directions. Who will confess to not being a democrat? But democracy, at core, is to do with the people as a whole having the power to make decisions about the rules under which they are going to live. This, on the whole, is a good thing – for lots of reasons. Who is more likely to make good rules than those who have to obey them? Rules restrict people's freedom, but those restricted by rules they have themselves been involved in making retain a kind of freedom – at least when compared with those subject to rules made by others. It's fair – it treats citizens as political equals – if rules are made by citizens as a whole rather than by some subset of the population.

It's good for people's characters and personalities that they should take an active role in the public life of their political communities. These are four, different, weighty reasons that do indeed make a very strong case for democracy. Others could be added to the list. But even the weight of these combined does not mean that democracy is always a good thing, or that all good things must, because they are good, therefore be 'democratic'.

To think that a decision should be made democratically is to think that it should be made by the people as a whole. Do we really want all decisions to be made this way? Aren't some decisions better regarded as private, better left to individuals than to the political community? Imagine two societies. In one, there is a democratic vote on what religions people are to be permitted to practise. In the other, there is a constitution which grants every individual the right to practise the religion of her choice. Which society is better? The second. Which is more democratic? I think the first. To be sure, *some* individual freedoms can be regarded as necessary for democracy itself. Freedom of association or freedom of expression are like this. If a society denies its members the right to say what they think, or to get together with others who agree with them, then we may well judge that it is denying them things that are needed for that society to be regarded as democratic. This is because of the connection between expression, association, and political activity. So some constitutional rights may be necessary conditions of, not constraints on, democracy. But is freedom of religion like this? Suppose a society doesn't prevent would-be followers of a religion from putting the case for why they should be allowed to practise it, or from organizing with would-be co-religionists to advance their cause. It simply prevents them from practising it. Is there anything that should be called *undemocratic* about this? Or what about freedom of sexuality? One might well think freedom of sexuality to be a central human freedom. A society which allows its members to do what they like sexually – as long, of course, as they don't harm others – is, other things equal, better than one that doesn't. But I don't think we should say that it is also a more democratic society. In fact, we should say that it is *less* democratic. It removes an issue from the scope of democratic control.



If we judge that the individual has a right to freedom of religion, or of sexuality, then these freedoms can be regarded as central to social justice. A society which denies them treats its individual members unjustly – being willing to violate people’s rights and to impose the will of the majority on a matter that should be left to the individual. There is, then, plenty of room for conflict between justice and democracy. Both are good things. We are ultimately going to be looking for the best balance between the different values that they embody. But we are not helped in thinking about the real issues by the misguided idea that the two concepts must coincide. On the contrary, we make intellectual progress by focusing precisely on the places where they come apart.

A society could be perfectly just – everybody is getting what they have a right to and all are acting dutifully towards one another – without its being a perfect society. Perhaps the vast majority of its members are bored (or, worse, *not* bored) couch potatoes, spending vast amounts of their time watching daytime TV. Justice is one dimension along which we can judge societies as better or worse than one another, but it is not the only one. It matters also how people live their lives *within* the social institutions that embody principles of justice – what they choose to *do* with their various rights and their just share of goods. Where things get interesting, of course, is where we think that justice and other good things are in some sense competitive with one another. Then it really does matter whether we agree with Rawls about justice being the first virtue. There is a famous climactic scene on the big wheel in the classic movie *The Third Man*, where Orson Welles, as Harry Lime, sketches the relative merits of Switzerland and Florence under the Borgias. Florence was savage and violent – not much social justice there – and it gave us the Renaissance. Switzerland has been a model of peace, fair-mindedness and social solidarity – and it gave us the cuckoo clock. Lime’s thought, of course, is that this is not coincidence. It’s not simply that there are more good things than social justice, but, worse, that social justice is actually inimical to some good things. Justice, from this perspective, can start to seem a rather tedious, tame virtue. A virtue, to echo the German philosopher Friedrich Nietzsche

(1844–1900) fit for slaves, not for people capable of actions nobler and more heroic than the petty, cowardly concern to treat one another justly.

The idea that justice might be inimical to excellence has other, less drastic, incarnations. Some defences of inequality appeal not to the idea that inequality is just, but to the claim that disproportionately concentrating resources in the hands of the few is a necessary precondition for intellectual or artistic progress. Alexis de Tocqueville (1805–59), the French aristocrat who wrote about democracy in America, thought that the system whereby estates were divided equally between sons rather than passing intact to the first, as happened in France, meant that America would necessarily produce fewer, perhaps no, great thinkers. Great thinking requires people with leisure and an aristocratic culture committed to the cultivation of the intellect so that, for example, children are not expected to pay their way but rather devote many years, perhaps their whole lives, to the acquisition of intellectually valuable but financially useless skills. America's commercial culture and society of misters, though better in many respects, and, for Tocqueville, overall, was bound to lead to a kind of intellectual mediocrity. Similar arguments abound today. Is it right to spend large amounts of public money subsidizing cultural activities, such as opera, that tend disproportionately to be valued by the better off – especially if, as is the case with the UK's National Lottery, the money is disproportionately raised from those who are less well off? Can the British universities of Oxford and Cambridge justify the claim that the state should provide any of the extra resources required by their labour-intensive tutorial teaching methods – especially if it is children of the better off who are disproportionately likely to receive such an expensive education? We are surrounded by what, at least at first sight, are hard choices between social justice and other values.

## Hayek v. social justice

According to Friedrich von Hayek (1899–1992) the very idea of social justice is a ‘mirage’, or the kind of confusion that philosophers call a ‘category mistake’. Hayek, an Austrian, was Prime Minister Thatcher’s favourite intellectual, and a major influence on the development of the New Right in Britain and the US during the 1970s and 1980s. In his view, the idea that ‘society’ is something that might be just or unjust involves a misunderstanding of the concept of justice. Justice is an attribute of action, a predicate of agents. A person acts justly when she undertakes a just action. The aggregate distributions of resources that result from individuals interacting in the market are unintended by any individual agent, and therefore not susceptible of being judged just or unjust. The idea of ‘social justice’ involves a fundamental failure to see this point. ‘Society’, not being an agent, is not the kind of thing that *can* be just or unjust.

Hayek says other influential things too. He thinks any coercive redistribution by the state beyond the meeting of common basic needs involves an unjustifiable interference with individual liberty. The title of his most famous book, *The Road to Serfdom* (1944), conveys the key idea. For Hayek, the state’s ambition to realize ‘social justice’ implies a centralized authority making people do things they might not want to do, interfering with their freedom to do what they like with their resources – and all this in the name of a conceptual confusion! Relatedly, Hayek thinks that state policies in the area of welfare and redistribution necessarily involve the state making judgements about the criteria that should govern distribution. Should goods be allocated on the basis of need or merit? If merit, what counts as merit? And so on. Hayek is a sceptic on these matters. He is doubtful that there are right answers to such questions and thinks that the only thing to do is to leave judgements of this kind to individuals. Finally, Hayek thinks that, just as long as the state doesn’t stick its nose in and distort the process, individuals interacting freely will produce a ‘catallaxy’ or spontaneous order that crystallizes the information and wisdom dispersed in their individual heads. The free market

represents such a catallaxy – with the price signal supplying knowledge of a kind in principle unavailable to any central planner, and guiding individuals towards economic activity conducive to the general good. This critique of the planned, socialist economy – a variant of the Scottish economist and philosopher Adam Smith’s (1723–90) ‘invisible hand’ defence of the market – means that, for Hayek, attempts to plan the economy, or to redistribute resources in pursuit of particular distributive goals, are not just invasive of individual freedom, they also amount to inefficient distortions of market processes which, left to themselves, would tend, in the long run, to benefit everybody.

These are all big and controversial claims – too big to discuss here. But it is worth saying something about Hayek’s distinctive rejection of social justice as a mirage. To begin with, even if it were true that nobody intended the overall distribution of resources that results from the market, it doesn’t follow that nobody is responsible for it. People can be responsible for outcomes they don’t intend. Think of the man who fails to check his brakes and, as a result, runs over somebody. He didn’t intend to run anybody over but, because he could reasonably have been expected to have checked his brakes, he is responsible for having done so. He is negligent, culpably negligent. Now Hayek would say that there is no agent in the distributive case who can be held responsible, even in the sense of being negligent. But is that right? Surely we, as political actors, are capable of coming together and deciding that we are not prepared to permit certain kinds of distributive outcome – say that some members of our society, through no fault of their own, will live in poverty and without access to education for their children. If we accept that this is a matter of justice, not something that should be left to individual charity, then each individual is responsible for ensuring that she does her fair share of contributing to the prevention of that outcome, by agitating politically, and by bearing her share of the financial cost involved in its prevention. What matters is not whether anybody intends the injustice, but whether anybody is responsible for the fact that it exists. When governments devise their economic policies, they have a good sense of the distributive outcomes that will result. If they devise, and citizens vote for,

policies that can reasonably be expected to produce distributions that include avoidable and unjustified inequalities, then, whatever their intention, they are responsible for the existence of those inequalities. If those inequalities are unjust, then the act of voting for them is an unjust act. Hayek's attempt to sever the link between individual agency and aggregate distributive outcomes fails. He misses the fact that individuals can act politically, in concert with others, to prevent outcomes that, as individuals, may indeed be beyond their control.

### Rawls: justice as fairness

John Rawls has written two big books – *A Theory of Justice* (1971) being followed up by *Political Liberalism* (1993). These have a combined length of over 1,000 pages and goodness knows how many forests-worth of commentary and criticism they have jointly generated. A lot of attention has focused on whether and how Rawls changed his position between the two books, so answering the question 'What does Rawls really think?' is far from straightforward. In this section, concentrating on the first (though using elements of the second where that helps), I want to give the merest introductory sketch of what all the fuss has been about. More of Rawls's position will unfold as I compare it with the two other conceptions of justice – entitlement and desert – that come afterwards. (I will discuss *Political Liberalism* in part 4.)

The ideas at the heart of Rawls's theory of justice, which he calls justice as fairness, are the original position and the veil of ignorance. Rawls believes that the way to find out which principles of justice are fair is to think about what principles would be chosen by people who do not know how they are going to be affected by them. He thus imagines people choosing principles in an original position, behind a veil of ignorance. This is a thought experiment. The idea is to help us think about what would happen if people deprived of all knowledge that might serve to distinguish them from one another – such as whether they are clever or stupid, Muslim or atheist – were to get together and

decide how they wanted their society to be organized. Justice, for Rawls, should be understood as that which would emerge as the content of a hypothetical contract or agreement arrived at by people deprived of the kind of knowledge that would otherwise make the agreement unfair. The intuitive idea is the link between fairness and ignorance. If I don't know which piece of cake I'm going to get, I'm more likely to cut fairly than if I do. Depriving people of particularizing knowledge means that they will choose fair principles rather than allowing that knowledge to bias the choice of principles in their own interests.

There are two kinds of thing that the parties to this hypothetical contract don't know. First, they are ignorant of their talents – their natural endowments – and their social position. They don't know whether they are clever or stupid, or born into a wealthy or a poor family. Second, they don't know their conception of the good. They don't know what they believe about what makes life valuable or what is worthwhile (art, sport, watching daytime TV), whether they are religious or not (or, if they are, which religion they believe in) and so on. But there are some things they do know. Most importantly, they know that they have what Rawls calls 'the capacity to frame, revise and pursue a conception of the good'. Indeed, they regard this capacity as one of the most important things about them and are very concerned to protect it, and provide conditions for its exercise, when they engage in the process of deciding what principles should regulate their society. And they know that, to exercise that capacity, they need certain all-purpose goods, which Rawls calls 'primary goods': liberties, opportunities, powers, income and wealth, self-respect.

The original position, then, is a device of representation. It is a way of representing particular claims about how we should think about justice. Rawls's idea is that it models fair conditions by abstracting from people's natural endowments and social (class) position, and from their particular conceptions of the good. It models conditions under which people solely regarded as free and equal are to agree what he calls fair terms of social co-operation. Society, for Rawls, should be understood as a fair scheme of cooperation between free and equal citizens, and the original position models or represents that understanding.

One way of thinking about what is happening in Rawls's theory is that he is attempting to model – to capture by means of a thought experiment – what kinds of reasoning are and are not acceptable when it comes to thinking about justice. Suppose you met someone who favoured low tax rates and minimal welfare provision. You ask her why, and she says that, as a very talented businesswoman with children at expensive private schools, she and they would be better off in such a society. She might well be right about that. But it's hard to see how she could seriously present these reasons as having anything to do with justice – at least not if justice has anything to do with fairness. (There are other kinds of reason she could give which would, but we'll come to those later.) Doesn't she think about all the untalented people, or children whose parents cannot afford to send them to private schools? Doesn't it occur to her that she is lucky to be talented, that she might just as well have been born untalented, and that justice is about seeing things impartially, or from everybody's point of view? The Rawlsian way to do this is to imagine what distributive principles you would have reason to endorse if you didn't know who you were, thereby thinking of yourself and your fellow citizens as equals.

So ignorance about talents and social background models the sense in which people are conceived as equal. It is ignorance of their conception of the good which models the sense in which people are conceived as free. For Rawls, reasons arising from conceptions of the good should be kept out of the process of thinking about justice because allowing them in would imply not respecting people's freedom, spelled out as their capacity to frame, revise and pursue their own conception of the good. Suppose you are a Christian, the kind of wholehearted Christian who believes yours to be the one true faith. You might think that it would be a good idea for the state officially to endorse Christianity: to give it favoured status in schools, to allow only Christians to hold certain public offices, to protect it and not other religions from blasphemy. But, for Rawls, this would be to bias the state, which is the collective power of free and equal citizens, in a particular direction, and that would be unfair to non-Christians. The only way to treat all citizens fairly is for the state not to take a view on

how people should lead their lives (the same applies to art, or daytime TV), respecting their freedom – their capacity to choose how they live for themselves. This restriction on the kind of reasoning that may legitimately be invoked when thinking about justice is modelled, in the original position, by people’s ignorance of their conception of the good.

So what principles does Rawls think people behind the veil of ignorance would choose? These:

- 1 Each person is to have an equal right to the most extensive total system of basic liberties compatible with a similar system of liberty for all.
- 2 Social and economic inequalities are to be arranged so that they are both (a) to the greatest benefit of the least advantaged, and (b) attached to offices and positions open to all under conditions of fair equality of opportunity.

(1) is the principle of equal basic liberties. This has priority over (2) which is concerned with social and economic inequalities and itself has two parts: (b), the principle of fair equality of opportunity, which has priority over (a), the difference principle. (It is mysterious why – and rather irritating that – Rawls lists these last two principles in reverse order. Perhaps he wants to keep his readers on their toes.) Taken together these mean that a just society will, first and most important, give each of its members the same set of basic liberties or rights – freedom of expression, of religion, of association, of occupation etc. Then, if there are social and economic inequalities, it will make sure that all citizens enjoy equality of opportunity in the process by which they come to achieve (and avoid) the unequally rewarded positions. Finally, it will only allow such inequalities at all if they tend, over time, to maximize the position of the worst-off members of society.

Would people in the original position really choose these principles? Many critics say that they wouldn’t. In particular, a lot of attention has focused on Rawls’s assumption – essential to the difference principle – that they would behave as if they were risk-averse, concerned to make the worst-off position as good as possible (or, in Rawls-speak, to ‘maximin’ – to *maximize* the



*minimum*) for fear that they might end up in it themselves. But why should they be quite so pessimistic? Wouldn't it be more rational to choose principles that would maximize the *average* position, perhaps subject to some 'floor' level beneath which they would indeed not want to take the risk of sinking? (Empirical simulations of the original position suggest that this is in fact what real people do choose.) Rawls has offered various defences of 'maximin' thinking, though he has tended to back off the initial suggestion that this would be the technically 'rational' way for them to proceed given the uncertainty they face. One argument – which invokes what he calls 'the strains of commitment' – goes roughly as follows: 'It matters that all those living in a society endorse it in a way that means they will be committed to it – rather than seeking to change things. If the difference principle is in operation, those who are at the bottom of the pile will know that the rules are working to ensure that they are as well off as they could be. So even they will be committed to the society.' (One obvious problem with this move is that somebody could accept that those who are worst off are as well off as they can be without accepting that she should be one of the worst off. In that case, she may not have the kind of 'commitment' that Rawls is looking for.)

Another focus of objection is 'the priority of liberty' – Rawls's view that the parties to the hypothetical contract would not be prepared to trade off the basic liberties for the sake of economic gain. (The kind of 'priority' given to liberty is very strict. It's not just that liberty is given greater weight in any decision about trade-offs, it's that there can't be any trade-offs.) Here Rawls would appeal to his claim about the importance of people's capacity to frame, revise and pursue their conception of the good, and the way in which the basic liberties are essential to the exercise of that capacity. Would you be prepared to take the risk of not being allowed to say what you believed, or of not being allowed to associate with whom you liked, or of being forced to practise a religion you thought was nonsense, in return for more money? Your answer will probably depend on how poor you would expect to be without the extra. If the choice were liberty or food, we would all choose food. Rawls accepts this, explicitly

acknowledging his assumption that everybody in society has reached a certain threshold of economic well-being. Only once we have reached that level do the basic liberties acquire their clear priority. (This in turn raises the question of how universally – to what range of societies – Rawls thinks that his theory applies. That’s a big and difficult one that would take us too far off the current track.)

It is the last principle, the difference principle, that has attracted most attention in debates about distributive justice. How *could* inequalities tend to maximize the position of the worst off? Isn’t the obvious way to do that to pay everybody the same? Rawls’s thought is the familiar one that people may need incentives if they are to be motivated to work in those activities where they are going to be useful. Some inequality, so the argument goes, is necessary (sociologists might say ‘functional’) if the economy is going to be as productive as it might be. Without inequalities, people will have no incentive to do one job rather than another – hence no incentive to do the kind of work which it is most useful (for everybody else) that they do. Imagine all those brain surgeons and dynamic entrepreneurs who would rather be poets. Without the extra money that will induce them to forgo the pleasures of poetry, the rest of us will be deprived of their surgical and entrepreneurial skills. Generalize to the aggregate level and you have an inefficient, stagnant economy which, because it pays everybody the same, does not provide the kind of growth that benefits everybody – including, over time, the worst off. This, so the argument goes, is roughly what happened under state socialism in eastern Europe.

This justification of inequality is very widely accepted. It has led some thinkers to conclude that there is no reason to worry about inequalities at all. If what matters is the absolute position of the worst-off members of society, then we should be prepared to countenance any inequalities that improve that position. There is, on this account, no need to ‘mind the gap’ between rich and poor – our attention should focus solely on whether the economy is organized in such a way that the poor are, over time, becoming better off. I will say more about this line of argument later on, in part 3 on equality. For now, it is worth pointing out that Rawls’s

principle says only that inequalities are justified *if* they serve to maximize the position of the worst-off. It is quite consistent with this that, in fact, no inequalities are justified (because it is not true that any are needed to maximize the advantage of the worst off). We should (and will) think carefully about whether they are needed, and if so, why. Notice also that the principle is demanding: inequalities are justified only if they serve to *maximize* the position of the worst off. The odd bit of 'trickle down' is not enough to satisfy the principle. What matters is whether the worst off are as well off as they could be, not whether they are better off than they might have been.

Another major source of debate has been who is to count as the 'worst off'. Rawls initially suggested that we measure how well off somebody is by seeing how many primary goods they have got. Those with least primary goods are the worst off. The problem with this is that it pays no attention to the process by which those with least came to have least. Suppose they are bone idle – people who started out with a fair amount of resources but chose to consume them rather than to work productively. After a couple of years they have nothing left and are now, by Rawls's original measure, the worst off. Does fairness really require the hardworking – and hence better off – members of society to channel resources in their direction? Seeing the problem, Rawls amended his position to recognize that 'leisure' might be included in the index of primary goods. We will return to this issue when we look at justice as desert, and again in part 3, when we consider whether those who are poor because they chose idleness really are worse off, all things considered, than those who chose to work hard and became rich.

To end this quick introduction to Rawls's position, a couple of thoughts about the 'contract' aspect of Rawls's argument. This can cause the kind of deep confusion that really gets in the way of understanding what he's up to. Rawls himself refers to the great tradition of social contract theory exemplified by the work of Thomas Hobbes (English, 1588–1679), John Locke (English, 1632–1704) and Jean-Jacques Rousseau (Swiss French, 1712–78). This is the tradition that thinks about social and political organization – law and state authority – as the outcome of an agreement

between individuals who see that they will be better off under law than they would be in the state of nature. Or, rather, it thinks about it *as if it were* the outcome of such an agreement. It's not at all clear that any member of the tradition really believes that there was a moment in history when the state and law emerged as the result of a contractual agreement. The key idea is rather that it might have done; that, whatever its historical origins, it is in people's interests to submit to it – they should go along with it because they would have agreed to do so (because the alternative is the state of nature). On this interpretation, then, it is not just Rawls's contract that is hypothetical – the contract tradition as a whole is most plausibly understood as positing a hypothetical contract, the point being that that helps us think about what we can properly expect people to go along with (on the grounds that they would have agreed to given the chance).

A common objection to Rawls is that hypothetical contracts, unlike real ones, have no binding force. They are, so the joke goes, not worth the paper they're not written on. But this misunderstands the role of the contract in his argument. If somebody asks 'Why should I go along with Rawls's principles of justice?' the answer is *not* 'Because you agreed to, and are therefore under a contractual duty or obligation to do so.' That, as the objection observes, is not true. The answer is rather 'Because you have a duty to act justly and Rawls has correctly identified what justice requires of you.' The hypothetical contract comes into the story only because it is, for Rawls, the right way to think about and identify what justice requires. If there were other, better, ways, then we should use them, and we would still be obliged to comply with the outcome. So it is not a contract argument in the everyday sense that people are bound to go along with the outcome because they agreed to it. The hypothetical contract is simply a device for thinking about what principles are indeed just, and it's because they're just that one is bound to comply with them, not because one agreed to them. (It's true that, for Rawls, the way to see that they are just is to see that we *would have* agreed to them *under appropriate conditions*, so it's not surprising readers get confused.)

The contractual aspect of the argument sometimes generates

another misunderstanding. The normal way of thinking about a contract is as something voluntarily entered into by people pursuing their own interests, for mutual advantage, and Rawls talks about the motivation of the people in the original position in a way that suggests that he sees them as essentially self-interested (or at least what he calls 'mutually disinterested'). Each is concerned to end up as well off as possible, to protect her own interests. Her thought is: 'What principles are going to be best for me given that I don't know who I'm going to be?' All this is true. But that doesn't mean that Rawls's theory is one for people who are ultimately, or in any overall sense, egoistic or self-interested. It is a theory for people who see society as a fair scheme of co-operation, who care about treating their fellow citizens fairly, and who regard them as free and equal. That is why they will accept the original position – with its equalizing and impartializing veil of ignorance – as the right way to think about justice. *Within* the original position, people are indeed regarded as choosing principles by looking out for themselves, by thinking about how they, as individuals, will fare under them. But the moral content is already there by then. It is there in the way that the veil of ignorance is set up in the first place. The parties to the hypothetical contract look out for themselves, one might say, only after they have been deprived of all information that might enable them to look out for themselves.

One often reads that the liberal approach to justice – and to politics in general – assumes that people are basically self-interested or egoistic. This view used to be common in Marxist writings and is now most prevalent in communitarian and feminist circles. (I will examine it in more detail in part 4, on community.) Certain aspects of Rawls's theory may have done something to encourage that misunderstanding. But it *is* a misunderstanding, and must be discarded before one can begin to see what Rawls is really about. Liberals like Rawls do care that individuals should be free to live the lives of their choice, but they care that *all* individuals should be free to do so and demand a fair distribution of resources for that reason. Moreover, the lives people choose can perfectly well include concern for others. It is hard to see what is self-interested about any of that.

## Nozick: justice as entitlement

The American Robert Nozick (b. 1938) was Rawls's colleague in the Philosophy Department at Harvard, teaching alongside him when Rawls published *A Theory of Justice* in 1971. By 1974 Nozick had published his counterblast, *Anarchy, State and Utopia*, which is still the most coherent and systematic articulation of libertarian principles around, and one of the most fundamental critiques of Rawls's whole approach. For Nozick, justice is not about agreeing fair principles by imagining that we don't know how lucky or unlucky we have been in the natural or social lottery. It is about respecting people's right to self-ownership, and their right to hold property, leaving them free to decide for themselves what they do with what is theirs. The proper role of the state, for Nozick, is not to meddle with the distribution of resources so as to produce some ideally 'fair' distribution. That would involve unjustified intrusions into people's legitimate holdings of private property. Its role should rather be limited to that of protecting people from such intrusions by others. Where Rawls is a 'left liberal' (or an 'egalitarian liberal') advocating a substantially redistributive welfare state, Nozick is a 'right liberal' (or 'libertarian'), committed to the idea of self-ownership and arguing for a laissez-faire 'nightwatchman' state. Like Hayek, his views – or at least versions of them as filtered through various think tanks and policy units – were influential in the development of the New Right. (It's not clear that Nozick continues to hold the views he endorsed in *Anarchy, State and Utopia* – he moved on to other philosophical areas soon after and has made only the occasional cryptic remark about it since then. So when I attribute arguments to 'Nozick', think of that as being his book, not the man as he now is.)

Nozick attributes to Rawls, and objects to, the view that we can regard goods as 'manna from heaven'. Were it the case that we had woken up one morning to discover that the world was suddenly full of things that people wanted, then it might be appropriate to adopt Rawls's or similar principles to distribute them. In that case, after all, why should anybody get more than

anybody else? But that is not how goods came into the world. They are made by people. They are the result of individual people's work, sometimes in co-operation with others. People create things by combining their own abilities and efforts with the natural world, entering into voluntary agreements with one another for the mutually advantageous exchange of such abilities and efforts, and the things that they thereby create are *theirs*. They are not like manna from heaven, unowned and up for distribution in accordance with fair principles. They come into the world already owned, by the people who produced them (or by those who have paid for the labour of those who produced them).

Rawls objects to utilitarianism because it fails to take seriously the separateness of persons. Maximizing overall happiness is a mistaken goal partly because there is no overall person to enjoy that overall happiness. There are just lots of separate people, and it would be wrong to make some unhappy for the sake of creating more happiness in some others. This thought underlies the idea of the contract, whereby principles have to be agreeable to each individual considered separately – which Rawls thinks will rule out principles aimed simply at maximizing overall utility (or overall anything else). What if I am one of the people made unhappy for the sake of other people's happiness? But Nozick thinks that Rawls does not take the separateness of persons seriously enough. Rawls does not see that we are individual, separate people, each with her own talents and attributes, which belong to her and her alone, and which may not be used to benefit others without her consent. She can choose voluntarily to give the fruits of her labour to others, but the state acts wrongly, failing to respect her separateness, when it forces her to give up some of those fruits to others. Nozick, then, opposes all redistributive taxation. If the wealthy are to give to the poor, they must do so voluntarily, not because the state forces them to.

In Nozick's view, people can do what they like with what is theirs. And there are three kinds of thing that might be theirs: (a) their selves – their bodies, brain cells, etc.; (b) the natural world – land, minerals, etc.; and (c) the things people make by applying themselves to the natural world – cars, food, computers, etc. I'll say something about the idea of self-ownership – that my limbs

and brain cells are mine to do what I like with – shortly. And once people own bits of the world, and own themselves, it's easy to see how they might be thought to own what they produce by bringing them together. So let's start by seeing how Nozick thinks bits of the natural world might come to be owned by people. He identifies three ways in which people can acquire a legitimate property holding (or entitlement): initial acquisition, voluntary transfer, and rectification.

Initial acquisition refers to the case whereby somebody comes to appropriate – to make their own property – previously unowned bits of the world. Imagine people settling for the first time an uninhabited continent. In Nozick's view the land and natural resources of that continent do not belong to anybody, and may legitimately be acquired by individuals on a first come first served basis, as long as nobody is made worse off by their doing so. (This is Nozick's variant on Locke's famous claim – in his *Second Treatise of Government* (1689) – that people may appropriate property just as long as 'enough and as good' is left for others.) This view has come under substantial and sustained criticism, and it would be fair to say that most political theorists think that Nozick's account of initial acquisition is inadequate. What exactly does one have to do to make previously unowned property one's own: walk round it, draw a circle on a map, put a fence round it? How do we decide whether others are being made worse off? They're clearly worse off in the sense that they are no longer able to appropriate that bit of land. And, in any case, who says that the continent was unowned – up for grabs – in the first place? Maybe it, and all the natural world, is jointly owned by all of us, in which case anybody wanting to use any of it needs permission from the rest of us. If the world were collectively or jointly owned, then it might look appropriate for us to get together and decide, collectively, how we want to use and distribute it – perhaps in accordance with Rawls's or other distributive principles.

For Nozick, however, the world is initially unowned and comes to be the private property of individuals through legitimate acts of initial acquisition. That is the first way to acquire property. The second way is by being given it by somebody who, by owning it herself, has the right to give it to you. Once somebody



owns anything, she can do what she likes with it, including, of course, giving it to whomever she likes, on whatever terms may be voluntarily agreed between them. This, for Nozick, is what happens in the market. I own my labour. You own some land (which you acquired, let's suppose, by an act of initial acquisition). We enter into a voluntary agreement whereby I sell – or lease you – the use of my labour for a certain price, thereby coming to own some money, which I can in turn do what I want with. So those of us who missed out on the initial acquisition stage – who came into the world when everything had already been snaffled up – shouldn't worry too much. We own ourselves and are therefore in a position to lease ourselves to others. If we're lucky, the selves we own may command a high price in the market, in which case we can lease ourselves for lots of money and ourselves come to own substantial amounts of property.

So the history of the world should be one of legitimate acts of initial acquisition followed by legitimate transfers of property, through acts of voluntary exchange, the result being the just outcome that people own exactly what is theirs and nothing else. But Nozick knows that it hasn't really been like that. He knows that the history of the world is actually one of unjust, involuntary transfers, whereby those with better weapons have forced those weaker than themselves to give up what – in his view – was rightfully theirs. The most familiar examples of this would be the way that white settlers treated the native populations of North America or Australia, but world history has really been one long sequence of such unjust transfers. Nozick's third principle – the third way whereby one can come to have an entitlement over property – is meant to deal with this. It is the principle of rectification, which holds that unjust transfers may be rectified by compensating transfers that themselves create entitlements. In practice, of course, as Nozick is well aware, the difficulties raised by this idea of rectification are horrendous. There is no way that we can identify who would own what if there had been no unjust appropriations, hence no way of rectifying properly. At one point Nozick suggests that the best thing to do might be to give everybody, as a starting-point, equal amounts of property – that might at least be a closer approximation to a just set of property

holdings than the vast and structural inequalities (inequalities between different ethnic groups, for example) that have been built upon those unjust acts of appropriation.

It would be a mistake, then, to see Nozick as an apologist for the status quo. He can perfectly well insist that existing inequalities are unjust, precisely because they have not come about in accordance with his three principles. That said, what is really significant about his position is that, on his view, vast and structural inequalities *could be* just. People own themselves, but the selves they own are going to be worth vastly different amounts to others. Some will be born strong, healthy and with high levels of natural ability. Others may be born weak, ill, and without even the potential to develop those attributes that others are going to be willing to pay for in the market. Some will be born to wealthy parents who can spend on education and bequeath their wealth to their children, and so on down the generations, with more and more advantage accruing all the time. Others may be born to parents in poverty, with no means of helping their children get a start in life. Nozick thinks that this is bad luck – he might even concede that it is unfair – but it is not unjust. As long as people’s property rights are respected, which means no coercive state action except that which is necessary for the protection of property rights (the nightwatchman or minimal state), whatever distribution results, however unequal it may be, is just. People can, of course, give voluntarily to those less fortunate than themselves. Nozick may well think that they ought to do so. But there is no justice claim involved – and no justification for coercive state action directed against the better off. Justice is simply about respecting people’s property rights, about leaving people free to do what they like with what is theirs.

Nozick describes his three principles as ‘historical’ and ‘unpatterned’. A summary slogan would be ‘From each as she chooses, to each as she is chosen’. The contrast is with ‘end-state’ and ‘patterned’ principles – principles that prescribe a particular state that must be realized (such as that inequalities are benefiting the worst off) or require distributions in accordance with a particular pattern (such as ‘to each according to her need’, or ‘to each according to her deserts’). On Nozick’s view, what matters is that

people have stuff that is justly theirs, and whatever distribution results from voluntary exchanges between them is necessarily just. Whether somebody has a justice claim to something depends solely on the chain of events that led to them having it. Inequality could be just, equality could be just. That depends simply on what it is that people choose to do with their property.

One way that Nozick formulates his objection to the redistributive state is that it uses some people as means to other people's ends. He thereby leans on the thought famously formulated by the German philosopher Immanuel Kant (1724–1804) that morality requires us to treat others not as means to our own or other people's ends, but as ends in themselves. Treating people as means seems like a fairly accurate description of what is involved when the state coercively redistributes resources from some to others. Not all taxation, of course, is used for redistributive purposes. Some of it pays for street lights, and the police, and defence. Some pays for a public education and health-care system from which those who are taxed themselves benefit. But some of it does involve involuntary transfers from some to others. When we tax people on their income part of what we are doing is using their productive abilities, which they might otherwise use solely for themselves, to help others. They may not be forced to work, or to do any particular kind of work – so Nozick's claim that taxation is akin to forced labour looks a bit over the top. But, if they do work, we are using them – some proportion of the exercise of their abilities and efforts – as means to other people's ends. Though true, it's not obvious that this is an objection. It might be wrong to treat people *solely* as means (which is what Kant actually said) – to be willing to enslave them and generally make their lives a misery for the sake of others. That might indeed fail to take seriously the separateness of persons, each of whom has her own life to live. But, if some people are lucky enough to be productive, and others unlucky enough not to be, one might think it justified to use the former to help the latter – even if they have not consented to that use. That will partly depend on whether, or in what sense, people own themselves, of which more shortly.

Another core Nozickian thought is the idea that 'liberty upsets

patterns'. Nozick's objection to patterned principles of justice – those holding that the justice of a distribution depends on whether or not it conforms to a particular pattern – is that the preservation of justice will inevitably involve restrictions, in his view unjustified restrictions, on people's liberty. This is the point illustrated by his famous 'Wilt Chamberlain example'. Wilt Chamberlain was, in 1974, a very high-earning basketball player in the USA, the Tiger Woods of his time and place. Nozick thinks that, if people are willing to pay a lot of money to see him play (and assuming the money they are willing to pay is money to which they are themselves entitled), then he is entitled to the money. The clever bit about the Wilt Chamberlain example is that Nozick allows us to imagine starting with whatever distribution of resources we like. Suppose we start with an equal distribution of resources. All members of society have exactly the same amount of money. Now some people so enjoy watching Wilt Chamberlain play basketball that they are willing to pay a bit extra to see him in action. So his club, as well as charging the normal ticket price, asks for an extra 25 cents specifically for Wilt. Millions of people watch him during a season, and he ends up a very wealthy man. There is no longer an equal distribution of resources, but nothing objectionable has taken place. People have simply freely chosen what they want to do with what is theirs. The general lesson is that liberty upsets patterns. If the initial distribution was just – whatever pattern it conformed to – then whatever emerges from voluntary exchanges must also be just. Any alternative conception of justice restricts people's freedom to do what they like with their just share of resources.

In its own terms, the Wilt Chamberlain example is very effective. If people really own property in such a way that it is theirs to do what they like with, then that must include it being theirs to give to others. If they want to give it to somebody else, like Wilt, with the explicit condition that it should thereby belong to him in the same way that it belonged to them (i.e. so that he could do what he liked with it), then it must be illegitimate for the state to come along and take any of it away for the sake of others. So anybody who wants to challenge the conclusion – that vast inequality could be just and that the state would be

acting wrongly if it engaged in any kind of redistributive taxation – must challenge the premise. She must deny that anybody ever owns things in the sense that Nozick requires. The force of the Wilt Chamberlain example comes from Nozick's saying that the initial distribution of resources can be whatever one likes – and showing that vast inequality may result even from an equal distribution. But this involves a sleight of hand. For Nozick assumes that the initial distribution, whatever it is, must be a distribution of full or absolute property rights: 'full or absolute' in the sense that they imply that people can do whatever they like with their property. If this were granted, the rest would indeed follow. Lots of the critical literature on Nozick's view is concerned to challenge the idea that we can ever have that kind of ownership claim over property. Ownership is a complicated idea. I can have the right to use my work room without having the right to bequeath it to my children. I can have the right to use the office's shared photocopier without having the right to sell that right to others. If people have absolute rights over what they produce, why can't parents sell their children into slavery? Nozick, it is widely thought, needs to do more to establish that property rights of the kind his argument presupposes are valid.

What about ownership of the self? Surely people at least own their own bodies – including their natural talents – in this 'full, absolute' sense? On this issue Nozick contrasts clearly with Rawls. Remember that, for Rawls, the original position models the idea that people as citizens are free and equal, and the idea that they are equal is partly captured by their ignorance of their natural abilities. This represents Rawls's view that the possession of talents is 'arbitrary from a moral point of view'. It is just luck whether one is born less or more strong, or clever, and so it would be unfair for people to be worse or better off than one another on that basis. At one point, Rawls says that his conception of justice treats people's natural talents as 'common assets'. It is easy to see why Nozick would object to this apparent failure to take seriously the separateness of persons, and the idea that people own themselves. Nozick doesn't deny that people's possession of natural talents (like the social class of the family into which they are born)

is a matter of luck. But that is neither here nor there. Even if it is luck, people nevertheless own themselves.

Most people accept some kind of self-ownership thesis. To test your intuitions, imagine how you would feel if the state argued as follows: 'It is just luck that some people are born with two good eyes, and others with none. To create a fairer distribution of eyes, we have decided to hold a lottery which will identify in random fashion some individuals who will be required to give up one of their good eyes to those who have none.' Most people, while accepting that the distribution of eyes is unfair, would nonetheless insist that their own eyes belong to them in a way that would make the state's proposal illegitimate. 'Look. These things are mine, they are part of me. If I want to give one of them to somebody who needs it more than me, then I can do so. Maybe I should. But the choice as to what I do must be mine, because the eyes are.' Those who endorse redistributive taxation while rejecting the coercive redistribution of body parts – probably the vast majority of the population – agree with Nozick about self-ownership, but deny that ownership of the self implies ownership, in the same full sense, of the things – goods, money – we create by using ourselves. People generally believe that forcible redistribution of body parts would involve a violation of their selves – would violate their integrity as people – in a way that forcible redistribution of things made by using those body parts does not. (Applying pressure to the pro-self-ownership intuition, imagine a natural disaster which leaves many injured and needing blood. Voluntary donations aren't enough. Is it obvious that the state would be wrong to set up a programme of compulsory blood donation?)

Rawls agrees with some aspects of self-ownership. Even though who has what body is 'morally arbitrary', we still have a right to bodily integrity, and an area of personal freedom within which we must be immune from intervention. In Rawls's view, for example, the individual must be free to do the job of her choice. The mere fact that I could be a brilliant surgeon, and would best serve my fellow citizens by becoming one, does not justify the rest of you in ganging together to force me in that direction. This, for Rawls, has more to do with the importance of the individual's

capacity to frame, revise and pursue her own conception of the good than with a right to self-ownership in Nozick's sense. Still, it is important to see that Rawls's claim about moral arbitrariness still leaves room to accommodate some of the widely shared intuitions that Nozick tries to capture in his notion of self-ownership. The big difference between them is that Nozick wants to use those intuitions in a way that extends ownership of the self to include ownership of the products made by the self.

### Popular opinion: justice as desert

It's important to see that Nozick does not claim that Wilt Chamberlain *deserves* the money he gets. To care about people getting what they deserve would be to go along with a patterned distributive principle of precisely the kind that Nozick doesn't like. The only reason Chamberlain has a justice claim to it – is entitled to it – is because his fans were entitled to their individual 25 cents and they freely chose to give that money to him. Whether he is deserving or undeserving is neither here nor there. If basketball fans for some bizarre reason decided to pay a bit extra to see some completely hopeless player, that player would still be entitled to whatever extra they paid.

Apart from wanting to get Nozick right, getting this clear matters because it helps us see how those who defend market outcomes on justice grounds tend very commonly, and completely illegitimately, to run together what are in fact quite different arguments. One argument holds that the market is essential to individual freedom or to respecting people's self-ownership. Forced redistribution of resources away from the outcome resulting from individual exchange violates people's freedom to do what they like with what is theirs. (I'll say more about this argument in part 2, on liberty.) Another, quite distinct, argument claims that the market gives people what they deserve. Talented, hardworking people deserve more than untalented, feckless ones, and the market makes sure that they get it. These justifications may coincide, in particular cases, but defenders of

the market shouldn't slide from one to the other without being aware that they may not.

So Nozick is not offering a defence of market outcomes that appeals to the idea of justice as desert. Rawls, too, from a completely different direction, is hostile to the idea that those whose productive activities can command a high price in the market deserve the money others are willing to pay them. In Rawls's case, this is essentially because luck plays too great a role in determining how much people can sell their productive activity for. The distribution of natural ability is 'arbitrary from a moral point of view', so those blessed with lots of the abilities that others are willing to pay for cannot claim to deserve greater rewards than those who are not. Rawls is thus hostile to what might be called 'conventional desert claims', claims such as 'Tiger Woods deserves to earn more than Jean Mason because Woods is a hugely talented golfer who gives great pleasure to millions around the world and is thereby able to sell his labour for a very high price whereas Mason is a social worker'.

Such claims are indeed 'conventional' in the sense that most people endorse them. We know that popular opinion is on Woods's side. It may not think that Woods deserves as much as he gets, but on the whole it is sympathetic to the idea that those who can do (and do do) things others are willing to pay for deserve to be better off than those who don't (even if the only reason why they don't is because they can't). We thus have the interesting situation that the two most influential political theorists on social justice – Rawls and Nozick – disagree with each other about whether it's just that Woods gets what he does. (Rawls says it isn't, Nozick says it is – indeed Nozick thinks that he shouldn't even pay any redistributive tax on it.) But they agree with each other that achieving social justice is not about making sure that people get the value of their productive activity on the grounds that they deserve it. (Rawls because of the 'moral arbitrariness' objection, Nozick because distributing according to desert is a patterned principle.) And, in agreeing this, they both disagree with popular opinion, which is largely sympathetic to conventional desert claims of this kind. Political philosophers are, on this issue, significantly out of step with the woman in the street.



To clarify our thinking about desert, let's distinguish three positions which I'll call the 'conventional' view, the 'mixed' view, and the 'extreme' view. The conventional view holds that one person can deserve to earn less or more than another even if this is due to factors that are beyond their control. Suppose that Jean Mason works as hard being a social worker as Tiger Woods does being a golfer. She worked just as hard at school and college, acquiring the skills she uses as a social worker, as Woods did acquiring his current skills. Her job now is at least as demanding – in terms of the effort it requires of her (emotionally demanding, long hours, short holidays) – as his is. The difference between their earnings cannot be attributed to any difference in their efforts, either past or current. Most people think that, in this case, Woods deserves to earn more than Mason. Not because he currently works harder, or worked harder to get where he is, but simply because his having been blessed with exceptional golfing ability enables him to do something that is more valuable – at least as measured by other people's willingness to pay – than what she is able to do. It's not her fault that she can't do what Woods does, and Woods can take no credit for the fact that he can and she can't. He's just lucky. Even in this case, the 'conventional' view holds that he deserves to be better off than her.

Contrast this with the 'extreme' view. This says that people do not deserve to earn less or more than one another even if they are exerting – or have in the past exerted – different amounts of effort. Somebody who works hard does not deserve to earn more than somebody who does not. What could possibly justify such a view? Answer: how hard somebody works is itself something beyond their control. People's character and psychological make-up are a function of their genetic constitution and their childhood socialization. Some people are born with a will to succeed, or to try hard. Others have that attitude instilled in them by their parents or other formative influences from an early age. Some are not so lucky. Why should those who have the good luck to be the kind of person who works hard deserve to earn more than those who have the bad luck not to be?

The 'conventional' view accepts the idea that people might deserve less or more than one another for deploying skills and

abilities which they are simply lucky to have or unlucky not to have. The 'extreme' view thinks that luck undermines differential desert claims and, because it thinks that effort is itself a function of luck, denies even that those who work hard deserve to earn more than those who do not. The 'mixed' view is the half-way house position. People don't deserve to be rewarded differently for things (or 'circumstances') that are genuinely beyond their control, like being born clever or stupid, or into a wealthy or poor family. But they do deserve to be rewarded differently for things that are genuinely a matter of choice – which include things like how hard you work, or what job, from those available to you, you choose to do. Rawls is right to think that it's unfair for people to be better or worse off than one another simply as a result of how they do in the natural and social lottery, but wrong if he thinks that people's choices should also make no difference to how well off they are.

Rawls is sometimes presented as holding the extreme view. He is not altogether clear on this point, but a plausible reading of what he says would have him acknowledging a role for free will, not claiming that every supposed choice an individual makes is actually determined by genetics and socialization. He believes rather that the choices people make about their level of effort are so influenced by factors beyond their control that it would be unfair to reward them simply in proportion to that effort. 'The idea of rewarding desert is impracticable', as he puts it, because it is impossible, in practice, to disentangle choices in the appropriate sense (i.e. choices uninfluenced by morally arbitrary characteristics) from the arbitrary characteristics that tend to influence them.

This seems plausible. Even if one believes that people do make choices for which they are responsible, and can deserve less or more than others on the basis of those choices, it is going to be very difficult to separate out anybody's current earnings into (a) that due to factors for which they can be held responsible and which they thus deserve and (b) that due to factors for which they cannot be held responsible and thus do not deserve. An important consideration here is that the abilities that adults possess reflect, to a great extent, how hard they tried when they were children.

Some adult abilities reflect natural talent. (Anybody who has seen the TV pictures of the 3-year-old Tiger Woods hitting a golf ball knows that he was blessed with prodigious natural talent.) But what isn't natural talent mainly results from people's habits as children. Some kids try hard, don't give up after the first attempt, develop the capacity to make what Rawls call a 'conscientious effort'. Some don't. But it is surely implausible to think that children are *responsible* for choices such as these. Their characters as children depend – when not on their genes – on their parents, their teachers, and other influences over which they have little or no control. It may be that, as adults, we are capable of making responsible choices about what to do with our abilities – and can be said to deserve greater or lesser rewards depending on the choices we make. But the very abilities we have as adults – where they result from choices at all – result largely from choices we have made as children, and for which we cannot be held responsible.

The most important thing to keep in mind, however, is that the market makes virtually no attempt to disentangle these various components of people's marketable skills. I say 'virtually' because two identically skilled people will tend to earn less or more than one another depending on how hard they work. But the marginal return to that marginal effort is trivial compared to the return to the skills they possess, and the market couldn't care less how they came to have those identical skills. Perhaps one was born lucky – high levels of natural ability, wealthy parents hence good education – while the other is less naturally gifted, and has had to struggle to better herself despite an unhelpful school. The market doesn't care. It is blind to distinctions of the kind I have been outlining here. It rewards people as a function of their ability to satisfy the preferences of others (actually – to satisfy the preferences of those others who have the money to pay to have their preferences satisfied). It pays no attention to the process by which people come to have that ability. And most of us have colleagues who are just as good at their jobs as we are even though they don't work as hard as us.

Even someone, like Rawls, sceptical about conventional desert claims might think that there are *some* things that you can indeed

deserve on the basis of attributes that you are just lucky to have. Suppose one thought that Seamus Heaney deserved the Nobel Prize for literature. That judgement need have nothing to do with any view one might hold about how he became able to write that poetry – whether through effort or natural ability or propitious upbringing. Even if there were minimal effort involved – he just happened to have been born with a gift for writing poetry and an unusually propitious upbringing – one could still say that he deserved the Nobel Prize. But that is because the Nobel Prize is awarded to the person who wrote the best literature. Since Heaney did that, he deserves the prize. So even the sceptic about conventional desert claims is likely to acknowledge that there are some contexts in which they are valid. The disagreement between the sceptic and the person who defends the market as giving people what they deserve turns, it seems, not on whether *any* conventional desert claims are valid, but on their proper *scope*. The sceptic says: ‘Why should some people have more resources to devote to their life plans than others just because they are luckier than those others? Sure. If somebody wants to offer a prize for the best poet, then the best poet deserves to win it – however lucky he is to be the best poet. But the money people get from their jobs is not like a prize. It is too important to be left to chance.’ The thoroughgoing sceptic might even say that Heaney deserved to be called the Nobel Laureate but did not deserve the money. Why should he have all that extra money to spend on his life just because he happens to be a great poet? On this view, conventional desert claims extend to symbolic rewards, like prizes, but not to rewards like money.

Like many concepts in this area, the term ‘desert’ is sometimes used rather loosely. In line with my commitment to drawing nitpicking (but clarifying) distinctions let me end by explaining how the idea of desert that I’ve been talking about here differs from other ideas which are sometimes formulated using the word ‘desert’.

First, there is a difference between desert and ‘legitimate expectation’. Imagine an institutional structure, a firm or the market economy as a whole, in which, as a matter of fact, people are rewarded unequally depending on their possession of certain

qualifications. We might then say that somebody who acquired those qualifications ‘deserves’ the reward just because the institutions were set up in such a way that the person acquiring the qualification has a legitimate expectation that, by acquiring the qualification, they would receive the reward. This is sometimes called an ‘institutional’ conception of desert. The important thing to see is that it is a completely separate question whether the institutions should have been set up the way they are in the first place. We can perfectly well say: ‘Since we are operating within a system that typically rewards people with good money if they get an MBA, and she has made various choices that have resulted in her getting an MBA on the basis of that assumption, her expectation that she should get good money is legitimate. In that limited sense, she “deserves” to get good money. Nonetheless, a system which rewards people with MBAs more than those without – indeed any system which pays people differently depending on their ability to pass exams of any kind – is fundamentally unjust, and certainly doesn’t give people what they really deserve.’ It is easy to formulate claims about legitimate expectations in ‘desert’ terms. Indeed, there’s nothing wrong with doing so – as long as one is clear that somebody can have a legitimate expectation of (hence ‘deserve’ in an institutional sense) a reward which they do not really deserve (because institutions are set up unjustly and do not reward people in accordance with their ‘actual’ or ‘brute’ or ‘pre-institutional’ deserts).

Second, some people use the term ‘desert’ when they are talking about compensation or equalization. Suppose I think people whose work is dangerous, stressful, dirty, boring, or inappropriately stigmatized should, other things equal, earn more than people whose work is safe, comfortable, interesting, healthy or prestigious. I might well say that they *deserve* to earn more. There’s nothing wrong with this kind of desert claim as long as it is clear how it differs from the kind I was discussing above. That kind was specifically to do with the issue of whether people might deserve less or more than others on the basis of their various attributes, and to what extent responsibility for those attributes was relevant. What we are talking about now uses a desert claim essentially as an equalizing claim. We can think of it in terms of

the idea of ‘compensating differentials’. In order to ensure overall or net equality between different people, we take into account the different characteristics of their work – interestingness, prestige, danger, etc. – and try to compensate for anything that would otherwise take them above or below some norm.

Again, there’s no real problem using the term ‘desert’ in this kind of case. It’s important, though, to see that is unlikely to justify the claim that Tiger Woods deserves to earn more than Jean Mason. It is completely implausible to think that the inequalities generated by the market in our society can be justified by appeal to the idea of desert as compensating differentials. (Some economists and political theorists think that the inequalities generated by an idealized perfect market could be. In that case the money people earned – the price for the job – would reflect nothing other than the net balance of advantages and disadvantages involved in doing their job. Employers would then have to pay more to get people to do unpleasant work than pleasant work – whereas the reverse is often the case at the moment.)

The third and last thought to be distinguished here can be, but need not be, related to this idea of compensating differentials. This is the idea that it is justified for some people to earn more than others because there will be bad consequences if they were not to do so. Sometimes this is formulated in terms of the idea of desert. Suppose we ask: ‘Do brain surgeons deserve to earn more than nurses?’ Somebody might reply: ‘Yes, they do. Because if we didn’t pay brain surgeons more than nurses nobody would want to be a brain surgeon. Since it’s clearly important that some people are brain surgeons, they deserve to get more money just so that we can make sure that some people choose to be them.’ This is a claim about incentives – about the need to induce people to do socially useful tasks and the justifiability of paying them more if that is the only or best way to get them to do those tasks. Does it have anything to do with desert?

Not as it stands. It is not, in itself, anything to do with the relative deserts of brain surgeons and nurses. It is simply a consequentialist observation, an observation about consequences, about what would happen if we didn’t pay them more. As it stands, we don’t know why, to get brain surgeons, we need to

pay them more than nurses. Perhaps it's because potential brain surgeons are more selfish than nurses and, realizing the value of their work to society, are prepared to hold the rest of us hostage, blackmailing us into paying them the extra. If that were the case, we would hardly want to say that they deserved that extra. (Any more than we would say that kidnappers who will only release a hostage if we pay them a ransom 'deserve' the money – even if we think we are justified in paying it to them.)

It can, however, be turned into a desert claim – at least a desert claim of the 'compensating differentials' kind. If we ask why we need to pay brain surgeons more than nurses if people are going to choose a career in brain surgery, the answer *might* be that they have high levels of responsibility and stress, or that they need to undergo many years of training – forgoing money they could be earning in other jobs and going through the arduous process of learning skills that most people don't need to worry about. So if we pay them above the average wage, to get them to do the job, this is just a compensating differential – money they 'deserve' given all the negative aspects of the job. The thought, now, is not simply that we have to give them extra money in order to get them to do the job – which is consistent with the blackmail scenario. It is that they actually deserve the extra, deserve it in the sense that it compensates for all the stress, long hours, training, or whatever and so provides the necessary inducement for them to take up brain surgery. Otherwise they'd be worse off, all things considered, than nurses. Construed this way, this is a genuine justice claim, and one that can be allowed appeal to the concept of 'desert' – even if it is a different conception of desert from the main one I've discussed. (Of course, such a claim could well be contentious. The kind of university education that some would present as investment, to be compensated for by higher pay, might well be enjoyable and valuable in itself. Just because someone tells us that something is a cost deserving compensation doesn't mean that we have to agree with them.)

## Conclusion

Each of the conceptions of social justice that we have looked at can be thought of as presenting a different justification of inequality. Hayek thinks that the whole idea of seeking social justice involves a philosophical mistake, so that inequality doesn't really need justification in the first place. Rawls holds that inequalities are justified if they conform to the principles that would have been chosen in the original position, most controversially the difference principle which holds that inequalities must serve, over time, maximally to promote the well-being of the least advantaged members of society. Nozick rejects this kind of thinking in favour of a principle of self-ownership that leaves people free to do what they like with property that is theirs – a principle that could justify extreme inequality. All three of these thinkers reject the popular view that people deserve differently depending on their productive contribution.

It is very common to find people defending the justice of the kinds of inequality we see in our society by appealing to some mish-mash of these different ideas. That is the reason carefully to distinguish between them. How could it be just that Tiger Woods, or Bill Gates, or any corporate lawyer, should earn more than a social worker, or a schoolteacher, or somebody who is involuntarily unemployed? Does the question involve a category mistake? Is it because their earning more – and that much more – serves, over time, to help the poor? Is it because they own their talents and whatever people are willing to give them for exercising them? Is it because they are in some way more deserving? These justifications *can*, in special circumstances, coincide – but they won't always do so. Those who would defend the justice of existing inequalities – or anything like them – need to think hard about which way they want to jump when they come apart.

### Further reading

Alan Ryan (ed.), *Justice* (Oxford University Press 1993) is a helpful collection, including key snippets from Hayek, Rawls and Nozick.



Tom Campbell's *Justice* (2nd edn. Macmillan 2000) is the best overview textbook.

On Hayek, the key work is *The Mirage of Social Justice* (Routledge & Kegan Paul), first published in 1976, and incorporated as volume II of his *Law, Legislation and Liberty* in 1982. 'The Atavism of Social Justice' in his *New Essays in Philosophy, Politics and Economics* (Routledge & Kegan Paul 1978) is short and to the point. John Gray's *Hayek on Liberty* (2nd edn. Blackwell 1986) and Chandran Kukathas's *Hayek and Modern Liberalism* (Oxford University Press 1989) are the two best critical accounts of Hayek's work as a whole.

Rawls's *Justice as Fairness: A Restatement* (Harvard University Press 2001) is the user-friendly version of his theory. From the mountain of secondary literature, it would be worth trying the introduction to Stephen Mulhall and Adam Swift's *Liberals and Communitarians* (2nd edn. Blackwell 1996) and Chandran Kukathas and Philip Pettit's *Rawls: A Theory of Justice and its Critics* (Polity 1990).

Nozick's *Anarchy, State and Utopia* (Blackwell 1974) is – as political philosophy goes – an entertaining read; the middle section on distributive justice is the most relevant. The best critical commentary is Jonathan Wolff's *Property, Justice and the Minimal State* (Polity 1991).

*What do we Deserve?* (Oxford University Press 1999), edited by Louis Pojman and Owen Mcleod, is a useful collection of papers on desert. Chapters 7–9 of David Miller's *Principles of Social Justice* (Harvard University Press 2000) defend the view that the market can (though it currently doesn't) give people what they deserve. Chapter 8 of Gordon Marshall et al.'s *Against the Odds? Social Class and Social Justice in Industrial Societies* (Oxford University Press 1997) is more sceptical.

— |

| —

— |

| —