

Philosophy, Business & Society

A basic introduction to ethics in professional life

Hallvard Lillehammer
Birkbeck College London
h.lillehammer@bbk.ac.uk

Table of Contents

Lecture 1: Why ethics is everywhere

Lecture 2: Who cares?

Lecture 3: Ethics for adversaries

Lecture 4: Jobs for your friends?

Lecture 5: Sell!, sell!, sell!

Lecture 6: Blah, blah, blah..

Lecture 7: Do you want to be trusted?

Lecture 8: The integrity question

Lecture 9: Who's to blame?

Lecture 10: In search of the corporate citizen

Why Ethics is everywhere

These lectures are about some of the ethical issues that arise in organizations, whether public or private; domestic or international. The aim of the lectures is to help the student to get a better handle on some of the basic ethical issues that arise in professional life. For whether we realize it or not, business managers, civil servants, advertisers, sales representatives and employers are all practical philosophers. We may not think explicitly in terms of philosophical arguments and theories, but every important decision we make is based on ethical assumptions that can be articulated and assessed. During the course of these lectures I give a critical introduction to some of the central philosophical issues that arise in the course of professional life, including truth; manipulation; trust; freedom; integrity; responsibility; and discrimination. In doing so, I will introduce the student to some of the most common ways that these ideas have been understood in moral philosophy and suggest how ordinary people can get a critical perspective on these ideas and theories without getting lost in endless debate and useless abstractions. By treating the discussion in the following chapters as a toolkit for exploring ethical issues as they arise in the course of their own professional experience, whether in business or elsewhere, I hope the students following this course will be able to go on to use the ideas and theories I will shortly introduce to think about these issues for themselves.

Before going on to discuss these particular ideas and theories we should be clear about the questions we are asking. In particular, we should be clear about what is meant by the claim that some professional issue is 'ethical'; how questions of ethics relate to practical questions of how to conduct business and professional activities effectively and successfully; and what, if anything, thinking systematically about ethical issues can contribute to professional life in practice. We should also respond to the obvious objection that the idea of 'ethical business' or 'the ethical professional' rests on a conflation between theory and practice that makes the whole project pointless and irrelevant. In this first lecture I address these questions by explaining how ethical questions are not only ever present in professional life, but are basically inescapable – whether we consciously think about it or not. We may not know what our professional ethics is, but once we start thinking about it, we will discover that we have one.

Ethical issues arise at various points in our thinking about professional life. We make use of ethical ideas when we try to understand what we are doing, either from a critical distance, or from the perspective of being engaged in the everyday business of doing our job. For example, we may consider how much information we are expected to give to a prospective customer. Would it be more sensible to provide more, or less? What if we are accused of being too vague, or misleading? What should we make of that complaint? These are questions that can arise either prior to action, or afterwards. Thus, we may wonder how precise to be in describing a product we are about to advertise. Alternatively, we may wonder in retrospect whether the customer who bought it really knew what they went in for. In asking and answering these questions we are likely to both be describing what we do (as misleading, for example) and

evaluating it (as less than what can reasonably be expected of us, for example). In the course of answering these questions we will effectively be committing ourselves to some kind of ethical classification of what we are up to, either in terms of concepts of better and worse (or good and bad), or right and wrong (or obligatory, permissible or impermissible). We will shortly have the opportunity to discuss how such ethical classifications are frequently disguised, camouflaged, or otherwise hidden from the surface of discussion by the use of words and arguments that make it seem that they are completely absent from professional life. But just because these ethical ideas are well hidden from the surface of discussion does not imply that they are not there at all.

There are at least three reasons in particular why ethical commitments are sometimes hidden beneath the surface of discussion. The first is that often our primary concern is not ethical, but rather self interested, instrumental or pragmatic. Thus, we may be more focused on whether our negotiation strategy is going to be successful than on whether it will meet with anyone's approval. Having said that, every negotiation strategy embodies *some* values, and people are usually quick to appeal to those values when challenged, e.g. after having withheld information on grounds of privacy, efficiency, or self-protection. The second reason is that our interest in ethical commitments is frequently descriptive rather than normative. Thus, we may want to understand whether the average negotiator thinks that lying is OK if it will help to get a good price, or whether a product for sale has been produced by people who work in conditions we would describe as slavery. Yet even a descriptive interest in ethical commitment can embody our ethical values, at least insofar as our ethical values will have an effect on what descriptive questions we are interested in in the first place. Thus, it may be controversial among colleagues whether lying is always OK provided it will help to get a good price. And few people would want to defend commercial arrangements that are agreed to amount to slavery. Furthermore, sometimes what appears to be a purely descriptive claim about ethics actually turns out to be an ethical claim, as when a long term customer complains that 'Nobody feels any loyalty to their high street bank anymore', or when a politician who defends the deregulation of financial institutions responds to a reformist policy proposal by saying: 'The markets would never permit that'. A third reason is that in the hustle and bustle of everyday decision-making we want answers and solutions to particular questions and problems, such as whether to lie to the electorate, or whether making bankers swear an allegiance to a banking code is likely to reduce the amount of irresponsible trading or will just add another layer of inefficient bureaucracy. We do not always either want or need abstract lessons in ethical theory, or arguments for ethical business practices derived from philosophical first principles. Even so, whenever we approach particular questions and problems we do so by thinking of reasons for and against the various alternatives. Some of these reasons will be ethical reasons. We do want those reasons to be informed by the relevant facts and to hang together in a coherent way so that we have something intelligent to say when our decisions are challenged. It is partly a philosophical task to work out how the different ethical considerations hang together and why. To that extent, we are implicitly engaged in a philosophical project virtually every time we act for a reason, whether in a private or a professional capacity.

Of course, ethical norms and values are not the only norms and values we appeal to in deciding how to behave in our professional lives. We also appeal to legal norms and values, such as the laws and regulations embodied in statutes across different legislative domains. For an international trader, for example, the question whether a

given trade is actually legal in a given jurisdiction is likely to take priority over the question whether it would be an ethically advisable trade in principle. After all, getting things the wrong way round between ethics and law may cost us our jobs, or even get us incarcerated. In some cases, we may have good reason to lobby, or otherwise argue, for a change of policy. Yet even if the ethical professional – just like any other citizen - is also a member of a polity and so free to engage in politics, he or she is normally best advised to observe the policies that are actually in effect, as opposed to any policies she would like to see implemented. Looking beyond the law, our professional lives depend on understanding a vast range of conventions and social expectations that will never be explicitly written down or endorsed by statute, and which may vary enormously from institution to institution, or from country to country. Thus, effective communication and negotiation depends on some kind of shared understanding of who speaks when, and for how long, during business meetings, for example. While parts of this understanding may be justified by participants on ethical terms (such as not loudly interrupting other people when they are speaking), other parts (such as how long meetings go on for without a break) might be accepted by everyone involved as just the way we do things ‘around here’, and so easily replaceable by other conventions in different circumstances (sometimes even in the same organization).

Just exactly where do legal or conventional reasons end and ethical reasons begin? That is, indeed, a very delicate question. Fortunately, we do not need to answer it here. One question we do need to answer, however, is why we should be confident that once we have subtracted all the purely conventional reasons in play, in addition to any legal considerations, there is anything ethical left to say. Why not think of professional ethics as a social code that tells us to follow whatever local conventions exist in a given institution, so long as these conventions comply with the law? To see why this - perhaps initially tempting - answer will not do, it will help to have a concrete example to hand. There is arguably no better example to make the point than talk about *rights* and *duties*, and the way that such talk both informs and transcends the rights and duties embodied in law and other conventions.

Let’s say that someone has a right when they have an interest or power that justifies imposing a duty on others to respect that interest or power in some way. Thus, I have an interest in making voluntary arrangements with others and the basic power to do so by giving and accepting promises. If that interest and power are important enough to justify the imposition of a duty not to interfere with my ability to make voluntary arrangements by giving and accepting promises, then I can be said to have a right to do so; at least as long as my doing so does not interfere with the same right as enjoyed by others. Now some rights are embedded either explicitly or implicitly within a legal system as institutionally enforceable protections, privileges or constraints. Two obvious examples of direct relevance to professional ethics are rights to property and rights arising from contracts. A devil’s advocate might argue that law exhausts the domain of rights as far as professional ethics is concerned. Consider, for example, the never-ending debate about the ethics of exploiting so-called ‘tax loopholes’. Yet no law is self-interpreting, however long and complex the legal documentation we may be able to produce. For example, even legal professionals will need to exercise their judgement in borderline cases. Furthermore, no law is perfectly determinate in every possible eventuality. For example, existing laws are likely to be silent about some applications of future technological innovations that no one has thought of yet. (This

is one of the most challenging aspects of the – still quite recent - invention of the Internet, and all the opportunities for monitoring and information capture that this technology has recently made possible.). And even where the law is determinate, it will not always be enforced; not only because of discretionary powers held by judges or right-holders, but also because the universal enforcement of some laws would be prohibitively expensive, or time consuming. (Tax law is an interesting case in point here.) Yet even where law is both determinate and enforced, laws and legal systems are changeable and do frequently change, e.g. when there is a change of government. And government legislation tends to be subject to criticism, from representatives of businesses and other institutions as well as the electorate, at least some of which may appeal to extra-legal considerations (such as effects on the non-human environment.). Last, but not least, not everything that is agreed to be undesirable, bad, or wrong is either productive or even possible to legislate against. Thus, a business may be badly affected by the poor morale or high turnover of its employees, but the best solution is unlikely to be to legislate against unhappiness. Or a business may be badly affected by the lack of solidarity among its staff, but incentivizing a better company ethos is unlikely to be best handled by making friendship compulsory. Similar considerations apply to such informal obstacles as openness, attentiveness or failures to ‘go the extra mile’. In each case, both identifying and addressing the problem is likely to require the exercise of evaluative judgements that do not reduce to matters of existing law or accepted convention. Not all of these considerations are interestingly ethical, of course. But some of them definitely are. To this extent, ethics is everywhere, both in our private and in our professional lives.

If ethics is everywhere in professional life, how should we handle things when ethical issues arise? To get clearer about this question is the main task of the lectures that follow. To get started, it may help to briefly introduce two different models of how ethical considerations might be embodied in an institutional context, each of which is likely to give a somewhat different answer to the question of how ethical challenges are most helpfully addressed when they arise.

According to the first model, ethical considerations are best considered by the individuals working within an institution asking themselves ethical questions, going through ethical arguments and reaching their own ethical conclusions. This is the model that J. R. Boatright calls ‘the moral manager model’, as applied to the case of large scale business corporations, such as banks, insurers and industrial corporations¹. According to the moral manager model, institutions ought to incorporate ethical considerations into managerial decision making and train their employees accordingly (e.g. in the ethical issues of tax evasion versus tax avoidance). According to the second model, ethical considerations are best considered by the individuals who design and regulate the institutions in which people work, leaving owners and employees to pursue the strategic objectives of their institution (such as profit maximization) without having to take into account the ethical ‘fall out’ of what they do. This is the model that Boatright calls ‘the moral market model’, as applied to large scale business corporations; a model according to which it is the responsibility of ‘society’ (e.g. legislators, regulators and those responsible for institutional design) to

¹ J. R. Boatright, ‘Does Business Ethics Rest on a Mistake?’ *Business Ethics Quarterly*, Vol. 9, No. 4, 1999, pp. 583-59.

develop markets that produce ‘ethical outcomes’, even without ‘ethical managers’. Although Boatright is primarily concerned with a very specific professional role (i.e. the business manager) in the context of a very specific institution (i.e. the modern business corporation), his two models easily translate into two different ways of distributing ethical responsibility across institutions quite generally. Put in slightly different words, the question is whether ethical considerations are better applied at the level of individual decision makers within an institutional setting, or alternatively at the level of those who design and regulate these institutions. With respect to our focus on ethical judgement, the question is whether we should primarily be concerned with the ethical qualities of the institutional structure, or with the ethical qualities of the individuals working within that structure. This is a question that arises for all of us, whether happen to be occupying a professional role within a given institution or whether we are judging the behaviour of the institution and its members from the outside. Moreover, it is a question that is unlikely to have a simple answer, or to have the same answer in every case. Thus, there could be institutions (possibly all real-world business institutions are like this) where some element of each model should be present in order to protect against downside ethical risk (as when making highly speculative investment decisions, for example). Furthermore, there could be institutions where each of the two different models is more suitable for different levels in the institutional hierarchy (e.g. between the level of executive decision making on the one hand versus the shop floor on the other). Either way, it would arguably be wrong to think of the two models as mutually *exclusive*. Finally, there may be some professions for which one model is more suitable than it is in others (e.g. investment banking versus medical care, for example). If so, it would arguably be wrong to think of the scope of these models as mutually *exhaustive*. In any case, it is probably safest to think of these two models as ‘ideal types’ that we can usefully appeal to in order to sharpen the focus of our thinking on a case-by-case basis. We shall begin to do so immediately in the next lecture, in which we discuss the frequently observed, but surprisingly poorly understood problem of professional, ‘detachment’, or indifference.

Who cares?

We are all familiar with complaints about the emotionally detached health professional, the callously selfish business person, and the faceless indifference of modern bureaucracy. What exactly is the problem? Is acting ethically never consistent with being, detached, indifferent or uncaring towards others? If so, then how and when? To what extent should professionals try to cultivate the same attitudes in all contexts, or 'roles' (e.g. in their 'private' versus their 'public' life)? If some measure of professional detachment is sometimes justified, what are its limits? How can institutions be designed, or professionals trained, to avoid the ethical downsides of professional detachment? These are some of the questions raised in this lecture.

What exactly do we want from this discussion? Why should we theorize about concepts such as 'indifference' or 'detachment', given the many different ways they are understood in different walks of life? There are at least four reasons for doing so. First, we would like to avoid misunderstanding or talking past each other when these terms are used, especially when we use them to criticize each other. We want to know who is 'really' indifferent, detached, or otherwise uncaring, and how. Second, we have an interest in better understanding professional practices of that encourage attitudes of indifference and detachment, for example when we shift our attention from one profession to another. Third, we need to engage in an informed and intelligent way with instances and opportunities for professional indifference and detachment, whether we are just thinking about it, practicing it, or trying to avoid it. Finally, we have an interest in being able to judge ourselves and others, either because we are actually faced with a decision that requires judgement in practice, or because we are considering a mere possibility, such as to introduce a new business practice (such as online trading) that removes the face-to-face interactions between people on the basis of which so many of our ethical judgements are normally made.

The kinds of indifference and detachment we are interested in can usefully be distinguished from the following three kinds of indifference with which they are sometimes confused. First, we sometimes talk about indifference when we say that something doesn't matter. Indifference in this sense is lacking in importance, or not *making* a difference. This is usually a matter of what the thing we are indifferent to is like, and so not a matter of our attitudes to it. Thus, I could care enormously about something that, on reflection, I agree really does not matter. Likewise, I could become indifferent, or cold, towards something that really does matter, such as my health, or what is happening to my friends and family. It may be tempting to think that the ethics of indifference and detachment is a matter of caring about the things that matter and not caring about those that don't, but in fact that is way too quick. In fact, many of the most difficult ethical problems in professional life arise precisely because we are encouraged not to care about things that clearly do matter in order to care more effectively about something else. We will return to this important thought again and again throughout these lectures.

One way of being indifferent or detached from something is not to form beliefs about it, or to make up one's mind about it at all. While not itself an attitude towards anything, this kind of indifference or detachment is actually ethically quite important because what people believe and think about is often a sign of what they do, or do not care about. Thus, one cause of culpable negligence is a state of cultivated ignorance. I am sure that most readers will be familiar with the temptation to avoid knowledge of uncomfortable truths, whether at work or at home. Another way of being indifferent or detached is to remain neutral between things, either in the sense of not having a preference, or not wanting to choose. This is also an ethically important sense of indifference and detachment, partly because retaining an attitude of neutrality is sometimes a symptom of a refusal to take responsibility and so live up to the demands of the situation. No one likes to be called a coward. Yet if we push our attitude of neutrality too far, that is exactly what we are likely to be called. A final kind of indifference or detachment is the kind of attitude associated with someone who is simply not concerned about something; either by not caring about it; not paying attention to it, or not being aware of it at all. In most cases when people complain about indifference and detachment, whether in professional life or elsewhere, it is this kind of indifference or detachment that they have in mind. To get clearer about what this complaint amounts to, it will help to make one more distinction.

So far I have been talking of 'indifference' and 'detachment' as if they basically mean the same thing, and in ordinary conversation the two terms are often used interchangeably. Even so, it can be useful to keep them apart, even at the cost of some artificiality. This is because uses of 'indifference' and 'detachment' are ambiguous between a 'binary' (as in 'detached' versus 'semi-detached') and a non-binary (as in 'I found the surgeon rather detached') interpretation. In what follows, I will use 'indifference' to capture the binary use and 'detachment' to capture the non-binary use that admits of gradations (for example when someone grows 'more detached' over time.). Although by no means perfect, making this distinction allows us to keep apart the question of whether what is wrong with someone is that they simply fail to care about something, or that they don't care about it enough. For the sake of simplicity, we will focus primarily of the first issue, and so on indifference as just defined.

Thus understood, we can say that a *subject* (e.g. an employee) is indifferent with respect to some *object* (e.g. a client), by not displaying some *orientation* (e.g. friendliness) towards that *object* in a certain *context* (e.g. when answering a customer call). Our question is when, if ever, an indifferent orientation to some object is ethically appropriate on the part of some subject, and in what contexts this might be so.

To answer that question we need to take account of the fact that states of indifference can vary along two separate dimensions. First of all, they can be more or less *dynamic*, for example in the sense that they serve our purposes or ends. Thus, I may be indifferent to the carbon footprint of my company in order to gain advantage over my environmentally more friendly competitor, or just because I don't care about the environment. Second, states of indifference can be more or less *object sensitive*, for example in the sense that they result from a negative judgement about their objects. Thus, I may be indifferent to the female employees in my company because I hold a (prejudiced) view of their abilities, or just because I am unaware of my own bias or

prejudice. Most of the ethically interesting cases of indifference in professional life combine the feature of being highly dynamic with being acutely object-sensitive, as when the intense pursuit of commercial targets blocks out from consideration the visibly high levels of stress imposed on employees. Such cases do not reduce to the question of whether the people involved either care enough in the first place, or whether they know what they are doing. On the contrary, such cases normally involve a complicated choice between different institutional priorities and a sophisticated knowledge of what is going on. It is this fact, not a basic failure to ‘give a shit’ that makes thinking about the ethics of indifference and detachment in professional life so difficult.

There is a place for targeted, dynamic and object sensitive indifference in professional life. In the context of a social world that is relatively stable and institutions that are basically decent to those affected, by excluding some aspects of a situation from our concern, we are sometimes able to better address the aspect of that situation which it is our professional role to manage or control. For example, a competent bureaucrat may practice a kind of dynamic and object sensitive indifference when he or she brackets all consideration of religious or ethnic differences in distributing a universally available public good, such as health services free on delivery to all citizens or residents of a certain community. A professional recruiter may practice the same kind of virtuous indifference when anonymizing job applications in order to prevent selection on the basis of racial or gender bias. Likewise, a business manager may practice the same kind of virtuous indifference by refusing to consider personal connections in distributing opportunities and responsibilities in order to resist the common temptation to let the soft corruption of ‘friendly favours’ take over the otherwise healthy ethos of a listed company. None of this implies that no-one in the relevant organization should ever think about the objects of virtuous indifference in question. On the contrary, a system designed to prevent the development of a culture of ‘friendly favours’, for example, may have people devoted specifically to consider these issues and to impose governance structures to keep them in check. To take just one among many examples, most modern institutions have guidelines about who should see what information about their employees at any given time, such as what different employees are paid. The fact that some of this information is not available to everyone working in the organization does not mean that there is no one in the organization whose job it is to care about that very thing. Selective indifference is a ubiquitous feature in professional contexts that depend on a ‘division of labour’, including (as in some cases involving delicate personal information) *ethical* division of labour.

A second kind of virtuous indifference results from the simple fact that sometimes we do better as a collective by each of us ‘minding our own business’. Indeed, in some situations, each of us cultivating a state of object sensitive indifference may be a basic necessity for economic survival or just keeping the peace. This may occur when there are aspects of each other’s beliefs, values or ways of life that people on the other side vehemently disagree with, find disgusting or otherwise disapprove of. The religious or sexual orientation of the home delivery courier is not something it makes sense to worry about when signing off on a dispatch of groceries or furniture. The same point arguably applies to a wide range of everyday multicultural transactions, whether they take place in the private or professional ‘sphere’. Yet the adoption of an attitude of selective indifference is not merely a defensive strategy. We frequently bracket our

differences also in the service of mutually beneficial cooperation, as when two competing companies buy and sell goods and services from each other. Once again, selective indifference is a ubiquitous feature in institutional contexts that depend on a 'division of labour', including the so-called 'invisible hand' scenarios that economists appeal to, where mutually disinterested competition has the indirect consequence of being mutually beneficial overall. We shall have other opportunities to consider the case for such appeals in later lectures in this course.

If we fail to monitor downside risk, or carry our lack of concern beyond its proper institutional context, a state of virtuous indifference can easily transform itself into a professional vice. This can happen if we focus all our attention on the ends of our action, forgetting about the means. For example, by chasing targets or giving into pressures of efficiency it is easy to forget about the side effects of our success, whether they involve failures of compliance, 'dirty hands', or harm to third parties or the environment. It can also happen if we abdicate our own responsibilities as individuals by giving in to the authority or demands of our superiors. A third possibility is that the tasks in an organization are so fragmented that by focusing on our own part we forget about the significance of either the other tasks, or the whole. For example, a set of management accounts consists of rows of numbers added together and subtracted to produce a balance sheet. The numbers themselves do not speak, and therefore say nothing about the damage caused or the lies that were told in order to produce them. We can think of the vice involved in ignoring this possibility as one of blinkered indifference.

While blinkered indifference is highly dynamic, it need not be very object sensitive. Other forms of vicious indifference are. Thus, it is all too easy in a competitive context to focus all one's attention on the 'in-group', with exclusionary or harmful effects on members of the 'out group', whether they are business competitors, members of other ethnic or religious groups, or just people in general (as when someone says: 'We've won! Forget about everyone else!'). In a variety of cases, we might argue that favouring 'our own', whoever they may be, is part and parcel of fair competition and the division of labour. Yet this will not be true in all cases. This kind of exclusionary indifference can also occur when people just forget about 'the rules of the game', as in some cases where companies are found to have been in breach of legally enforceable rules of financial transparency or fair competition.

The negative effects of vicious indifference are frequently observed in organizational behaviour. One obvious symptom is the all-too widespread toleration of ethical or legal abuse. Another symptom is the temptation of employees to conceive of business as an 'amoral', or 'ethics free', zone; or to 'stonewall' ethical questions. A third symptom is the personal stress that is sometimes felt when denied the opportunity to think openly about ethical questions that are mutually known to arise. A fourth symptom is to misdescribe oneself, or to deceive oneself, into thinking that ethical considerations have no role to play in our decisions even when they actually do. Thus, it may be tempting in a cut-throat environment to present oneself as merely 'wanting to seem ethical' for purely egoistic reasons even if there are obvious limits to what we would do to others even if we could get away with it (i.e. where our motivation is

really at bottom ethical after all.) These are only some of the symptoms collected by Bird & Waters under the provocative heading ‘the moral muteness of managers’.²

The obvious causes of ‘moral muteness’ include competitive social arrangements in which the success of individuals and institutions depend on their ‘positional’ advantage over others, where for you to do well is for you to do better than others. They also include the complex hierarchies of authority and sub-tasks that characterize most large scale institutions, as well as the bias and prejudice found in virtually every minimally complex institutional environment. Of course, there are other causes of moral muteness as well. In each case, we are arguably dealing with what is an inescapable aspect of professional life in institutions of any minimal size. To think that we could control or eradicate moral muteness with a simple mechanism, or in one fell swoop, would therefore be hopelessly naïve. As we saw in the previous chapter, the most effective way to approach these ethical risks is arguably to identify the problems and address them case by case by a combination of ethical deliberation and organizational regulation, or design. In the next lecture, we shall consider one especially vivid example of this problem, namely the question that arise when by doing well in your won task is either likely, or certain, to cause loss or harm to others.

² F. B. Bird and J. A. Waters, ‘The Moral Muteness of Managers’, *California Management Review*, vol. 32 (no. 1), pp. 73-88. 1989.

Ethics for adversaries

Once upon a time, the corporate giant *Google* had as its motto ‘Don’t be evil’. The obvious problem with officially adopting such a motto is that any corporation acting in a highly competitive environment will frequently do well only at the cost of others doing badly, or at least not as well as they could. The same issue arises for company employees competing for their jobs when threatened by redundancy; lawyers litigating on behalf of clients who claim a right to take over all of someone else’s assets; police and prison personnel who take away people’s freedom, and millions of ordinary people struggling to maintain themselves in conditions of scarce minimal resources. The problem is that in a wide range of institutional contexts it is not practically possible to be successful without causing significant harm to others. How can that not be evil? If we think, as most of us do, that doing well in such environments is perfectly compatible with acting virtuously, or well, what we need is what Applbaum has aptly labelled an ‘ethics for adversaries’.³

There are three things we would like an ethics for adversaries to help us with. First, we want to understand what it is to occupy a professional role, and how to set the limits of its authority. Thus, with few exceptions we generally don’t want our friends and family to take their professional role home with them. Second, we want to understand how the roles in question work in their institutional context so as to be able to engage with them, occupy them, and criticize them as necessary. Thus, we are most of us familiar with the situation where our occupation of a professional role presents us with a dilemma (e.g. when we are dealing with a family member in a professional capacity), or where institutional boundaries are challenged (as when someone argues that in this particular case we should make an exception). Third, we want to understand how institutions and policies work in society more widely, so that we can form a view about how we want them to be and express those views either as citizens, policy makers or other participants in a democratic process.

Following Applbaum, we can distinguish three different ways of thinking about professional roles. On one way of thinking about roles, their nature and function is defined by their *essence*, or what they really are. It is tempting to think of some professional roles that way, such as the role of a doctor, the essence of which has surely got something to do with preserving or promoting their patient’s health in the face of death or disease. Yet once we move beyond this obvious example, this way of thinking about roles quickly collapses. Most professional roles are historically contingent, temporally variable and socially constructed to the point that even if they had a universal essence it would arguably be so vague and indeterminate as to leave most interesting practical questions unanswered. Consider, for example, the role of being a ‘head of state’. If its essence does not inform us whether a head of state

³ A.I. Applbaum, *Ethics for Adversaries: The Morality of Roles in Public and Professional Life*, Princeton 1999.

should be a monarch or democratically elected; have executive power or be purely ‘ceremonial’, it hard to see why we should be particularly interested in it. Either way, it will frequently fail to help us with most of our ethical questions.

Another way of thinking about roles is to define them in terms of how they are actually occupied and understood at some given time and place. One obvious advantage of this ‘positivist’ view is that what some role actually consists in is a question it is in principle possible to settle by empirical investigation. This, it is definitely the case that some countries are monarchies and others not. The problem about this way of thinking about roles is that all professional roles are subjects to interpretation and criticism, which normally involves at least some minimal normative element. (Compare the discussion of legal rights in Chapter 1 above.) Thus, we may all share a view of doctors as being essentially involved in the activity of ‘healing their patients’. Yet where does ‘healing’ end and ‘enhancement’ begin? Is it up to my doctor to worry about my happiness even if my failure to be happy is not a disease? These are some of the very questions we want to have answers to, but which merely describing what doctors currently do is not going to answer for us.

The obvious fallback is to think of roles in terms of what we consider to be their *proper* function, either everywhere or just in some given time and place. Once we start thinking about roles in this way, however, we are already doing ‘ethics’ – at least to some minimal degree. While some theorists would rather avoid this normative element in our understanding of professional roles, there is arguably no better way of doing so; at least as long as we bear in mind that all normative claims about proper function can be, and frequently are, hotly contested. This, however, is just as it should be, given that most professional roles are historically contingent, temporally variable and socially constructed. Having said that, thinking of professional roles in terms of their proper function does raise a number of important challenges given the fact that we cannot assume that every occupant of a professional role will conceive of that role in terms of what we consider its proper function. Thus, in the case of adversarial practices there may be people who occupy ethically contested roles who conceive of themselves and the harms they cause in ways that do not correspond to the way we think of that role. Consider, for example, the case of corporate lawyers who make a fortune representing known criminals. It is hardly an advance in our ethical understanding to claim that people in this position have misunderstood the nature of their role. Nor is the claim that they are not ‘really’ occupying that role (as it ought to be understood) very likely to make them think again. Very quickly, we will just be arguing about words.

Furthermore, professional role-holders face ethical challenges in the here-and-now, and given how things *actually* are, not merely in the abstract and given how they ought to be. Thus, a ‘straight’ police officer working in a corrupt and dangerous police force may endanger both her career and her life by acting in every situation *as if* her colleagues were occupying the role of police officer as defined in terms of its proper (as in ‘not corrupted’) function. Any sensible police officer in this situation will need a way to translate her ideal conception of police work into the currency of a non-ideal world.

Finally, understanding professional roles in terms of their proper function may not always be possible. Thus, there may be professional roles that for ethical and other

reasons, we think that on reflection ought not to exist. Consider, for example, the role of an executioner, a torturer, a whale-hunter, or a *torero*. If what the execution of the role essentially involves is something that we think had better not exist, there is little to be achieved in thinking about what a 'good' instance of its kind would be. For practical purposes, therefore, we shall sometimes need to fall back on the kind of empirical descriptions appealed to by the 'positivist' view.

Is causing harm consistent with the effective occupation of a professional role in light of its proper function? As we noted at the outset, a vast range of individual and institutional behaviour seems to presuppose that it is. If so, what's the explanation, given that we normally think of causing harm as bad, or wrong?

The most common argument for the claim that certain harms are ethically permissible in adversarial contexts appeals to the idea of *consent*. Following Applbaum, the basic idea is that insofar as the harms in question are a result of the normal operation of rules and regulations to which people have voluntarily agreed, they have no complaint if it turns out that they are the ones who ends up on the losing side of the bargain.

Once again, the appeal to consent can be understood in three different ways. When understood in the first way, permissible harms are constrained by the *actual* consent of participants in adversarial practices, where these participants are assumed to know what is in their own best interest and have agreed to participate on that basis (e.g. by entering the marketplace in order to make a better living). While no doubt true about some of those who loose out in one or another version of 'the rat race', it is not true about everyone. First, and foremost, not all participants in adversarial practices have consented to those practices. You can be sued and loose all your possessions without even having signed up to the judicial system in your territory, or even ever having heard about it in advance. Second, no participant in adversarial practices has consented to every aspect of those practices, not only because of their sheer complexity but also because of how aspects of how the system is designed is normally delegated to institutions over which they have no control. Third, not all participants can always be assumed to know all the 'rules of the game'. Not only is much predatory advocacy and salesmanship based on this fact, it also explains why we all depend on various kinds of 'experts' in our formal dealings with each other. Fourth, partly because of their relative ignorance of complex systems, participants cannot always be assumed to know what is best for them. Fifth, for some participants in adversarial practices the results of those practices are definitely not best for them. Lack of knowledge, expertise, or simply lack of time will often stand in the way. Sixth, adversarial practices frequently have deleterious effects on third parties whose consent has not, or could not in principle be given. Thus, if my lawsuit brings you financial ruin, I may also in the process ruin the life of your family, including your children. Finally, even actual consent arguably only legitimates the result of an adversarial practice if its participants have acceptable alternatives in signing up for it. This is not always so, as can be vividly seen in documentary portrayals of 'sweat shop' industries or child labour. The harms caused by existing adversarial practices cannot therefore be justified everywhere on the basis of actual consent.

It is sometimes argued that participants in adversarial practices *tacitly* consent to those practices by participating, complying, and accepting the benefits of those practices, although they never explicitly agree to them. No doubt this is sometimes

true. If I go along with your suggestion of a business venture for long enough without protesting, you may arguably form the reasonable expectation that I will deliver when push comes to shove. Yet not everyone who goes along with a practice will actually approve of it, once more because they might be ignorant of relevant facts or consider that they have no acceptable alternatives. Furthermore, where consent is merely tacit rather than explicit it is even more vague which aspects of the practice I consent to and which I do not. Does paying your bus fare imply your consent to the current standard of public transport? Perhaps it does, perhaps it doesn't. Either way, the prospects for justifying the harms caused by adversarial practices are not much brighter when we move from actual to tacit consent.

In response, it is sometimes argued that the consent that matters is not a matter of what participants in adversarial practices actually agree to, but what they *would have agreed to* if they were asked in conditions where they could rationally decide on the basis of being able to consider all the relevant facts. Thus, it is undoubtedly true that we sometimes accept the excuse of someone who interferes with us that we would have approved of their action if we were thinking clearly or knew what was going on. (Imagine a colleague who cancels a trade on the grounds that she knows that you have decided to make it on the basis of having received deliberately misleading information from a customer.) The most general problem with appeals to *hypothetical consent* of this and similar kinds is that they often seem to presuppose what they are meant to establish. What are the 'ideal' conditions for giving valid consent if not those conditions, *whichever they are*, that issue in consent to permissible actions? Even if there is a satisfactory answer to this question (which for all we have seen there might well be), there are two further problems with appeals to hypothetical consent; at least the second of which is undoubtedly serious. The first problem is that in a vast range of cases, actual consent (or its absence) is general agreed to trump hypothetical consent, both ethically and legally. Thus, it is normally not OK for me to take possession of your car on the grounds that you would have lent it to me if I asked for it. Unless I actually ask and you consent, by taking it I am committing a theft. Second, for a wide range of people who will be predictably harmed by adversarial practices, accepting those practices is not in any way rational. For at least those people, the justification of these harms - assuming they can be justified - are unlikely to derive from their hypothetical consent.

Because the appeal to consent is too limited to capture all the cases we want to cover, it is natural to supplement it by appealing to substantial benefits that adversarial practices bring to its participants; whether they stand to win or lose from its operation in any given transaction. Building on this idea, it has been popular in political philosophy to appeal to the idea of 'fair play'. The basic thought behind this idea is that where people engage in a mutually advantageous practice by restricting their freedom to yield benefits for all, those submitting to these restrictions have a right to similar submission from others who benefit from the practice whether those people have actually consented to it or not. When applied to the present case, the idea is that since some adversarial practices (such as a system of formal litigation) yield benefits for everyone by means of practice participants restricting their freedom in accordance with the 'rules of the game', those participants are permitted to impose those rules on everyone who stands to benefit from the existence of the practice, even when imposing those rules will have negative or harmful effects on people in some of its individual applications. To think otherwise is to license a culture of 'free riding' on

universally beneficial practices, such as the legal system, health care, public taxation or gun control.

Sometimes the argument from fair play is run together with the case from hypothetical consent to yield the dubious argument that we would all agree to basic rules of fair play if we knew the relevant facts and were fully rational. Yet even when considered in isolation from the case from hypothetical consent, it is clear that the potential for the argument from fair play to justify the harms caused by many adversarial practices is strictly limited. First, some people have no choice but to benefit from universally beneficial adversarial practices, even if they would prefer not to. Thus, I might prefer not to benefit from the environmental restrictions on private vehicular transport imposed by new legislation. This offer of protection is basically an offer I can't refuse. Second, some people, such as members of small minorities, may have no choice but to participate in some adversarial practices, their withdrawal from which would result in the destruction of their communities. Consider, for example, the need to access institutions of legal justice, however unfair or prejudiced that system might be considered to be overall. Third, not all participants in a practice may actually be advantaged by its effective operations. Thus, there are forms of competition that severely restrict the possibilities of minor investors to enter the market. Fourth, in some cases mutual benefits are achievable without the adversary aspect. Thus, not all legal systems favour a strictly adversarial model. Fifth, in some cases, the mutual benefits in question could be achievable without the adversary aspect being applied in a particular case. Thus, it is possible for a single player to have a dominant market share without driving their competitors to bankruptcy. Sixth, even if everyone technically benefits from a given practice, those benefits (and the correlated burdens) are frequently unevenly distributed to the point where the dice are systematically loaded in favour of some over others. Consider, for example, the costs in litigation and legal fees that can be afforded by some, but very few, individuals and corporations in developed economies.

The biggest problem with appeals to fair play, however, is that they tend to ignore the fact that for a significant number of people, even in modern societies, there is no clear case for the claim that they themselves actually do noticeably benefit from the practices in question. For people on the margins of society, it is not always obvious, for example, that it is better to abandon local practices of vigilante justice in favour of what other people call 'the rule of law'. Indeed, the fact that 'the system' is widely perceived as being 'against us' is arguably one of the main causes of so-called 'populist' complaints against the state, 'the elites', or other manifestations of the public order. Unless we dig deeper than the argument from 'fair play' encourages us to do in its standard variations, we shall never be able to understand the underlying force of these complaints.

The conclusion we should draw from this is not that the arguments from consent or 'fair play' inevitably fail to justify harms caused by adversarial practices in business, law, or politics. Subject to standard assumptions, they arguably do justify such harms in a non-trivial range of cases. The conclusion we should draw is that the scope of these arguments is limited, and that sometimes these arguments fail to address at least some of the harms that existing adversarial practices actually cause, especially where these harms are less than obviously foreseeable or well-understood. In these cases, those harms could be without justification. When they are, it is the responsibility of

the relevant institutional representatives (whether business managers, lawyers or politicians) to reconsider those practices in the light of their harmful effects.

Jobs for your friends?

Bias, discrimination, favouritism, nepotism, social exclusion... The list is virtually endless. The problems are widely recognized and universally subject to disapproval and censure. Yet they seem to come up even where you least expect them and to raise their ugly head even when we think we have got them under control. One of the most puzzling things about discrimination and nepotism in professional life is that they often seem to survive even our most strenuous efforts to avoid them.

What is discrimination? What, if anything, is wrong about it? Those are the two main questions we discuss in this lecture. We would like to get a better idea of what are the distinctive problems of discrimination in different institutional contexts, and to make some progress on the question how to respond to discrimination and nepotism in professional life. In particular, we would like to get a better handle on kinds of response to discrimination and nepotism that are ethically advisable in different circumstances (e.g. in different kinds of institutions) so as to be better prepared to handle public concerns about discrimination and official responses to it, whether in terms of legislation, institutional design, or the ever contested distinction between what is often called the ‘public’ versus the ‘private’ domain.

To discriminate between people is to treat them differently on the basis of some *feature* with respect to which they are different; either themselves or on relation to others (including yourself). For example, the state will treat people differently depending on whether they are citizens or not. To ‘discriminate’ in this sense is ethically neutral. The default assumption is that ‘being more discriminating’ (e.g. more perceptive) is a standard sign of intelligence and social awareness. In most cases where discrimination is controversial, the differences in question relate to the membership of people in some social group. What makes discrimination controversial in these cases is the question of whether membership of the relevant social group is a fair basis for treating people differently.

To discriminate between people *wrongfully* is to make a difference between them on the basis of some feature they have that does not justify that difference in treatment. For example, it is wrongful (and illegal) to discriminate against a job applicant the basis of their class, gender or ethnicity. Speaking more precisely, to wrongfully discriminate between people is to *wrongfully* impose a relative disadvantage on some of them in virtue of their membership of some social group, where the wrongfulness of the discrimination rests partly on the fact that they are treated differently *because* they are members of that group. For example, it is wrongful to discriminate against applicants for a senior role on the basis of the (ethical) judgement that people of a certain ‘class’ are universally unsuited for positions of major responsibility.

The law distinguishes between discrimination that is ‘direct’ and discrimination that is ‘indirect’. To discriminate *directly* against someone is to impose a relative disadvantage on them when the imposition of that disadvantage is *aimed* at, or *intended*, by those imposing it. For example, it is directly discriminatory to exclude someone from career opportunities on the basis of misguided paternalism, as when a

male manager omits to expose their female colleagues to certain challenges because those challenges are thought to be ‘too tough for girls to handle’. To discriminate *indirectly* against against is to impose a relative disadvantage on them, where that imposition of disadvantage is *not aimed at*, but is a foreseeable or systematic or consequence or side effect of the action that is strictly intended, or aimed it. For example, it is indirectly discriminatory to impose a dress code on employees that members of a minority group in society is unable to comply with for religious reasons.

Both direct and indirect discrimination is unlawful in most developed countries. But even if many important arguments about discriminations are legal and many of the most important ways of dealing with discrimination are a matter of actually enforcing the law, other important arguments and measures are distinctively ethical. This is not only because there is always the question of *what the law should be*. It is also because much discriminatory behavior takes place in situations where questions of law and its enforcement do not apply. For example, when colleagues go out to socialize after work they will often do so with the same people most of the time. In many cases, this pattern of behaviour will have indirect effects on who, among the employees of a certain company, are happy at work and likely to prosper in the company as a result. Yet there are obvious limits on what kinds of ‘after work’ behaviour it makes sense to legislate for, or otherwise formally regulate. So the question of wrongful differential treatment has an ethical dimension that extends far beyond the limits of the law.

When treating people differently is wrong, why is it wrong? Unsurprisingly, philosophers have offered a large number of apparently competing answers to that question.⁴ When pressed a little further, however, it turns out that few of these answers are mutually exclusive. It is therefore possible that, in at least some cases of wrongful discrimination, there is more than one thing that is wrong with it. And that, for those of us who are not dogmatic, is just fine. Just as the same kind of event can have more than one kind of cause, so the same kind of wrong can have more than one explanation. So let us now consider a sample of what those explanations are.

One complaint about some cases of discrimination is that it treats people differently on the basis of features that are not within their control. Wrongful discrimination of people with inherited disabilities is an obvious case in point. In many cases, this complaint is plausible. In other cases, however, it fails. After all, some characteristics beyond our control are justified grounds for discrimination. For example, at the current state of technology there is no justified cause for employing blind airline pilots. We should also bear in mind that just because something is within our control that does not make it a justifiable ground for discriminating. To insist that someone should bleach their skin in order to fit in with the rest of the crew, for example, would be a grossly offensive way to treat them as a colleague.

A closely related complaint against some forms of discrimination is that it fails to treat people according to desert, or merit. This is one of the guiding ideas behind the claim that careers should be ‘open to talent’. This is also a plausible idea, at least in theory. In practice, however, appeals to desert or merit are normally faced with at

⁴ See e.g. Altman, A. ‘Discrimination’, *The Stanford Encyclopedia of Philosophy*, 2011. Available at: <http://plato.stanford.edu/entries/discrimination/>

least two complicated factors. Probably the most important of these is that what counts as ‘desert’ or ‘merit’ is always defined relative to criteria or standards that are themselves contestable and potentially subject to criticism. For example, when faced with the fact that some people are unable to complete a professional task due to a physical accessibility issue (e.g. if they are confined to a wheelchair), there is sometimes the possibility of changing the physical environment so as to make it accessible, as opposed to deciding that the people in question are unable to perform the relevant task. There is also the further complication that the question of ‘desert’, or ‘merit’ is not always the overriding ethical consideration in deciding how to distribute benefits and burdens within a given institution. In extreme emergencies, for example, we may just need to get some job efficiently done. In other cases, not focusing on the question of ‘desert’ or ‘merit’ is arguably a defensible way of addressing the problem of wrongful discrimination itself. Perhaps the most obvious case to illustrate the point is the (ever controversial) practice of ‘affirmative action’, according to which members of historically disadvantaged groups are given preferential treatment to reduce the incidence of structural inequalities. It would seriously stretch the bounds of plausibility to argue that the members of the historically advantaged groups who loose out as a result of affirmative action loose out because they, individually, don’t ‘deserve’, or ‘merit’ the opportunities in question.

An equally common diagnosis of wrongful discrimination is that it is a product of *prejudice*, or *stereotyping* or other forms of unacceptable bias. Wrongful discrimination is a product of prejudice when the members of some social groups are treated as though they are not entitled to the same respect and concern as are members of other social groups, based on either ignorance, false beliefs or arbitrary assumptions about what distinguishes members of one social group from the members of another. Stereotyping occurs when individuals are discriminated against on the basis of features they either do not (or need not) have; because members of their social group are somehow associated with those features. Prejudice and stereotyping is a common explanation of wrongful discrimination against ‘foreigners’, immigrants, or members of other ethnic or religious groups. They are also frequently visible in institutional contexts where members of one social group are assumed to be incapable of some tasks, regardless of their individual capacities. A very good case in point is traditional assumptions about male versus female physical strength, and the barriers these assumptions have created for women seeking a career in some manual professions, or the military.

Differential treatment caused by prejudice and stereotyping is obviously wrong. The problem in practice is that both prejudice and stereotyping are associated with practical ‘heuristics’, or ‘rules of thumb’, which in certain contexts can be valuable in promoting the *efficiency*, *predictability* and *protection against risk*, and which sometimes *actually do* correctly pick out individuals as unsuitable business partners or sources of danger. Consider, for example, travelling alone to a foreign country and being approached by total strangers offering you easy travel from the airport to the city centre. There are obvious ‘heuristics’, or rules of thumb, that may help you to avoid ending up either broke or in serious danger in such situations, even if considered as a universal generalization these ‘heuristics’ may leave something to be desired from an ethical point of view. Furthermore, it is generally agreed that the ethical constraints on who to choose as our colleagues or business partners do not

extend to the question of who to choose as our friends, or romantic partners. Very few, if any, human societies have ever denied some form of that distinction. In most societies we know at all well, individuals and families have been granted some form of 'refuge' from the calculating eye of impartial justice. For these reasons, and because we are finite creatures with limited knowledge and limited time to consider the causes and effects of our various beliefs and motivations, these causes of wrongful discrimination are close to inescapable in practice. Which is not to say that we should not try to understand, control, and minimize their effect as much as we possibly can.

Wrongful discrimination is differential treatment that is *unjustified*. The lack of justification in question can be a matter of drawing arbitrary distinctions between people; wrongly evaluating people's characteristics as negative when they are not; treating people on the assumption that they have features that they do not in fact possess; or otherwise restricting opportunities or prospects on the basis of false generalizations or unfair expectations. Sometimes when people wrongfully discriminate they will do so with ill will, bad intentions, or while thinking of their victims as inferior in some way. Yet this is not always so. It is perfectly possible to wrongfully discriminate (e.g. 'indirectly') while having 'the best will in the world'. This could well be the most important lesson to draw from our discussion of the ethics of discrimination in professional life. As for the small sample of views we have discussed, they each seem to get something right, although it is probably unwise to think we can make do with either just one of them, or with only the ones we have considered here.

Sell!, sell!, sell!

In recent years we have been treated to horrifying stories of credit card customers being mis-sold redundant payment protection insurance, individual investors enticed to purchase investment packages designed for financial institutions only, investment banks encouraging their customers to buy shares that those banks are themselves getting rid of at the same time, and much, much worse. In all these cases, what the customers got was not exactly ‘what it said on the tin’. On the contrary, they were subject to adverts and sales tactics that made them part with their money on false pretenses. On the other hand, the entire basis of the market system they were aiming to benefit from is based on the legitimacy of voluntary exchange, and a wide range of techniques of advertising and persuasion. And after all, few, if any, of these costumers were ‘forced to buy’. In fact, at the moment of purchase, many of these customers were actually getting exactly the product what they wanted. So what exactly is the problem?

It is obvious that the problem has something to do with the value of truth and transparency in commercial communication. What this actually amounts to, however, is a tricky issue. In this lecture and the next, we explore this question; focusing first on the ethics of sales and advertising before we turn to the general issue of truth and truthfulness in the next lecture.

What attitude should an ethical professional have to the desires and needs of their colleagues, customers or clients? To what extent can institutions and their representatives legitimately aim to ‘create’, or ‘mould’, the desires of colleagues, customers or clients? Why do we think it matters exactly how they do so? There seems to be a difference between *influencing* a colleague, customer or client on the one hand, and *manipulating* them on the other, but what exactly is it? These are some of the questions we explore in this lecture. What we hope to gain by this exploration is partly a understanding of what influencing colleagues, customers and clients actually consists in. Once armed with understanding, we may also be able to better influence colleagues, customers and clients by communicative means that we actually approve of, and to praise or criticize individuals and institutions when they either do, or do not, live up to our legitimate expectations.

To sharpen our minds, it will help to focus on what we think goes wrong when someone’s decision is based on *poor*, *undue*, or *the wrong kind of influence*. What counts as poor, undue, or the wrong kind of influence is normally sensitive to the nature of the transaction, who the stakeholders are, and context in which the decision is made. For example, it normally matters whether the transaction in question involves a transfer of economic resources (as when I make an investment on the stock exchange), whether there are any asymmetries of power among the people involved (as when I invest on the advice of a professional stock broker), and what the stakes are for those involved (as when what I am investing is my entire pension savings portfolio).

Advertisers and sales representatives have an interest in presenting their products in such a way as to ‘tick the right boxes’ of their customers in terms of how those customers understand the decision they are making.⁵ For example, sometimes we have a high personal investment in buying something that we also feel very strongly about emotionally. Consider, for example, buying a wedding gift or a romantic holiday. A successful sales pitch in this case is likely to involve a fair amount of detail and an attractive presentation.

At other times, we have a high personal investment in buying something that we don’t feel very strongly about emotionally. Consider, for example, the pension investment mentioned in the previous paragraph. In this case, a successful sales pitch may be likely to focus more on the details than on the attractiveness of the presentation.

In other cases, we may have little personal investment in buying something that we do feel strongly about emotionally. Consider, for example, deciding between different kinds of sweet or entertainment, the primary purpose of which is to make us feel good. In this case, the detailed facts of the case are likely to matter less in a successful sales pitch than that the product is attractively presented.

Finally, we sometimes have low personal investment in buying something we don’t feel very strongly about. For example, there may be little more at stake for me either way when I decide between basic office supplies or different brands of paint. A successful sales pitch in this case may therefore simply focus on the question of price.

Having said that, most products on the market are sold in a competitive environment. And all the different products of the same kind are rarely sold all at the same price. So companies have an interest in making their product more attractive by presenting them in such a way as to make the customer *more* invested in them, and to attach a *higher degree* of emotional significance to their choice. It is by targeting their customers by designing, packaging, advertising or selling their product in a way that often draws attention away from the basic features of those products that successful advertisers are able to make fortunes both for their clients and themselves. Yet what could possibly be wrong in convincing someone to buy something by making it look more, rather than less, attractive so long as everyone involved is clear about the fact that this is what we are doing?

Following Roger Crisp, we can distinguish between three different ways of presenting the attractions of a product, depending on whether we are presenting it in light of features it actually has (‘informative advertising’); features it does not actually have but the customer is likely to desire (‘persuasive advertising’); or presenting it in ways that are not readily transparent to consciousness at all (‘subliminal advertising’).⁶ Of the advertising in the first category, consider an advert that informs you of the engine capacity of a motorized vehicle; the length of warranty of a camera; or the ‘period features’ of a residential property. Of advertising in the second category, consider an

⁵ See e.g.

http://www.adcracker.com/involvement/Consumer_Involvement_Theory.htm

⁶ R. Crisp, ‘Persuasive Advertising, Autonomy and the Creation of Desire’, *Journal of Business Ethics*, vol. 6 (no. 5), 1987 pp. 413-418.

advert that associates a brand of deodorant with being thought sexually attractive; a brand of tobacco with life in the 'great outdoors'; or a brand of detergent with familial and domestic bliss. Examples of subliminal advertising are, by definition, harder to come by. Maybe this does not matter. Most of us will be familiar with the after-effects of aspects of our experience (such as rapidly changing images, background colours, or barely audible noises) that we fail to pick up on at the time, or the subtle effects of the background computations that track, respond to, and transform your online presence across various devices. The thing about subliminal influences, if and when they occur, is that we are supposed to not be aware of them, except indirectly.

There is little ethical controversy about informative advertising as such. After all, the informative advertiser presents the product in light of features it is said to possess, and which it can at least in principle be verified as possessing by the person buying it. It is practically impossible to present a product for sale without revealing at least some of the aspects it is supposed to possess. So unless that presentation includes either deliberate or negligent falsehood, informative advertising is ethically above board, even though we may sometimes feel there is much too much of it.

Having said that, some products are widely considered unsuitable even for informative advertising. This is normally because they are considered *harmful*, *addictive*, or not suitable for *sale on the open market*. Pornography, tobacco and vital medicines are sometimes placed in one or more of these categories. Furthermore, some products are considered unsuitable for advertising to selected audiences, such as children or people unqualified to manage them. (Perhaps with respect to our smart phones, there is some danger that we are all turning into children in this respect.) Some weapons are obvious candidates to place in this category. In addition, some information is widely considered unsuitable for use in advertising even if the product advertised is not in itself regarded as unsuitable to advertise. Thus, there are quite strict norms in place for the way in which defense equipment is presented even to the most knowledgeable customers. Most importantly, however, even informative advertising is selective in the aspects of the product it informs about, and how. Every advert is an exercise in selective emphasis. One way to confuse the customer is to bombard them with excess information. And every advertising image represents a choice of subject, angle and context, for example. So if we are worried about forms of advertising having undue influence on customers, we have every reason to study even informative advertising quite closely.

Persuasive advertising is a different matter. Once again following Crisp, we can identify three potential sources of undue influence in making a customer decide to buy a certain product. The first is where the decision is induced without the customer's knowledge, e.g. where the main reason they prefer it is one they are not fully aware of. Thus, you may be convinced to buy a more expensive house than you need because the agent selling it has been able to tap into your status anxiety. The second cause of undue influence is where the decision is induced for a reason that will not stand up to reflection. Once again, the case of status anxiety could be a perfect case in point. A third cause of undue influence is where the effect of the advertising is precisely to prevent the customer from doing the rational reflection that would show whether or not the product is worth buying, or at that particular price. Thus, by tapping into your status anxiety the agent may be causing you to be unduly anxious, or to forget what your real priorities are (such as the person who is constantly

checking their online status for updates in search of the ever so elusive ‘Like’). In this, and in many other ways, a persuasive advertiser can engage your thoughts and feelings in such a way as to make you decide to buy, or stay with, a sub-optimal product for reasons that you yourself would not recognize as good if you knew what was really going on. If these forms of undue influence are wrong, and if all forms of persuasive advertising depend on at least one of them, then all forms of persuasive advertising are wrong.

On second thought, it is not obvious that we have to be this pessimistic, nor even that we *can*. First, not all cases of persuasive advertising actually depend on causes that are hidden from our reflective consciousness. Indeed, some of it seems to depend on their being easily accessible, e.g. when the ‘attractive’ features they are associated with are presented ironically, or as a joke. Thus, few people actually believe that bottom-of-the-range deodorants make people sexually irresistible. Part of the point of presenting them that way is to draw the customer in by making him or her laugh (potentially at him or herself).

Second, some of the attractive associations included in persuasive advertising actually happen to be real, whether we like it or not. Thus, status anxiety may be undesirable in and of itself, but there are very good prudential reasons to seek the kind of status that makes other people treat us with respect. In other words, not all customer preferences produced by persuasive advertising are actually harmful.

Third, some preferences formed by exposure to persuasive advertising would survive reflective scrutiny because in the wider scheme of things they don’t really matter. This is arguably what happens in the case of some products, such as inexpensive washing liquids, that are often virtually indistinguishable in themselves and therefore often advertised by the use of attractive imagery that has nothing to do with their essential, or functional, features.

More interestingly, some aspects of persuasive advertising are inseparable from the product itself. Consider, for example, the elegant but non-functional aspects of design that distinguishes one brand of consumer good from another, whether it is a telephone, a laptop, a multimedia system, a car, or the visual presentation of a residential building. The associations of being seen to carry something elegant, for example, are normally very different from the associations of being seen to carry something ugly. Indeed, some persuasive element is arguably present in every single communicative transaction where one party attempts to put something in a favourable light, from a basic invitation (as in ‘I think you will like it’) to the most sophisticated sales pitch (as in a video or a luxury catalogue). The interesting question about persuasive advertising, therefore, is arguably not whether it is ‘right’ or ‘wrong’. Instead it is a question about when, and how, it amounts to undue influence. And the answer to that question, as we have just seen, depends largely on the facts of the individual case.

A subliminal advertiser would present their product in light of features it may or may not possess, but crucially in ways that are not accessible to the consciousness of the prospective customer. As a result, the customer will be making their decision based on a lack of knowledge, where that lack of knowledge is likely to affect the rationality of their decision-making. Whether the product they decide to buy is actually ‘worth it’ or not would therefore seem to be ethically beside the point, insofar as they have

clearly been *manipulated* in a way that fails to show respect for them as autonomous, or free, rational agents.

Subliminal advertising, assuming it can actually be made to work, is clearly ethically problematic. To induce someone to make a decision precisely *because* they are unaware of one of the main causes for their making it would seem to be one of the most obviously wrong ways to manipulate someone, on a par with making them buy something by telling them a lie. We shall return to the ethics of lying and other forms of manipulation in the next lecture. Before we do so, however, we need to register one important caveat that should always be kept in mind when thinking about the ethics of communication, whether in the context of advertising or more generally.

The issue is that, subliminal or not, all communication involves at least some causal elements beyond the conscious awareness of its recipients, and regardless of whether these elements are deliberately chosen by the persons communicating or not. Colour schemes may affect our mood. Frequency may affect our sense of urgency. Tone of voice may affect our levels of trust. Whether we like it or not, we don't normally communicate as fully autonomous angels protected by fully anonymizing screens. Not even the internet (or perhaps - given the algorithms on which it is based - especially not the internet) affords us this level of idealized luxury. These reflectively non-transparent aspects of communication exist everywhere. If there is anything like a soluble ethical problem here, what is it?

Setting aside the issues of addiction, children and other vulnerable customers, perhaps the best we can say is this: manipulative advertising is ethically problematic to the extent that it is *deliberately* designed to be effective on the *condition* that its suggestive aspects are not made transparent to reflective consciousness (e.g. when catering to someone's status anxiety). The real problem with manipulative advertising arises when the customer would not have made a purchase if they knew what was really going on. In other words, their purchasing decision was based on deception, or false pretenses. They were basically sold a lie.

Blah, blah, blah..

None of us like to be lied to. We hate to find that we have parted with our money based on deception or on false pretenses. Yet most of us actually lie quite a lot. And many of us frequently vote for politicians who we want to represent us, knowing full well that much of what those politicians say is insincere, misleading, obfuscatory, bullshit, spin or just plain false. To what extent is ethical behaviour compatible with such lying, obfuscation, bullshit, or lack of insincerity? What is the value of truth and truthfulness in the context of professional, as opposed to private, life? And when, as would sometimes seem to be the case, are the norms of truth and truthfulness that we all know and love permissibly defeated by other practical considerations? Those are some of the questions we discuss in this lecture.

In order to think more clearly about these questions we need to have a basic understanding of what truthfulness and deception consist in. In particular, we need to have a clear idea of the various different forms that lies and deception can take. If we are lucky, this will help us to better detect and classify different communicative acts as either basically truthful or deceptive in practice; to form a view about which forms, if any, of deception are ethically acceptable in different circumstances; and to better be able to engage with laws and regulations that concern the ethics of communication, such as the considerations that go into our intuitive understanding of what goes wrong when someone is ‘mis-sold’ insurance, or when a minister is accused of misleading parliament.

It will help to get started by considering some of the core ideas at work in claims and counterclaims about truthfulness and deception.

The first idea we need is that of *truth*. A statement, such as witness testimony, is true just in case things are the way the statement says they are. Philosophers often put this point in terms of the idea of truth as ‘disquotation’: a statement ‘P’ is true if and only if P. We can therefore say that disquotation is the ‘core’ of the truth predicate, and that is enough for us to get on with here.

The second idea we need is that of *truthfulness*. Roughly speaking, a person (such as an expert witness) truthfully makes some statement ‘P’ if and only if they *believe* that P, and P is actually *true*, or otherwise correct. Truthfulness therefore has two parts, namely the *sincerity* of the speaker, and the *truth*, or accuracy, of what they say.

Having got a handle on the ideas of truth and truthfulness, we are now in a position to start thinking about *deception*. Roughly speaking, one person deceives another person with respect to some statement, ‘P’, if and only if the first person *intentionally misleads* the second person with respect to P. There are many ways of misleading someone, of which here we should mention only some. Thus, I could mislead you about the fact that the car you are about to buy has defective breaks by *omitting to mention it*; by *suppressing the test report*; by *talking about the car in an obfuscating way*, or – more simply – by *lying* about it. If you call me out on my deception I may go further and respond by *bluffing* (as when I pretend that my deceptive behaviour

was really just a miscommunication, or a misunderstanding.) Or, if the levels of trust between us are really that low, by *double bluffing*, as when *by telling you the truth I make you discount it* because I know you will think that *I must be lying*. (Apparently, double bluffs are the stable diet of Mafiosi and secret agents.)

Not wanting to complicate things more than necessary, we shall focus here on *deception by lying*. There are at least two plausible ways of understanding what lie is, and since they are both obvious forms of willful deception, we shall make no attempt to choose between them here. According to the first way of understanding it, a lie occurs when some person makes a statement he or she believes to be false to another person with the intention of deceiving that person about the facts *in some way*.

One concern about this definition is that a person can make a statement he or she believes to be false to another person with the intention of deceiving that person about *any fact whatsoever*, as opposed to the facts about what he or she *believes*. Thus, a politician may intend to deceive their back-stabbing colleague about their *intentions*, not their *beliefs*, as when they say to their back-stabbing colleague: ‘I know you would never betray me’ before mounting a plot to have their back-stabbing colleague either sidelined or deselected. On some ways of understanding the word, this kind of willful deception is not a lie, *strictly speaking*.

According to a second way of understanding it, a lie *strictly speaking* occurs when one person makes a statement he or she believes to be false to another person with the intention to deceive that other person about *the truth of that statement*. Thus, a politician may intend to deceive their back-stabbing colleague about whether or not some third colleague as betrayed them both. There is little or no controversy about whether or not a speech act that meets these criteria is really a lie. So in what follows we shall allow ourselves to treat it as our official definition.

What we *communicate* by saying something is not just a matter of what we *literally say* (or even *assert*), but also of *how* we say it, and thereby what we ‘con conversationally *implicate*. Consider, for example, the various different interpretations (admiration, disgust, ridicule, etc.) that could be given to the question ‘Did you see that shirt?’ as an interview candidate leaves the room. The ethics of communication does not boil down to a simple norm of *truthfulness*.

Indeed, *speaking truly* is consistent with *lack of sincerity* (e.g. when a recruiter sinks one of the applicants by saying: ‘No doubt somebody could conceivably think that he meets all the essential criteria for the job’; manipulation (e.g. when a politician ‘spins’ a news story to draw attention away from an inconvenient truth); willful deception (for example in the double bluffs mentioned above); power abuse (e.g. when consistently stressing the negative aspect of an employee’s performance), or just a general lack of good will (as when writing in someone’s character reference: ‘As a student, she always demonstrated her talent for nice handwriting’.) Whatever may be wrong with a lack of truthfulness, we should not fetishize truth for its own sake.

Nevertheless, we all agree that truthfulness matters, and that there is something iffy about someone who makes statements about all sorts of things, apparently regardless of whether they are true or not. This is the person that Harry Frankfurt calls ‘the

bullshitter'.⁷ On Frankfurt's (partly stipulative) definition, one person *bullshits* another person with respect to some statement, P, if and only if the first person asserts P for reasons *unrelated to the truth of P*, for example because the first person does not care whether P is true or not, or because by getting their audience to believe that P, they will be more successful in getting what they want.

The 'bullshitter' can easily be identified among some of the manipulative advertisers we discussed in the previous lecture. They are also an all-too-frequent occurrence in politics and the entertainment industry. Having said that, we should be careful not to overstate the point. If neither the person speaking nor their audience really care about whether what is said is true or not, they could both be 'bullshitting' in the perfectly innocent sense of having some fun. We should be careful about our commitment to truthfulness lest it prevents us from ever 'having a laugh'. When bullshit is really a problem, this is normally because it is used as a tool in deception, manipulation or the abuse of power, all of which we should worry about; both in our private lives and as professionals.

So what is wrong with deception in the cases where we think it is wrong? As we have seen in previous chapters, there might be several plausible explanations for the wrongness of the same action, and we should not think that we have to choose just one of them in every case.

Perhaps the most well-known explanation of its kind is the notorious claim about lying, historically associated with Immanuel Kant, that since all morally acceptable principles apply universally to everyone, and since by lying we are making an exception of ourselves given that a lie can only be successful against the background of people being truthful, the principle 'go ahead and lie when it is convenient' cannot be universalized, and so must fail as a morally acceptable principle. If we wanted to generalize the point to deception in general, the idea would be that since successful deception depends on a background of truthful communication, all acts of deception involve the deceiver essentially making an exception of themselves, and this is morally problematic, or wrong.

Although it does clearly have a powerful intuitive appeal, it is hard to accept this explanation as a diagnosis of the ethics of deception in general. First, Kant himself distinguished between active lying (which he thought was morally impermissible) and passive failure to tell the truth (which he thought could be morally innocent) on the grounds that telling the truth could hardly be a universalizable duty since it would seemingly give us all an impossibly demanding duty to tell the truth about everything to everyone whenever we can. Yet it is perfectly obvious that failures to tell the truth (e.g. by avoiding a difficult issue) can be at least as harmful as telling lies, so even if Kant were right that there is an interesting distinction here, it is hard to believe that it fully captures what's ethically problematic about lies and deception in every context we want to criticize it. Furthermore, some practices involving lying or other forms of deceit can be embedded in institutional practices the basic norms and principles of which are themselves universalizable. Thus, a rule licensing limited public deception in matters of national security could be justified against the background of a political

⁷ H. Frankfurt, 'On Bullshit', in his *The Importance of What We Care About*, Cambridge: Cambridge University Press, pp. 117-33.

system of democratically accountable government. Alternatively, a rule licensing bluffs and counterbluffs (as in ‘This is our final offer’) could be justified against the background of a negotiating practice the basic terms of which has the *consent* of all participating agents. Thus, people enter deception-involving practices voluntarily both in order to promote their own interest (e.g. for fun, or their own security), and in order to preserve their capacity to voluntarily consent (e.g. where secret intelligence functions to protect the integrity of the state against external aggression). It follows that any plausible ethics of deception should be sensitive to whether, and if so how, acts of lying and deception are licensed and controlled within an institutional hierarchy of licensing and control. It should also be sensitive to the comparative harms and benefits that truthfulness, lying and deception are likely to produce.

A different explanation of the wrongness of deception appeals to how lies and deception often undermine the practices of which they are a part, and potentially end up causing more harm than good even for those who stand to benefit in the short term. Thus, it is frequently observed that lies and deception can undermine the trust and transparency on which stable voluntary exchanges depend, and can therefore be in conflict with the contractual market transactions on which markets and entire economies depend. An obvious example of this is the protective divestment of multinational companies from troubled regions of the world where there is a high risk of data being falsified or previous agreements not being honoured.

Once more, this explanation fails to capture the full complexities of the ethics of deception. There are at least two reasons for this. The first is that some practices are arguably be enhanced by the licensing of limited deception. The example of business bluffing just mentioned is one example. The way in which information is managed in the context of adversarial legal practices is another. (Compare our discussion of legitimate harms in Lecture 3.) And even if we are attracted to the idea of ‘full disclosure’ as some kind of ideal in all professional communication, even full disclosure is constrained by the fact that it is simply not possible to fully disclose everything, all at the same time. Whether we like it or not, some relevant facts are always likely to be ‘up our sleeve’, whether we like it or not. The second reason is that some ethically legitimate practices are *defined* by their licensing of limited acts of deception. Perhaps the most obvious case in point here is the game of Poker, which is doubly apt for our purposes, insofar as it is both a *game* with clearly definable rules and normative expectations (such as the prohibition of aggressive or intimidating behaviour) and also an international *multimillion-dollar industry*.

A third possible explanation appeals to the idea that communication is essentially a co-operative activity, and so even if deception is sometimes ethically innocent this is only insofar as it takes place against a background in which ‘cooperative principles of good communication’ are *generally* observed. According to the philosopher P. Grice, for example, these norms imply that you should ‘[m]ake your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged’.⁸ Given the

⁸ P. Grice, ‘Logic and Conversation’, in his *Studies in the Way of Words*, Cambridge: Harvard University Press, 1989.

vastly different contexts and practices in which communication occurs, and given the vastly different kinds of harms and benefits that lies and deception may produce, it is hardly a big surprise on this view that we should find it difficult to come up with a simple formula of why deception is wrong when it is, even less so why deception should always be thought to be wrong (which, as we have seen, it isn't).

In his widely influential paper, Grice articulates four generic rules of successful communication. The first is the rule of *quantity*, which says that you should make your contribution as informative as required, and no more. The second is the rule of *quality*, which says that you should try to make your contribution one that is true. The third is the rule of *relation*, which says that you should be relevant. The fourth is the rule of *manner*, which says that you should be perspicuous (i.e. avoid obfuscation, ambiguity, excessive elaboration, etc.). One way of reading these rules is as setting an ideal standard for any speaker who wishes to make herself understood to her audience, *all else equal*. Another way of reading the rules is as giving a framework for how to interpret a speaker *on the assumption* that they are trying to make themselves understood. A third way of reading these rules is as articulating a standard which, unless we *tend to approximate to* it in practice to a significant extent, will make all communication impossible. (This point is also relevant to the Kantian point about universalizability we discussed above.)

It obviously does not follow from our acceptance of these rules thus interpreted that no communication is possible unless the fourfold standard is always fully met. Nor does it follow that speakers always aim to make themselves understood. Much less does it follow that they always have an overriding reason to want to. Indeed, the kind of adversarial contexts and practices that we have discussed in this lecture (and in Lecture 4) arguably include cases where often they do not. And even taking Grice's rules at face value, we must remember that what information is 'required' is always sensitive to context (e.g. in the context of negotiation); that some legitimate practices either require, or are enhanced by, the withholding of information, even where this is likely to mislead (e.g. matters of state); that some ethically legitimate practices effectively suspend the 'quality' rule (e.g. matters of politeness or courtesy); that not all aspects of ethically legitimate practices are aimed at a simple end relative to which 'relevance' can easily be established (e.g. mingling, or 'networking'); that some ethically legitimate practices effectively suspend the 'manner' rule (e.g. protective obfuscation in the context of hostile or dangerous environments); and, finally, that what counts as 'good', or 'successful' communication overall depends on assumed *background ends* (e.g. we are engaging in a process of therapy or market competition) and *background conditions* (e.g. conditions of mutual trust vs. general distrust). The different ways in which each of these factors can apply will frequently make the question of whether to lie or otherwise deceive a highly complex question, not only in the context of business or professional life, but across the whole range of our 'private' lives as well.

We have seen that some forms of lying, deception and bullshit are ethically acceptable, but only against a background in which 'cooperative principles of good communication' are both *somewhere in play* and generally recognized as a *default aspiration* of *most* communicative acts. We have also seen that whether or not to license acts of lying, deception and bullshit depends on what harms and benefits are likely to follow, quite beyond the harms or benefits associated with the acts of

communication themselves. We shall be returning to the issue of benefits and harms that acts of deception produce both inside and outside professional institutions in later lectures. Before we get to that, however, we need to know more about the background conditions that make professional acts of deception permissible.

Do you want to be trusted?

Many people worry about the absence or loss of trust in contemporary society. And increasingly interconnected as we are, we are also ever more frequently exposed to examples of breaches and abuses of trust, with all the anxieties this entails. At the same time, millions of people go to work everyday without having their valuables stolen; millions of drivers negotiate crossings without having suicidal maniacs crashing into them, and more and more people are publicly storing information about themselves on Facebook, Instagram, Linked-in, and other online media. If trust is ever diminishing, there definitely be quite a lot of it still around. Either that, or we are very confused about the nature of trust.

In what follows, we shall try to get a little bit clearer about what trust is, or at least some of the things that trust is often said to be. We shall also consider the relationship between trust and trustworthiness. After all, trustworthy partners is one of the things we are constantly looking for - and not only in professional life. Finally, we shall consider how different kinds of trust can be important in different situations, from one-off interactions in the intimacy of our homes to long-term relationships across multinational institutions. At the end of our discussion, we shall hopefully have a better understanding of what trust and trustworthiness consist in; be in a better position to classify people as trustworthy (or not, as the case may be); be better able to form a view about what forms of trust and trustworthiness are ethically desirable in different circumstances; and have a better understanding of public concerns and about 'crises of trust', the dangers of 'audit culture', and other ways in which the notion of trust has been variously targeted in moral and political discussion.

Let's start by getting clearer about what is meant by the crucial ideas in play.

Let's say that a *subject* (such as a person) *trusts* an *object* (such as an organization) with respect to a *content* (such as not lying about their product) in a certain *context* (such as in a presentation at a trades fair). In the same way, I may trust you to hand me back the wallet that I left on the counter; I may trust you not to shoot me while we go hunting for moose; I may trust my mechanic to fix my car if he agrees to do it; I trust my dry cleaner to give a fair price; I may trust my personal shopper to pick the best vegetables; I may trust the Chancellor of the Exchequer to tell us the truth about the deficit in the annual budget; or – if I am lucky - I may trust the government not to assassinate me (pretty much no matter what).

All these trust relationships would seem to involve an *attitude* of *confidence* or *reliance*, which may or may not be consciously entertained, such as a positive feeling among friends; a belief that someone will deliver on a promise; a desire or expectation that this someone will deliver (in cases where the trust relationship is mutually advantageous); and a belief that the someone who is expected to deliver (e.g. your friend) is favourably disposed towards you, at least with respect to *the content of the trust*.

The last condition is of crucial importance to understand the nature and ethics of trust, and also a source of much confusion on this score. Thus, I don't trust my car mechanic to fix my kitchen plumbing. I don't trust the Chancellor to make me rich. Nor do I trust my personal shopper to go on delivering fresh vegetables to my house long after I have stopped paying for them. Moreover, in a game of squash I may trust you to try your hardest but certainly not to let me win. None of these facts imply that I don't trust you, nor do they imply that you fail to be a generally trustworthy person.

There are many different theories of 'the trusting attitude', but fortunately we don't have to definitely decide between them here. The fact is that different people interpret the term 'trust' in different ways, and although some of these ways are directly relevant to the ethics of trust and trustworthiness in professional life, not all of them are. (I once came across a document explaining how to re-establish the communication between two different pieces of software on a computer entitled 'Reestablishing the trust relationship'.)

It is common, however, to distinguish at least the following three families of competing views about what it is to trust. According to the first view, to trust is just to have a belief or expectation that someone or something is predictable or reliable. Thus, it is a crucial commitment of the postal service that customers can rely on their letters and parcels being delivered on time.

There can be no doubt that this 'predictability' view captures an important aspect of many cases of trust, including cases involving the behaviour of individuals and institutions in the context of professional life. One of the things we want from our business partners is to rely on them. Having said that, there are ways of caring about trust that this view fails to speak to. One example is the case where the only reason your partner is reliable is because she is being held hostage by a coercive threat (e.g. of legal proceedings, or even violence). Another example is where reliability is purchased only at extortionate cost (such as astronomical legal bills). A third example is the case of trusting someone at a first, or only, encounter in the hope that by trusting them they will decide to behave in a trustworthy way. In light of these, and other, examples some people prefer to think of predictability and reliability as relevant *evidence* that someone is trustworthy rather than something that goes to define what it is to trust. In sum, it looks as if someone could meet all the conditions for trust set out by this family of views, but still not trust. Be that as it may, if someone were to use the word 'trust' as defined by this family of views, it is not as if we would completely fail to understand what they were on about.

On a second family of views, the attitude of trusting is one we undertake voluntarily when we place ourselves in the power of another who, but for their trustworthiness, is thereby able to harm our interests. Thus, I normally trust my doctor to use my reports on where it hurts to find a way to cure my ailment rather than to find a way to kill me as quickly as possible, or (which is more likely) not pay attention to them whatsoever.

This 'voluntariness' view does a good job of making sense of cases where someone *places*, or *ventures*, trust by exposing their own vulnerability. It does less of a good job in speaking to cases where no such act is involved (e.g. the trusting attitude of a small child towards its parents), or where (if we agree this makes sense) a person has a trust (or basic confidence) in themselves. It also fails to make sense of trusting

relationships in business that, although they do involve an element of ongoing risk or vulnerability, do not imply that the trusting party must place him or herself in someone else's power, nor that the party trusted need have any favourable attitude towards them on that account. Thus, a wide variety of trusting relationships in professional life involve continuous monitoring or audit in a context of rough equality of recourse to sanctions (such as litigation). Finally, when someone places themselves in the power of another in the context of competitive or conflicting environment such as business or politics, this is as likely to be a result of the person in question acting out of fear or embarrassment, a sense of duty or obligation, or as part of a long-term commercial strategy, as opposed to being an expression of trust. If so, it is tempting to conclude that someone could meet all the conditions for trust set out by this family of views as well, but still not trust. Perhaps, then, the 'voluntarist' family of views is best understood as describing a valuable kind of dependency that sometimes involves trust, but does not need to.

A third kind of view, defended by G. G. Brenkert, arguably has a better case for describing a kind of attitude we would naturally associate with trust in a business or professional context. According to Brenkert's 'encapsulated interest' view, *to trust is to accept the risk of harm from another agent on the basis of an uncertain belief that the other does not intend to do one harm, even though they could, because of their mutual affirmation of common values or aims.*⁹ Thus, a sales representative might offer a prospective customer a free test drive, therefore entrusting the customer with an expensive vehicle, on the assumption that both the seller and prospective buyer are ultimately aiming to complete a purchase that will be to the financial advantage of both parties, without resulting in a major loss to either.

There are potential examples of trust that this 'encapsulated interest' account is poorly placed to explain. One is trusting oneself. After all, the encapsulated interests required would seem to be the interests of at least two different agents. Another example is blind, or certain, trust, such as the complete confidence that some people claim to have in their Leader, or the Deity. A third kind of example involves cases where the harm caused by a breach of trust will not involve a harm to the person trusting, as when I trust you to take care of my garden gnomes after I am dead. Finally, and most relevant for our purposes, it is not obvious that all trusting relationships involve 'common values or aims' in any interesting sense. Examples that don't may include certain adversarial relationships, one-off transactions, or trusting attitudes (if such there be) towards technology or machines.

Having said that, none of these examples show that any of the cases that meet all the conditions of the 'encapsulated interest' view are ones we would not intuitively classify as genuine trust. In that sense, at least, the 'encapsulated interest' view is arguably an improvement on the 'reliability' and 'voluntarist' views we just discussed. In any case, there is no need for us to insist that one of these conceptions of trust (or any others, for that matter) must speak to all possible cases which we would intuitively say involved a trusting attitude. So for present purposes at least, we can make do with a 'family of partially adequate families'.

⁹ G. G. Brenkert, 'Trust, Morality and International Business', *Business Ethics Quarterly*, vol. 8 (no. 2), 1998, pp. 293-317.

Let's say that a person is an *object worthy of trust* in a certain *context* just in case the *subject's* trust in that person is *well placed* in that context (and where the context in question will normally extend to more than a single, or one-off, transaction). Thus, you are a trustworthy friend on the condition that you do not betray me; you are a trustworthy mechanic on the condition that you know what you are doing with my equipment; you are a trustworthy judge on the condition that you apply the law with justice and integrity; and you are a trustworthy business partner on the condition that your decisions are guided by what is in our mutual advantage. What we want from our partners, in business as well as in private life, is to trust them *because they are*, in fact, *trustworthy*.

Trust and trustworthiness do not always go together. I can trust someone even if they are not trustworthy, as when I rely on the loyalty of a duplicitous colleague. I can trust someone who is trustworthy even if I am not trustworthy myself, as when I abuse your loyalty for my own duplicitous ends. You could be trustworthy even if I don't trust you, as might happen if we are both trapped in a professional culture of mutual suspicion. Or you could be trustworthy even if you don't in fact trust yourself or consider yourself trustworthy, e.g. as a consequence of losing your self-confidence after having made a big mistake. Finally, you could be trustworthy by never reveal that fact, or never yourself venture trust, because you are worried about the consequences and are holding back as a matter of self-protection.

So what, then, is lost by not trusting; or by not investing in supporting the trustworthy dispositions that would vindicate trust? By failing to trust, where trust would be merited, both you and your colleagues will stand to lose out on at least the following three advantages. First, by developing trusting relationships you will be able to pool your resources, for example by dividing the work of knowing what's what about different areas of activity and relying on each other's word, or testimony, instead. Second, you will be able to protect yourself and others against a wide range of harms and risks of harm by having your trustworthy partners look out for you when you are not able to do so yourself. Third, you will be able to save energy and resources on conflict, threats, monitoring and policing, such as the endless contracts that people come up with in situations where they are desperately trying to cover every loophole that their assumed adversary is assumed to be looking for. When put this way, the advantages of trustworthy relationships do not only seem obvious, they would also seem to describe some of the most basic presuppositions of any minimally decent society.

It would be wrong, however, to conclude from this that more trust is always better. We know from experience that excessive trust can itself generate vulnerabilities in response to potentially predatory behaviour, such as aggressive sales tactics. The informality and comforts of excessive trust have also been known to produce complacency or malpractice, such as the institutional tolerance of friendly favours or 'soft corruption' in environments that suffer from the absence of real, or effective, accountability. Furthermore, the fact that many institutional contexts are essentially either *competitive* (such as markets) or *conflictive* (such as politics) suggest that there are significant limits on the extent to which *specific forms* of trust should be cultivated across the entire range of professional life. Thus, it would normally be unwise to trust our political enemies with knowledge of all our internal conflicts and weaknesses, even though it might still be reasonable to trust them to 'play fair', or

follow ‘the rules of the parliamentary game’ in the context of democratic debate or competition. Once more, the crucial thing to bear in mind is the precise *content* of the trust relationship in question, as it is found in a particular institutional *context*.

To draw out the contours of how different kinds of trust, we could do worse than follow Brenkert’s distinction between what he calls ‘basic’, ‘guarded’ and ‘extended’ trust. By ‘basic trust’, Brenkert means the trust that others, with whom we may have only impersonal relations, will not take undue advantage of us. This is the kind of trust that most of us are willing to show every day while negotiating traffic or sharing public transport. To benefit from basic trust, we require little more the *consistency* and *mutual common motives* from each other, such as the fact that we all need to restrict our physical movements in order to make it safely home.

Guarded trust goes beyond basic trust and relates to direct interactions between people, such as business transactions, negotiating settlements and mutual assistance. What people guardedly trust is that unclarity, ambiguities and uncertainties can be resolved without conflict, sanction or a breakdown of relations. Apart from consistency and mutual common motives, guarded trust requires *competence* and relevant *knowledge* of the interactions involved, such as the nature and prices of the products being traded or the content of the laws and regulations that govern the exchange. Guarded trust may well be the default attitude between colleagues and competitors in business and politics, at least so long as the system is not irredeemably corrupt, or otherwise in danger of breaking down.

Extended trust is the reliance that others will not take advantage of you even in the absence of monitoring, sanctions, and other self-protective practices. It is the kind of trust that involves information sharing; open access, voluntary contributions and the absence of records. Beyond consistency, mutual common motives, competence and relevant knowledge, extended trust requires *openness* and *informality* between truster and trustee. It is the kind of trust that many of us associate with relationships between friends. According to Brenkert, it worth considering extended trust as an aim in professional relations, not only on ethical grounds, but on grounds of cost and efficiency. The thought is that by removing some of the barriers to cooperation that characterize relationships of guarded trust, we can improve communication between competitors and strangers, overcome linguistic and other cultural barriers, and cut at least some of the costs associated with seemingly ever expanding legal provisions and control. To this extent, it is tempting to think of extended trust as sitting on top of an extended pyramid the pinnacle of which we shall only be able to reach once we are able to let our guard down’ and cooperate like friends.

Tempting as this may be, however, we have already noted some of the dangers involved in removing the safeguards associated with guarded trust. In highly competitive environments there will always be some risk of extended trust leaving us vulnerable to various forms of opportunistic or predatory behaviour by others. And as the history of institutional corruption clearly shows, the informality and lack of monitoring that extended trust involves may severely undercut our ability to trust even ourselves. It is probably unwise, therefore, to imagine extended trust as some kind of universal professional ideal. What we have instead is a menu of options, some of them more ambitious than others, which we must choose between depending on the comparable risks and benefits in the case at hand.

The integrity question

Public figures sometimes risk their reputations or their jobs by visibly failing to behave with professional integrity. Some public figures leave their jobs to preserve their integrity. Some of these jobs are very important and extremely well paid, and the people in question have sometimes worked for years to get them. Integrity is obviously a very important thing. Yet what exactly is it for someone to show integrity, and why does it matter so much? Is it always right to prioritize integrity even if it is bad for business, or if it conflicts with other professional aims? Do conflicting institutional incentives always have a tendency to conflict with integrity, or are all such conflicts in principle resolvable? And is it possible to be a person of integrity in one area of life (e.g. at work) but not in another (e.g. in private relationships)? These, and similar, questions are seldom far from the surface in discussions of whether or not some politician, say, should be forced to resign when they have been found to be dishonest in some personal capacity (such as plagiarizing their doctoral thesis, or having an extramarital affair).

In order to think clearly about these questions we need to have some view about what integrity actually consists in, and whether it means the same thing in different circumstances. As we will shortly see, it would be a serious mistake to always assume that it does. Armed with some idea of what integrity consists in, we may hope to be able to better classify each other as showing or not showing integrity in different circumstances; to form a clearer view about which forms of integrity are ethically advisable, and where (e.g. in politics versus private life); and be better able to evaluate the public concerns about integrity we constantly encounter in public life, from journalists talking of a ‘crisis of integrity’ in politics to reluctant citizens refusing orders from the state on the grounds of ‘conscientious objection’.

Any plausible account of integrity has to address at least the following two questions. The first question is what it means for an act or sequence of behaviour to show integrity. This is the kind of integrity we are interested in when we consider the behaviour of professionals who refuse to take a bribe; admit to faults in their risk management strategy; or refuse to victimize their colleagues for personal gain. The second question is what it means to someone to be a person of integrity. This is the kind of integrity we are interested in when we consider the behaviour of professional who are consistently able to show resistance to temptation; reasonableness in their dealings with colleagues and clients; a good understanding of their professional role; honesty in communication; and consistency in behaviour (their treating like cases alike in an ethically sensitive way).

A plausible account of integrity in professional life needs to make sense both of *acts* and *persons* displaying integrity; the *relation* between the two; and how the attribution of integrity to acts and persons occasionally seems to come apart in some situations. Thus, we do not want to rule out the possibility that a person who has been historically lacking in integrity could show at least some integrity on a particular occasion. After all, this could be their first step on the road from being a dishonest

crook to being an upstanding citizen. Likewise, we do not want to rule out the possibility that a person of integrity could lapse on a single occasion. After all, there are very few people, if any, who never give in to any temptation, ever. When thinking about integrity in a systematic way, we should therefore bear in mind the need for flexibility that is necessary in order to make us able to apply the term ‘integrity’ to real human beings.

As with all contested terms of ethical appraisal, there is quite a variety of minimally plausible theories of what integrity consists in.¹⁰ As always, we do well to remember that that no-one can force us to pick just one of them. Different accounts may have their uses in different circumstances. The following sample arguably gives a reasonable overview of what our options are:

At one end of the spectrum, there is a family of views that is best thought of as targeting the person of integrity as a whole. Thus, according to Greg Scherkoske, to show integrity is to live a ‘convincing life’, where living a convincing life involves being a person of sound convictions, who shows due respect for their own capacity to form sound convictions, a disposition to act resolutely on them, while also being disposed to consider reasonable challenges from others, e.g. in the context of discussion or disagreement. This definition has all the standard strengths and weaknesses displayed by in this family of views. On the one hand, it describes what is undoubtedly a species of admirable, or even virtuous, person. On the other hand, it fails to speak to the ethics of integrity in some particular situations, such as the case of the manager who shows integrity in some contexts not others; the dutiful professional who is able keeps her moral flaws away from the office; or the obvious significance of everyday and run-of-the mill transactions such as your banker refraining from pocketing your deposit.

Much the same applies to similar views, such as the idea that showing integrity is *being true to yourself*, or showing a *coherent* (and favourable) *disposition* towards what you value or approve of over time. As important and admirable as we all agree it is both to be true to oneself and to coherently display one’s values over time, many professional choices concern issues that do not affect our sense of self in any deep way. Whether you truthfully own up to the fact that you broke my whiteboard marker would normally not be thought to have that kind of profound significance for either of us.

Furthermore, the issue of consistency is a tricky one to handle in the context of professional and other social roles, where the values in play will often conflict quite significantly as you move from one role to another. It is by no means obvious that every professional role holder either should, or could, aspire to some maximal level of integration.

Third, a person whose deliberative focus is consistently on their sense of self, or on whether their behaviour is observable to be consistent with their values over time is a

¹⁰ See e.g. G. Scherkoske, ‘Two Cheers for Integrity?’, in his *Integrity and the Virtues of Reason: leading a convincing life*, Cambridge: Cambridge University Press, 2013, pp. 1-32.

person who is therefore to a very large extent concerned with themselves, as opposed to other features of their circumstances (such as the number of livelihoods that are at stake) that could also be ethically relevant. It is far from obvious that this is always a mark we should associate with integrity. What we need to make sense of these cases is a view that speaks directly to the integrity of *acts* or multiply changing individual transactions, as opposed to the quality of someone's personality as a whole.

At the other end of the spectrum is a family of views that speaks more directly to the idea of integrity as it applies to individual actions. According to one such view, to show integrity in a certain situation is not to be vulnerable to temptation, weakness of will, or actions that are contrary to our character or considered judgement. An obvious example from a professional context would be the official who refuses to take a bribe, or the business colleague who refrains from joining in a scheme of deception or embezzlement. Among the virtues we admire in professional life, the ability to hold one's ground and not fall into temptation are clearly among the most important ones.

The obvious problem with these views as just stated is that having a strong will or standing by one's judgement is equally compatible with wrong, or otherwise inadvisable, action; so on this view integrity is in principle consistent with bad character. It follows that a ruthlessly corrupt judge and a manipulative turncoat could both be paradigm examples of people who act with integrity. Somewhat more esoterically, it is sometimes argued that a person can display integrity in virtue of their *weak will*, as in the case of a soldier who refuses the order to shoot an innocent civilian, but who can't help but feel deeply guilty about it afterwards. (Philosophers have labelled this 'inverse akrasia', which basically means weakness of will in reverse. Whether we should describe this case as a straightforward example of acting with integrity, or instead as some other kind of ethical achievement (assuming it is ever an achievement at all), is an open question we do not have to settle here. Either way, there is arguably something about integrity that makes it different from merely showing strength of will or standing up for one's judgement, whether in professional life or elsewhere.

Perhaps we can improve on this family of views by imposing an ethical constraint on the kind of steadfastness in will or judgement that we are willing to associate with genuine integrity. On this view, to show integrity in a certain context is to stand up for one's own judgement above some threshold of ethical decency in that context, such as a baseline determined by what we think that other people could reasonably expect. In other words, integrity does not require ethical perfection, moral virtue, or even a stable and coherent disposition to behave in accordance with one's values over time. It may, however, require you not to be a 'selfish bastard', or a sociopathic monster. Examples of the kind of behaviours we would normally consider as falling beneath the threshold would be smears, threats, or negative briefing of colleagues. Bribes and embezzlement would also be candidate behaviors for failing to meet the threshold, although in certain contexts (involving extreme danger or severe deficits of trust) we might have to adjust the threshold to be compatible with survival in the circumstances. So one of the trickiest features of this view is how to decide on the ethical baseline, and on what grounds.

As different as the views of integrity are that we have just discussed, most of them have at least one thing in common; namely the high value they assign to the person

who is able to stand up for their judgement. When that judgement is sound, perhaps they can always be said to act with integrity, and if it generally is, then maybe we should think of them as a person of integrity as well. On balance, this not be an entirely hopeless view; as long as we bear in mind that standing up for one's judgement, however sensible, can sometimes also be an expression of arrogance, stubbornness or coercive behaviour. Indeed, in certain kinds of circumstances, standing up for one's judgement no matter what is a recipe for individual or collective disaster. No doubt the reader can think of some real life cases for her or himself.

Integrity, then, is an important value in professional life. Integrity is *intrinsically* valuable, or valuable in itself, in virtue of expressing the ethical standing of those who display it. Yet as the occasionally negative effects of resolutely standing one's ground amply demonstrate, even if integrity is an *intrinsic* value, that does not make it a *supreme* value, or valuable without qualification. The value of different forms of integrity is both *context* and *role* specific. There is an ethical difference between being a conscientious objector and an opportunistic free-rider, even though both may be associated with consistently standing up for one's beliefs. Likewise, there is an ethical difference between being willing to compromise with a business competitor and with someone on your team. The fact that you are more disposed to 'bend' in some contexts rather than others need not undermine your claim to value integrity for its own sake.

Of course, integrity is also instrumentally valuable, or useful, in professional life, as it is (almost) everywhere else. In an institutional environment where integrity is rewarded, showing integrity may further your professional advancement and serve your business objectives. And just as with dispositions of trustworthiness, discussed in the precious lecture, actually showing integrity is often the best way of acquiring a reputation for it. Indeed, if we 'instrumentalize' integrity too much, or are thought by others to do so, we can thereby undermine its usefulness, thereby making the entire exercise practically self-defeating. Indeed, the complete instrumentalization of integrity is arguably paradigmatically non-universalizable project, insofar as it is likely to *presuppose* the non-instrumental valuation of integrity on the part of at least some of the relevant stakeholders. (Why attempt to gain a reputation for something that people don't value?) When thinking things through, therefore, we are often likely to find that we, and our colleagues or competitors, actually value professional integrity for its own sake *as well as* for its instrumental value.

Some professional roles arguably require more by way of personal integrity than others. (Medical professionals, judges, religious representatives and (even possibly) politicians are examples that often spring to mind.) Perhaps it is true that in some of these roles the value assigned to integrity approaches something like that of a personal ideal. (After all, for some of our professional their colleagues, their line of work is also their vocation, or 'calling'.) Yet as our discussion in this lecture has hopefully shown, a wholly person centred conception of integrity is arguably too demanding to be plausible in all situations where integrity matters, and therefore abstracts excessively away from at least some of the institutional contexts or professional encounters where the integrity question matters.

9 Who's to blame?

If lie to you, or I am otherwise dishonest, this may be a cause of serious grievance on your part, and on the part of our colleagues and clients. But at least you can lay the blame squarely at my feet, at least if we assume that I was acting voluntarily and not in a state of confusion or duress. Where groups and institutions are involved, however, things are often not that simple. Thus, there might be an accident on the shop floor that happens to take place on my watch. But the fact that it was me who was there could be a matter of bad luck and the damage caused could be way out of my control. Further up the institutional hierarchy, managers and executive officers come and go, collective decisions are made by committees with significant numbers of dissenting voices, and orders and instructions are passed down along complex networks of authority where it is simply impossible for every participant to know what every other participant is doing. Yet when organizations do well or badly (in economic as well as ethical terms) we find it natural, indeed irresistible, to attribute praise and blame to at least some of the participants involved. How, if at all, is this possible? And does it really make sense? These are some of the questions that will occupy our discussion in this lecture and the next.

One of the difficulties we immediately have to face is that what we the behaviour of organizations as though both those organizations and the individuals who work in them act. Yet it is not an entirely obvious matter how organizations are supposed to 'act', over and beyond the actions of the individuals who make them up. Furthermore, if we are happy to say that organizations as well as individuals act, should we thereby also be happy to say that organizations can themselves be held ethically responsible? And assuming that we are happy to say this, should we also be happy to say that individuals and organizations should be constrained by the same rights and duties? A careful look at the history books is likely to show that the answer to none of these questions is entirely straightforward. After all, people are frequently loyal to, or angry with certain institutions (such as the anger felt towards a multinational following an industrial disaster). And representatives of institutions (whether companies or states) frequently collect awards or apologize on their behalf. At the same time, most readers of these words would probably balk at the thought that registered companies have a right to life or that sports clubs should be granted the same rights or legal protections as ordinary citizens. So there is definitely a difference in how we regard the ethics of individuals and institutions, even though the details of how we do so clearly varies considerably across time and space.

In order to focus our minds on these questions we need to better understand what corporate (as opposed to individual) agency consist in across a range of different cases. Only when we have a better way of assigning actions and outcomes to individual or corporate agency in practice can we get into a position to form a view about which forms of corporate agency are ethically admirable or problematic, and thereby understand the nature of public concerns about corporate behaviour, for example in discussions of corporate social responsibility, or 'corporate citizenship'. We should therefore start with the question of how institutions such as business corporations can ever truly be said to act at all.

According to a *reductionist* view of corporate agency, genuine agency basically consists in intentional, purposive behaviour, explicable in terms of reasons that can at least in principle be self-consciously formulated and understood; embedded in a psychology of beliefs, desires and emotions; embodied and controlled by a single individual psycho-physical entity, such as a human person. While at least some human individuals arguably meet these criteria, but corporate entities arguably do not. It follows that all group, collective and corporate agency can therefore only be understood in terms of individual agency and what individual agents do together, such as a group of business executives engaging in a business transaction, or a group of likeminded people getting together to play a game.

The problem with reductionist views is that they are too narrow, or demanding, along too many different dimensions. First, some forms of individual agency arguably fail to meet the above criteria because they are pre-rational (as in the actions of some non-human mammals); pathological (as in the actions of some humans with mental disabilities); or artificial (as the case of in robotic agency). Second, not all agency obviously involves anything like self-conscious or controlled deliberation because it is performed by groups that are loosely and improvisationally structured and lacking a central source of control. (A rioting mob or a hiking party would be common examples of unstructured agency of this kind.) Third, in many of the cases that really matter, the structure of institutional agency fails to resolve itself into a neatly structured series of steps that is ultimately traceable back to the reasoned and voluntary transactions of rational individuals. Not even democratic politics really works like that. And where large scale and hierarchical bureaucracies are involved (such as in business corporations or government departments) the individuals involved at any given time are frequently replaceable by (almost) arbitrary others in an institutional culture that often seems to have a 'will' of its own. The most reasonable conclusion to draw from these, admittedly inconclusive, considerations is that reductionist views are implausible because they ask too much of agents for them to count as 'genuine' agents.

According to *non-reductionist* views, on the other hand, there clearly can be such a thing as genuine corporate agency. On this kind of view, all genuine agency really requires is the presence of systematic, deliberate and controlled behaviour truly describable as performing some intentionally designed purpose or function, and where this behaviour would normally *supervene* (i.e. strictly depend) on the underlying actions of individuals acting within an intentionally manipulable structure, such as a club, society or a business organization. (The non-reductionist view is consistent with the view that the individual actions on which corporate agency depends also satisfy the conditions of agency proposed on the *reductionist* view.) Thus, non-reductionist views effectively relax the criteria that agents have to meet in order to count as agents, and are thereby able to accommodate such central cases of corporate agency such as a board of directors deciding on an implementing a business strategy consistent with the purpose of the firm; an electorate reaching a policy decision by means of a (direct or indirect) voting procedure that is mediated by officially sanctioned representatives; and a company producing and selling a product (such as an investment bank designing and marketing an investment vehicle designed for private pension plans).

The criteria for agency implied by non-reductionist views are arguably met both by

many human individuals and by structured groups such as governments and business corporations. The obvious problem with non-reductionist views is that in the real world things rarely happen exactly according to plan, or in accordance with how the behaviour of the relevant organization is 'ideally' described, or modeled, in the abstract. Thus, few, if any, institutions are entirely 'pure' in how they execute their management structure. Does that mean that we must deny that either hybrid institutions (because some decisions accord with the model and others not); evolving institutions (because not clearly fitting clearly into a defined category at a given time); or 'corrupt' institutions (because non-compliant with the relevant 'design specifications) really act? That would be a pity, if the point of being able to speak of corporate agency was to be able to hold corporations ethically responsible. Furthermore, some of the social causes and effects involving institutions could be better explained by attributing them to economic and other social forces that are structural, impersonal, or otherwise non-agential (e.g. cyclical market fluctuations or environmental changes). To think that by lowering the bar of agency we shall thereby be able to account for all the hard cases that seem to escape reductionist analysis is clearly much too simple. Having said that, we should not be tempted to conclude that just because the domain of genuine corporate agency might be smaller than we hoped for we might just as well give up and say it is completely empty. In some cases, no doubt, actual institutional behaviour does approach the ideal model to such an extent that assigning agency to the relevant institution is perfectly reasonable. (Consider the unanimous decision of a board or directors to accept a complete take-over by a rival firm.) And just as with purely individualistic (and so reductionist) explanations in terms of the behaviour of 'rational and unitary selves', the inevitable presence of some element of idealization in explanations involving corporate agency does imply the falsehood of the explanation involved. (In this respect, explanations of corporate agency are no different from other explanations in the human or natural sciences that appeal to idealized models, such as vacuums, full information, or rational consistency.)

The nature of corporate *agency* is an interesting issue in its own right. Yet our main interest in it here, and the interest that has driven most of the literature on the subject, is its role as a necessary condition of corporate *responsibility*, a condition in the absence of which all responsibility attributed to corporations would ultimately have to be distributable without remainder to the individuals who compose them. We therefore need to ask: does it really make sense to hold corporations ethically responsible for anything, over and above the ethical responsibility we attribute to their members?

Let's call the view that corporations can be held ethically responsible *the social view*, and the view that only individuals can be held ethically responsible *the individualist view*. In common with most of the literature on this topic, let's assume for the moment that the individualist view is the default view, to which we must retreat if it turns out that the social view is faced with insuperable difficulties. After all, the social view does not say that individuals cannot be held ethically responsible. What it says is that corporations can be held responsible as well as individuals. We should therefore examine the pros and cons of the social view and decide whether corporations as well as individuals can be held ethically responsible, or whether we should retreat to the individualist view.

There are two kinds of arguments in favour of the social view, the first of which appeals to how we actually do relate to corporations as a matter of fact; the second of which appeals to how we should relate to corporations in order to promote and preserve what we think of as ethically important, or valuable.

The first kind of argument appeals to the following facts about how we actually do relate to corporations. For example, corporations are punished (e.g. by being fined or excluded from certain activities); they are praised and blamed (e.g. by grateful or grieving customers); they claim and recognize praise and blame (e.g. when advertising their ethical profile, or offering apologies to persons harmed); they take responsibility for critical events (e.g. by responding to emergencies or investing in ethical products); and they revise their ethical profile and reputation in light of reputational effects and ethical arguments inside the organizations themselves (e.g. about how to formulate and realize their 'business values'). In sum, there is a very strong descriptive, or 'phenomenological', case for the claim that we relate to corporations as *bona fide* ethically responsible agents.

The second kind of argument appeals to the *ethical* (and other practical) costs of not relating to corporations as *bona fide* ethically responsible agents, the conclusion being that rejecting of the social view involves an ethical cost we should not be prepared to accept. The claim that corporations need to be held ethically responsible mainly derives from the idea that there are cases where ethical responsibility needs to be distributed, but where there are no individuals available on whom the burden of responsibility can fairly be placed. Among the most important considerations here are the undesirability of imposing burdens on individuals that they, alone, cannot realistically bear (which argues for some kind of alternative arrangements whereby people working together can pool their risk); the need to protect employees from bad circumstantial luck (such as the case of the damaged product that just happened to come through on your shift); the need to protect employees from the consequences of bad luck (such as claims for compensation from the clients who have bought a damaged product); the need to distinguish the reputation of the company 'brand' from the person who happens to represent it at any given moment (such as a customer service agent working under a pseudonym from an off-shore call centre); and the need to incentivize individuals to take on jobs that are either potentially unpleasant or dangerous to their clients (such as invasive surgery or certain forms of high-speed travel).

If the division of responsibility between individuals and institutions is an ethically inescapable fact, that still leaves the question of how to distribute that responsibility and what kinds of responsibility we have in mind. By making 'the company', as opposed to one of its employees, pay for the damage caused we are no doubt registering our grasp of some idea of corporate responsibility, in the sense of *accountability*, but what, if anything, does this have to do with our intuitive grasp of *ethical responsibility* in particular? The answer to this question is much less clear.

If we start from the idea that the paradigm case of ethically responsible agency is the normal and mature human individual, it is easy to find the idea of corporate ethical responsibility obscure, or even incoherent.¹¹ For example, we might worry that

¹¹ M. Velasquez, 'Why Corporations Are Not Morally Responsible for Anything They

corporate agency is not *unified* enough to justify the attribution of ethical responsibility for its behaviour. The argument would be that only unitary beings capable of acting on their own beliefs and desires (such as human individuals) can be held ethically responsible. Corporations and other institutional ‘agents’, not being individually ‘minded; in the right way, are not unitary beings in this sense. Therefore, corporations and other institutional agents cannot be held ethically responsible for anything, although they can, of course, be held accountable in other ways, such as paying out financial compensation for breaking the law or the harms they cause.

There are at least two limitations to this argument; the first of which is obvious, the second of which less so. The obvious limitation is that how unitary is unitary ‘enough’ is itself a controversial question. Thus, it is often pointed out that many corporations and other institutions achieve very high degree of unity in their agency by being organized around highly specific internal decision making structures (such as management hierarchies or voting mechanisms), as a result of which their behaviour is both predictable and responsible to as a single ‘entity’ that exists over and beyond the activities of its individual members at any given time.¹² The democratic state, with its bureaucracies and decision-making structures is one example that is frequently used to illustrate this point. The modern business corporation is another.

The second, and less obvious limitation of the argument is that in drawing the contrast between individual and corporate agency, it is easy to overestimate the degree of unity of individual agency. Perhaps it is true, as some philosophers argue, that when individuals act there is at least one relevant description of what they are doing under which they know that they are doing it (such as walking through the corridor as quickly as possible). Yet we should not forget that there are often many descriptions of what we do under which we do not know what we are doing (such as being rude to our colleagues by ignoring them), and that we are very often divided against ourselves whether we like to admit it or not (such as when I explain my action of rushing through the corridor by saying that I am just very busy all the time). It is very far from obvious, therefore that appealing to the unity of individual agency is enough to establish that corporations and other institutional agents cannot be held ethically responsible for what they do.

A second argument is that corporations and other institutional agents cannot be held ethically responsible because they are not *free* in the same way that individual agents are (or that they are therefore not really agents in the same way as are mature and statistically normal human individuals). The basic thought behind his argument is that genuine ethical responsibility requires the freedom to do otherwise. Organizations and other institutional agents, however, are unable to act otherwise than they do because their behavior is determined by the operations of external forces and their internal decision-making structures. So corporations and other institutional agents cannot be held ethically responsible.

Do', *Business and Professional Ethics*, vol. 2, pp. 1-18. 1983.

¹² C. List & P. Pettit, *Group Agents: the possibility, design and status of corporate agents*, Oxford: Oxford University Press, 2011.

As tempting as it may be to fall for this argument from ‘corporate determinism’, we should really not fall for it as long as we hold on to the possibility that there is any such thing as free agency at all. (To settle that issue would obviously take us too far afield here, and would not in any case decide the question whether there is a crucial asymmetry between individuals and corporations in this respect.) First of all, we should remember that so long as the freedom of the individuals who make up the corporation is assumed, there is nothing to stop the corporation to act otherwise than it does in virtue of its executive officers collectively exercising their freedom to turn its activities in one direction rather than another. There is nothing either mysterious or particularly ethical about this fact. After all, businesses make decisions to change their activities all the time, for example when they decide to develop a new product, or to take an existing product (such as a malfunctioning passenger jet) off the market. To this extent, we can think of the freedom of corporations and other institutional agents as arising from the freedom of the individuals within it. Second, and more importantly, it is simply not true that the behaviour of corporations is strictly determined by external forces and the nature of their internal decision-making structures. In fact, corporate behaviour is frequently seen to diverge from both internal and external constraints, as when corrupt or nepotistic practices of promotion and reward first develop and then get addressed, or when business is restructured in response to a predicted change in the economic cycle. Appealing to freedom does therefore not offer any knock-down arguments against the idea that corporations and other institutional agents are ethically responsible.

Similar considerations apply to the argument that corporations and other institutional agents cannot be held ethically responsible because they lack the kind of ethical *understanding* that individuals have in virtue of our possession of emotions and affective dispositions (such as empathy) that put us in a position to adopt a distinctively ethical perspective on the world.¹³ The thought behind this argument is that organizations and other institutional agents, having no comparable inner life, or ‘soul’, are unable to exhibit ethical understanding and are to that extent comparable to persons with severe mental disabilities, such as the stereotypical sociopath who pursues her or his ends without any consideration of the benefits and harms their actions may have on vulnerable others.

The limitations of this argument are very similar to the limitations of the argument based on the idea of freedom we have just discussed. If corporate agency strictly depends, or ‘supervenes’ on the agency of the individuals who make up the corporation at any given time, then there is no reason in principle to deny that the ‘supervening’ entity can exhibit features of agency, relevantly analogous to the features of individual agency, as an ‘emergent property’ of those individual features when found in a certain dynamic relationship. To put the point in plain English, the capacity for the corporation to have the ‘inner life’, or ‘soul’, that allows it to display genuine ethical understanding derives from the display of genuine ethical understanding of the individuals that compose it. In this way, for example, we can describe an industrial company as showing ethical understanding of its effects on its

¹³ S. Wolf, ‘The Legal and Moral Responsibility of Organizations’, in J. R. Pennock & J. W. Chapman (eds.), *Criminal Justice, Nomos 27*, New York: New York University Press, 1985, pp. 267-86.

neighbours when it ceases to pollute the local water supply, in virtue of the empathy or concern shown by the company executives who make the decision to manager the company waste in an environmentally more responsible manner. In this respect, the case of corporate agency is crucially different from the case of any individual sociopath, stereotypical or not. For whereas the mind of an individual sociopath is not composed of a large number of interacting mini-minds each of which is perfectly capable of showing empathy and other psychological prerequisites for ethical understanding, at least some business organizations and other corporate agents most definitely are.

For these and similar reasons, any case against the social view based on the differences in capacity between individual and corporate agents is likely to be inconclusive. The real ‘action’ between the social view and the individualist view takes place within the ethical domain itself, and so is really a question about rights and wrongs, and what we value. There are at least four different ethical considerations that defenders of the individualist view might raise against the social view. Although none of these considerations are sufficient to refute the social view, they arguably do establish that any plausible ethics of organizations will need to strike a balance between ethical responsibility attributed to corporations and other institutional agents, and ethical responsibility attributed to the individuals who compose them.

The first consideration relates to what we might call the problem of *parasitism*. Holding corporations, as opposed to individuals, ethically responsible for the harms those corporations cause allows culpable individuals take serial advantage of organizations, by moving from one institution to another, and leaving other people behind to deal with the mess. (This is one complaint that is frequently made about high-level business executives who may leave one organization after a scandal with a sizeable pay-off, only to produce a similar kind of damage in another company, before moving on again with yet another sizeable pay-off, and so on.)

The problem of parasitism is, undoubtedly, one of the most serious ethical problems facing large-scale organizations. Yet, quite obviously, it does not follow from the fact that we attribute some element of ethical responsibility to corporations and other institutions that we cannot *also* attribute ethical responsibility to individuals (including obvious cases of ‘parasitic’ CEO’s). Furthermore, we should not forget that attributing responsibility to corporations certainly can, and frequently does, have palpable effects on the individuals within it. Thus the repercussions for a company of being involved in a scandal (whether in terms of its reputation or its ‘bottom line’) will in some cases transfer, even if indirectly, to some of the agents within it associated with that scandals (even if not always the right ones). Thus, not all attempts at corporate ‘parasitism’ are equally successful. And some individuals involved in serious wrongdoing will never work again.

A second consideration relates to the fact that individuals (and groups) are frequently *complicit* in serious harms or other forms of ethical wrongdoing that they have no direct involvement in, or that are otherwise not of ‘their own making’. Yet at least for some such individuals, if they refuse to approve of what a corporation decides jointly, they can change the pattern of corporate behaviour or, if not, they may be able to leave the corporation altogether. (Within limits of sustainability, this is one of the key differences between being an employee and being a slave.). On these grounds, it

might be argued that what we have so far called *corporate* responsibility should instead be thought of as a kind of collective individual responsibility, with the potential (at some ideal limit) that corporate responsibility is ultimately replaceable by individual responsibility without remainder.

The problems of complicity are arguably at least as important to corporate ethics as the problem of parasitism. It is also a problem that is comparably ill understood. Yet as important as the problem undoubtedly is, we should not be carried away lest we forget that attributing ethical responsibility to individual agents who are complicit in corporate harms is perfectly consistent with also attributing ethical responsibility to corporations considered *as corporations*. (We already made the same point about the problem of parasitism in the paragraphs above.) Furthermore, all institutional decision procedures have *unintended* or *unforeseeable* effects (such as changes in demand or the behaviour of competing organizations) the ethical burdens of which is not obviously best shouldered by individual employees as opposed to the corporation as a whole. And finally, it is not always obvious when either ‘protest’ or ‘escape’ is a realistic possibility for the individuals involved in causing a corporate harm. (Managers and employees of unethical businesses that offer their only hope for a basic income is an obvious case in point.) To place the full weight of corporate ethical responsibility on the idea that everyone involved in an organization is somehow complicitly ‘involved’ in its behaviour would therefore arguably be both implausible and self-defeating.

Having briefly considered the problems of parasitism and complicity, we are now in a position to evaluate a third consideration in favour of the individualist view, namely that the distribution of ethical responsibility to individuals is justified by the way in which ethical risk either can, and frequently is, priced into the salaries and benefits of employees and other participants in corporate activity depending, for example, on their place in the decision making hierarchy and other facts about their influence, power or economic safety. Thus, we are all familiar with the case of the CEO who is forced to leave a company after some large scale disaster, such as a massive oil spill, which it is quite obvious that they, personally, had no power to prevent. Yet making the axe fall on them is in no way unjust, we might think, if the risk of ‘taking the fall’ is duly built into their pre-agreed contract of employment. Perhaps we are here dealing with a case where the ‘responsibility’ taken is not intuitively best conceived of as genuinely ‘ethical’, as opposed to financial, legal, or otherwise ‘practical’. Whatever we may think about when the word ‘ethical’ is, or is not correctly applied, however, there is no doubt that even in these cases the ‘genuinely’ ethical is not too far away. After all, one way of being an ‘agent’ is to be the ‘agent of’ someone else. And in all cases where people are acting on the behalf of others, or on behalf of an organization, there will be a question of how the ethical responsibilities in question are most fairly distributed. It would be naïve, or narrow-minded, to think there can only ever be one answer to that question, and that the responsibility of the ‘fall-guy’ CEO whose ‘head must roll’ in the aftermath of an environmental disaster does not have a genuinely ethical aspect.

Having said all that, the point remains that there are cases where distribution of ethical responsibilities between organizations and the individuals who compose them could still seem perfectly appropriate. We have already noted, in the context of the problems of parasitism and complicity, that corporate actions have consequences

which are not intended, foreseen, or even foreseeable. There are also limits on how heavy the ethical burden a single person should be asked to take on is in light of a corporate harm, unless they can be seen to have been a true ethical monster. (The history of large-scale bureaucracies, many of them created and maintained by nation states, clearly demonstrate that there are sadly too many examples of this kind.) Finally, we should note that the very idea that all relevant ethical risk is built into the salary and benefits of the individual employee is obviously based on a massive idealization. In the real world, where we actually praise and blame individuals or organizations this is very frequently not the case. And when it is not, it is both natural and perfectly reasonable to shift the focus of ethical responsibility from the single individual to the corporations, institutions, structures or networks in which they play a part.

A final consideration in support of the individualist view, and one we should not underestimate, relates to the fact that where responsibility and duties reside there is normally also corresponding privileges and rights. Suppose we think that corporations and other institutions should be assigned the same *duties* as individuals, must we also think that corporations and other institutions should be assigned the corresponding *rights*? Adult citizens have a right to vote, yet we hardly hear those who wish to hold organizations ethically responsible arguing that we should enfranchise corporations (as opposed to blame, punish or abolish them). Individual human beings are said to be protected by rights to live, liberty and the safety of the person. Yet we hardly hear those who wish to hold corporations responsible arguing that when businesses or other institutions are abolished or 'go under' the people responsible for doing so have infringed their right to life. If rights are as tightly connected to duties for organizations as they are for individuals, the idea of corporate ethical responsibility is surely based on a mistake.

On balance, I think we should take this challenge seriously, but without losing sight of the wider perspective. As a matter of fact, many people have historically defended the idea that some institutions or structured collective have rights as well as duties that are at least on a par with the rights we unreflectively assign to human individuals. Whether we are talking of valued traditions (such as vulnerable minority cultures); treasured institutions (such as organized religious groups); or cultural artifacts (such as museums or the performative arts), we normally have little difficulty getting our heads around the idea that there are certain corporate, institutional or otherwise collective entities that are worthy of special protection in very similar ways to how individuals are. The issue is arguably not so much the coherence of this idea, as whether in enacting these protections we are assigning the corporate or institutional entities in question too much *power* over the individuals whose life they affect. (This is one reason why, in the context of democracy, individual citizens have the power to kick out the legislators at regular intervals if they do not like the kind of institutional agents they are creating or maintaining). Furthermore, even in the case of individuals, it is not obvious that every duty must correlate with a corresponding right. Thus, we might think, for example, that parents have duties to their children that extend far beyond their rights over them. (The obvious alternative is to think that a child is the *property* of its parents; an idea that, even if not historically unknown, is hardly a natural premise in an argument for corporate ethical responsibility.) Our conclusion, therefore, should be that while the problem of rights and duties does much to complicate our conception of corporate ethical responsibility, it does not completely

undermine it.

We started off this lecture by asking what the conditions are for something to be either an agent or a responsible agent, and then went on to ask whether corporations and other institutions ever meet the relevant conditions. Historically speaking, this is the standard way of addressing these issues. It is not, however, the only way of doing so. Another way of addressing the issues of corporate agency and responsibility is to observe that we undeniably think and talk in terms of corporate agency and responsibility, and then ask what agency and responsibility must be like in order for it to make sense for us to do so. The first approach to the question arguably carries with it the danger that corporate agency and responsibility takes on the appearance of something that is too mysterious to exist. The second approach carries with it the danger that we come to take the ideas of corporate agency and responsibility too much for granted. The most reasonable solution, I would suggest, is that we make use of both approaches together; and that we allow ourselves to make critical adjustments to our idea of corporate agency and responsibility when the results of the two approaches conflict. The outcome of this process is uncertain and open ended. Yet we should not assume in advance of the process that we must end up rejecting the idea of either individual or corporate ethical responsibility entirely in favour of the other.

In search of the corporate citizen

Regardless of whether we think of organizations as agents in their own right who ought to behave in some ways rather than others and can be held ethically responsible for the things they do, we certainly *want* them to *do* some things rather than other and to *be* some things rather than others. Moreover, organizations certainly do things, and occasionally change what they do and are over time, sometimes as a result of ethical criticism. So whether we should think of such ethical criticism as directed at the individuals within those institutions or at those institutions themselves, we can criticize organizations of ethical grounds. When should we? And why?

Some people claim that corporations have no ethical responsibilities beyond their duties to obey the law, or otherwise conform to local custom. (We already came across this view in Chapter 1.) In practice, however, most corporations do at least pay ‘lip service’ to ethical considerations, whether they do so in the name of ‘morality’, ‘social responsibility’, ‘responsible stewardship’, ‘corporate values’, or ‘corporate citizenship’. One cynical hypothesis is that this is mostly a case of ‘window dressing’, a reputation for corporate social responsibility having been identified as being ‘good for business’. Yet we should not allow ourselves to paint all corporations with the same brush from the armchair, especially if (as may seem to be the case) the commitment to social responsibility of a given company demonstrably hits ‘the bottom line’. Nor should we allow ourselves to simply assume without argument that companies engaging in practices claimed to be socially responsible are actually enhancing the ethical aims they claim to endorse.

In order to get clearer about these issues we need to acquire a better understanding of what the social responsibility of different corporations actually consists and to get better able to classify different kinds of corporate social responsibility where we find it. Only when we have made some progress in achieving those aims will we be in a good position to form a clear view about which forms, if any, of corporate social responsibility are ethically advisable or required in different circumstances, and be able to truly understand the ethical force of public concerns about the ethics of business and other professional organizations, whether in politics, the board room, or in our personal lives.

According to one common formulation, corporate social responsibility involves corporations actively balancing concerns of business profitability or success (such as the size of dividends) on the one hand, with wider social goods or harms (such as urban decay) and natural goods/harms (such as sustainable environment) on the other. Among different corporate practices involving considerations of social responsibility it also helps to distinguish between socially responsible practices that are pursued for their own sake (or *intrinsically*) and socially responsible practices that are pursued for the sake of some further end (or *instrumentally*). Those who claim that corporate social responsibility is merely a form of ‘window dressing’, or is pursued only as a profit-making strategy think that practices of corporate social responsibility always fall into the second, or instrumental, category. Frequently this is heard as a form of ethical *criticism*. We should be careful to remember, therefore, that even corporate responsibility pursued purely instrumentally is capable of preserving or promoting

genuine social and natural goods. (After all, that is the point of what is known as ‘sustainable development’.) Whether instrumental corporate social responsibility is always pursued openly as such is, of course, a very different matter.

Adapting a matrix created by C. Crook (no pun intended) for an article in *The Economist* in 2007 (just prior to the financial crash), it can be useful to distinguish four categories of CSR, depending on how the activities in question play out along two separate dimensions.¹⁴ In the first category, we can place CSR that both increases profit and success, *and* preserves or promotes wider goods. This is what is sometimes known as ‘win-win CSR’, or (in Crook’s words) ‘good management’. Win-win CSR includes socially responsible practices where there is no specific intention to preserve or promote wider goods (such as the design of energy efficient housing); where the socially responsible activity is purely instrumental to the generation of profit (such as ‘agro tourism’ or other forms of ‘cosmetics’); and where the aim of the practice in question is to preserve and promote wider goods for their own sake (such as landscape gardening). In each case, the activities in question accord both with ethical considerations and the profit motive. Genuinely win-win CSR is therefore not in any way ethically controversial.

In the second category, we can place CSR that both reduces profit or success and fails to preserve or promote wider goods. This is a kind of corporate practice that Crook labels ‘delusional CSR’, in that it is both bad for business *and* bad for society, even if it may be based on genuinely ethical or otherwise pure intentions. Examples of delusional CSR might include inefficient or polluting recycling schemes (including the transport of poisonous substances, such as batteries, to ‘far away’ places where the ‘recycling’ process results in the poisons in question being dumped unfiltered in the local environment), or environmentally targeted savings that increase company overheads without having any discernable result in practice (such as energy use or levels of waste). Thus defined, the ‘badness’ of delusional CSR is no more ethically controversial than win-win CSR is. (In fact, it is just the unethical flip-side of win-win CSR.) What is controversial, and sometimes hard to establish in practice, is whether a given case of CSR is really delusional in the way described.

A third category of ‘pernicious’ CSR involves apparently responsible practices that increase profit or success while actually harming or otherwise undermining the existence of wider social or natural goods. Pernicious CSR might well be ‘good for business, at least in the *short run*, but is bad for the social or natural world around it, and is likely to give rise to suspicion and controversy where companies are either ‘found out’, or are seen to be covering up the damage they cause to its victims or the natural environment. Among examples of pernicious CSR mentioned in the literature are market domination or monopolies by firms that claim to be socially responsible but whose market position is squeezing out even more socially responsible competitors and wage controls that dis-incentivize investment by potential market participants who are unable to pay their potential employees the ‘going rate’. Whatever we may think of the frequency these and similar examples in practice, pernicious CSR is clearly more than an abstract possibility, and one that undoubtedly

¹⁴ C. Crook, ‘The Good Company: a survey of corporate social responsibility’, *The Economist* 22 January 2005.

does have its defenders on the business world (e.g. among those who deny that the ethical responsibilities of corporations extends beyond the domain of the law and local custom). Yet even among the most stringent proponents of market freedoms and deregulation, the category of pernicious SCR ought to raise some ethical eyebrows, insofar as it is essentially *deceptive*, dishonest, and otherwise untruthful, which is very much in conflict with the idea that ‘free’ market activity is justified by the idea of free and voluntary consent, a condition that can hardly be said to be met in the case where the victims of pernicious CSR are basically being deceived.

The fourth category of SCR is one where the corporate practices in question do genuinely promote wider goods and prevent wider harms, but where by doing so they also reduce profits or make the corporation less successful in terms of its business, or corporate, mission. This category of CSR has been labelled ‘borrowed’ (or even ‘stolen’) virtue, and is arguably the kind of CSR around which much of the ethical controversy surrounding this idea has historically turned. Borrowed virtue, thus understood, may be good for society or the natural world, at least in the short run. But equally, by definition, it is also bad for business. Yet from giving charitable donations to public/private partnerships in the delivery of public services, many corporations do seem to operate in an environment in which their activities could both be more efficient and more profitable for their owners and shareholders if they did not have the wings of their capacity to generate corporate value constantly clipped by the perceived need to also produce social and other forms of value that extend beyond their corporate mission. Some people would therefore argue that insofar as any forms of CSR do fall onto the ‘borrowed virtue’ category there is something deeply wrong about them, and that the corporate aim should always be to turn any case of borrowed virtue *into* a case of win-win CSR.

There are at least four distinctive arguments for this conclusion, each of which can be traced to some highly influential contributions to the debate over corporate social responsibility, but each of which can also be shown to have clear limitations once we apply it to the corporate reality they are meant to describe.¹⁵

The first argument we can label ‘the Robin Hood argument’. According to this argument, all forms of non-profitable CSR is stealing other people’s money; primarily, although not exclusively, the money of shareholders who have invested in company stock in order to maximize their financial return. Yet stealing other people’s money is normally wrong. (We don’t need a separate chapter in this book to establish that conclusion.) So all forms of non-profitable CSR, including most specifically CSR of the ‘borrowed virtue’ variety is wrong.

There are at least two obvious limitations of this argument. The first is that some companies are specifically set up as ‘ethical’ businesses, from companies that advertise themselves as targeting the ‘fair trade’ market to multinational corporations with publicly advertised CSR programmes and highly articulated ‘business values’. No one is forced to invest in such companies, or to retain their investments in the knowledge that there are better returns to be made elsewhere. The second limitation of the argument is that owners and shareholders, as well as customers and the general

¹⁵ M. Friedman, ‘The Social Responsibility of Business is to Increase its Profits’, *New York Times Magazine*, 13 September 1970

public, have the power to influence the ethical profile of corporations, for example by approving or disapproving of the development or shape of different kinds of CSR programmes. The Robin Hood argument is therefore unconvincing in any economic environment that *allows* for investors to place their money in corporations the unambiguous aim of which is to maximize their profits (or *not* to do so, as the case may be).

The second argument against non-profitable CSR we can label the ‘Invisible Hand argument’. According to this argument, the more profits a business makes, the more value it produces. The more value a business produces, the more value there will be in the wider society in which the business operates. The best way to produce as much values as possible in the wider society as possible is therefore to make businesses be as profitable as possible. It follows, therefore, that the only ethically acceptable form of CSR is ‘win-win’ CSR, and that CSR of the ‘borrowed virtue’ variety is at best sub-optimal.

There are two problems with this argument. The first is an equivocation of the term ‘value’. In order for the argument to be sound we arguably have to understand ‘value’ in narrowly economic terms, such as in the claim that profitable businesses maximize the amount of *money* in the *economy*. Yet the kind of values that CSR is supposed to preserve or promote is frequently thought to go beyond what is measured by the amount of money in the economy (such as safe and stable communities), and is sometimes even thought to include values that conflict with increases in economic value, at least in the short or middle term (such as a clean environment or a stable eco-system). This is a version of what economists refer to as the problem of ‘externalities’. The second problem is that the fact that there is more of some good overall in the economy does not imply that this good is *justly distributed*. Thus, we might see an economy growing but only at the cost of severe inequalities or even the exclusion of entire demographics (such as non-favoured minority groups, immigrants, child workers, or slaves) from a minimally safe or decent quality of life. The Invisible Hand argument has little to offer in response to these problems apart from hopes, prayers and aspirational metaphors.

A third argument is aptly labelled ‘The Division of Labour argument’, and is in some sense a restricted version of the Invisible Hand argument and also a version of an argument we have seen at work in previous lectures in this course (e.g. Lectures 1, 2 and 3). According to the Division of Labour Argument, companies that engage in non-profitable CSR are effectively redistributing income by *taxing* their owners for good causes. Yet the systematic redistribution of income in the economy is really the job of the *state* (on top of which we can add the charitable donations provided by individuals from their own private, or personal, wealth.). Hence, by engaging in non-profitable CSR, businesses interfering with the role of the state, the constituent institutions of which are presumed to be more knowledgeable and effective in carrying out the redistributive task. Companies are therefore better off leaving the state to do its job and restricting its socially responsible activities to win-win CSR.

There are at least two problems with the Division of Labour argument. First, in the case of at least some very big corporations, their CSR programmes can be more effective than state programs (for example in building houses or transport links for employees and their families). And even if the corporations in question are neither

more, nor as, effective in their CSR programmes as the state, they are nevertheless be able to make some difference in terms of the wider goods they preserve or promote, just as the charitable donations of individuals can. In response, it might be argued that business managers are trained to run profitable businesses, not to run social programmes in the way public workers are. Yet this is clearly an entirely contingent historical matter, given that some business managers may previously have been public workers (and vice versa), and that there are plenty of examples of public institutions working alongside businesses to provide public goods, such as roads, railways, bridges, airports and hospitals. If skill-sets can be flexible within the corporate world (which they obviously are), then there is no reason to think that they cannot also be flexible across the private/public boundary.

A second problem with the Division of Labour argument is that in some contexts, there is no state (or no stable, or strong, state) to carry out the relevant redistribution that the argument assumes the state to be normally and legitimately engaged in. Indeed, many cases where the issue of corporate social responsibility is most acute involve the behaviour of large and powerful multinational corporations that operate in regions of the world (such as sub-Saharan Africa) where these organizations have more power than some of the states with whose legal blessing they are supposed to operate. In response, it is sometimes argued that leaving social programmes in the hands of private business will actually subvert the aim of creating the strong states we wish to support. Yet once again, this is a purely contingent historical matter, unlikely to apply in all places without restriction, e.g. in situations where privately funded CSR programmes (such as education, medical care or security) could themselves be one effective way to encourage or support the creation of a strong or otherwise stable state.

A fourth, and ethically more substantial, argument might be labelled ‘The Democratic Legitimacy argument’. According to this argument, non-profitable CSR is once more said to involve private businesses redistributing income by taxing individuals. The problem now, however, is not that the state is said to be a more *effective* vehicle for doing so, but rather that it is only in the hands of the state that such taxation is *democratically legitimate*. (The standard battle cry, which has a prestigious anti-colonialist pedigree is: ‘No taxation without representation!’) The democratic legitimacy of agents of the state is only guaranteed by the fact that its power to do so is vested in them by an electorate who vote for representatives whose policies include a view on redistributive taxation. Business executives and other employees of private companies have no comparable legitimacy in redistributing wealth across different populations in either their home territory, or, indeed, in any other legal territory. For businesses to engage on non-profitable CSR is therefore for those businesses to *usurp* the power and authority of the state.

There are two problems with this argument, one of which signals a limitation in scope and the other of which presents a fundamental obstacle. The first problem is that owners of private companies (such as interested shareholders) can give valid consent to being ‘taxed’ in this way, either *implicitly* by investing in firms that engage in non-profitable CSR, or *explicitly* by voicing their approval of these practices (e.g. by voting for governance proposals or the appointment of executive officers). In response to this objection, critics of ‘borrowed virtue’ may point out that the people affected by the ‘taxation’ involved in unprofitable CSR will normally include the

wider public beyond the ownership of the company, people who may benefit less from the company's business activity to the extent that the company is unable to afford to continue its operations in the manner to which they have become accustomed. (An airline cancelling a variety of routes to less busy regional destinations would be one example of these kinds of wider effects.)

The more serious problem with the Democratic Legitimacy argument is both more fundamental and much easier to grasp. This is that if a democratically accountable government has the democratic power and authority to regulate businesses at all (which it must have in order for there to be any kind of large-scale market in the first place), it has the democratic power and authority (as vested in it by the electorate) to regulate businesses in such a way as to not only *allow*, but actually *encourage* unprofitable CSR. (If government *required* companies to engage in non-profitable CSR, this might be considered a case of unjustified coercion, but 'permit', 'encourage', or 'incentivize' do not imply 'require'.) Indeed, it could be argued that insofar as the electorate has elected a government that either allows or encourages companies to engage in non-profitable CSR, that electorate has either tacitly or explicitly consented to any additional taxation that the resulting non-profitable CSR involves. Furthermore, not only are governments frequently elected that do allow private companies to engage in socially responsible practices that negatively affect their profitability, they also allow for the existence of organizations (such as charities and NGOs) that do not exist in order to maximize their profit at all. It would be a strange political culture that made the existence of not-for-profit organizations illegal, or that sought to delegitimize institutional practices of *aid, friendship, or the gift*. The crucial issue is not whether all organizations, whether business corporations, registered charities or public bodies should be forced to either maximize profit or not exist with the aim of making any profits at all. This is a false dilemma, based on a serious oversimplification. The issue is rather what the balance between profit and wider goods should be for any economic agent in a democratically legitimate social environment. There is no good reason to think that there can only be one reasonable answer to that question regardless of the circumstances or the ethical issues involved, even within what we know as 'the private economy'.

Crook's category of 'borrowed virtue' is the hardest case to handle for proponents of corporate social responsibility. We have seen that in spite of serious challenges, even this category of CSR is arguably permissible, or even required, in some circumstances, even on basically market-friendly terms. Yet the detailed controversy over the ethical status of 'borrowed virtue' should not detract our attention away from the fact that there are other forms of CSR, such as delusional and pernicious CSR, that are clearly to be avoided, and yet other forms of CSR, such as 'win-win' CSR (or just 'good management') that is clearly to be encouraged. In many, if not all, cases one might hope that companies would be able to carry out their social responsibilities within the boundaries of 'win-win' CSR. Yet even where they are not, there are no compelling arguments to show that their CSR activities are thereby illegitimate in themselves, even if it is true that in many cases what may pass for CSR is really something quite different in disguise.

Epilogue: No end in sight..

When I started teaching seminars on this material over ten years ago I approached the exercise in the spirit of an experiment. I had no designs for a fully worked out lecture course and no plans for future expansion. As the preceding lectures have gradually taken shape over the years I have regularly stopped to ask if the best way forward would not be to produce a thematically integrated and theoretically systematic presentation of ethics in business and professional life, explained and argued for on the basis of a small number of basic principles – in other words, a universal ‘system’ of ethics for business and professional life in general. Yet every time I have stopped to think about it I have concluded that this would not, after all, be the best way to proceed, at least for me. (Other teachers may reasonably have different ambitions.) Three considerations in particular have made me draw this conclusion.

The first is that presenting a ‘system’ of ethics in the relevant sense would require me to engage with a large number of abstract theoretical debates that, quite frankly, in my experience many of the students in my classes find an unwelcome distraction from what they are worried about, and so would be likely to make little use of in practice anyway. Instead, therefore, I have continued to adopt a ‘bottom-up’ approach, variously pursuing the theoretical presuppositions of different debates as far as necessarily to clarify a point at issue, but also feeling free cease pursuing theoretical sophistication at that point. If to a trained philosopher like myself this have sometimes have made me feel as if I hardly ever get started, it has also made me realize the strict limitations on the usefulness of what actually tends to happen when I do.

The second is that by focusing on the ‘issues’ rather than the ‘theory’, I have found that I am generally more successful in drawing in students from a variety of different educational, professional, theoretical and political backgrounds to focus on problems that in some shape or other they all tend to already recognize. For whatever they may end up thinking about the ethics of lying in business or the corporate social responsibility, for example, they all tend to agree that the *issues themselves* are of genuine interest and importance, even if they may vehemently disagree (or just not care) about the underlying *theory*.

The third is that by focusing on a number of separate issues and then drawing connections between them in the context of real-life cases (such as issues of trust, truthfulness and integrity displayed by our political representatives in recent elections) I have been more effectively able to honestly communicate my own conviction that one of the most useful things that philosophy can provide is a conceptual ‘toolkit’ to address the ethical challenges that arise in professional life, and our understanding of which will always to some extent remain partial and open ended. The competing view, that ethics is best thought of as a case of applying generally fixed and previously understood principles to ever more cases is one that I not only have very limited time for as a philosophical method in general, but one that I have also

seen reduce what could otherwise have been a creative and constructive discussion of a pressing ethical issue to barren, vacuous or dogmatic exchange of abstract and pre-programmed platitudes. To this extent, I am reluctant to concede that the palpable absence of system and rigour that characterizes these lectures is really a serious cost. I therefore much prefer to maintain their presentation as the partial, open ended, and 'living' document that they are in my classes. But then (to quote one of my own students), I would say that, wouldn't I?