The disunity of the moral

Adam Morton, University of Bristol

(This paper appeared in J.A.M. Bransen & M. Slors (eds.) *The Problematic Reality of Values* van Gorcum 1996, together with a sharp commentary by Menno Lievers.)

1. Naturalism and groundedness

I shall defend a kind of naturalism about morality. I have no ambitions of deriving a moral code from assumptions about the human species and its place in the natural world. In fact, I shall argue that natural facts about human action make the derivation of a single best moral code impossible. But I shall argue that the fact that we have moral codes and the force that these codes have in our lives, can be explained naturalistically. Moreover, by seeing why animals like us have the need we do for norms and ideals we can, I claim, go some way to seeing why we should accept their authority, why our lives would be poorer if we lacked this dimension.

Among philosophers the most controversial aspect of the view I defend is its denial of any deep unity to moral considerations. But this denial is consistent with an internalism about the good, that is, an insistence that values can give us reasons for wanting. In fact, it reinforces internalism, as I see it. And this conflicts with a widespread suspicion in the general non-academic culture around us, that a commitment to the authority of moral considerations is incompatible with a full naturalism about humanity. One of the few philosophers who has tried to articulate this suspicion is Charles Taylor. It is hard to be sure whether what I defending is exactly what Taylor attacks, because while he is usually very clear and expressive, if not exactly precise, he is not at his best about the word 'naturalism', for all its rhetorical importance for him. It seems to mean several different things for him. Footnote 1. For Taylor the root of morality consists in the creation of conceptions of our selves through which we can live and narrate our lives, and which tie the identity of each person to that of others. The creative and narrative element of this is seen by Taylor as entirely different from the process of describing facts about the world about us. Creating conceptions of ourselves is, for Taylor sometimes, creating selves, and these selves are not elements of the natural world.

Related to the issues about naturalism is a point about modernity. Taylor, together with Alasdair MacIntyre and some other modern writers, sees a crisis of morality in modern culture. Footnote 2. We lack the basis for a sense of heartfelt conviction in our values. Taylor puts the point thus:

> In fact, I want to consider a gamut of views a bit broader than what is normally described as 'moral' ...They concern.. what makes a life worth living. What they have in common with moral issues, and what deserves the vague term 'spiritual', is that they all involve what I have called elsewhere 'strong evaluation', that is, they involve discriminations of right and wrong, better or worse, higher or lower, which are not rendered valid by our own desires, inclinations, or choices but rather stand independent of these and offer standards by which they can be judged.
> Sources of the Self p4

This is one of the 'demands' on the concept of the moral, that Taylor thinks is hard to satisfy given the usual ideological slant of modern life. The connection with naturalism is that our desires, inclinations, or choices are natural facts about us, and the demand of morality is for something from which the natural may be evaluated.

This demand for meta-evaluation is, I am sure, an essential component of any adequate conception of what morality is all about. But I shall argue that it can be met naturalistically. Concepts of the self do enter into the story, but then making concepts is something humans naturally do. (Perhaps it is the thing they naturally do. There is no need for selves as non-natural created entities when we have natural human beings and the concepts they create.) There is a price to pay for this naturalism, as we shall see, and at the end of the paper we shall have to consider whether the price is too high. That price is what I call the disunity of the moral.

It is thus important not to phrase the demand for meta-evaluation so that it begs the question against naturalism. We need to be able to step beyond our values and ask whether they are the right ones. And we need some ways of

answering these questions once we have asked them. That does not mean that the answers cannot depend on our beliefs, our desires, and other contingent features of the human condition. Whether there are convincing answers in these terms is a question we should leave open until we see good arguments either way.

2. The moral landscape

Why do humans have moral codes. Why moreover do they take good care of their moral codes: evolve, maintain and improve them? One central reason has always been clear, though it became more visible with Hobbes and has become even more explicit with contemporary analyses of strategic choice. That reason is that we are social beings whose welfare depends on how we get on with one another. Humans from early on lived in groups, and needed to act as groups in order to flourish, even survive, and since they lived in groups needed to be able to outwit and not be outwitted by others. So the basic reason we need ideas about right and wrong - which may take the form of ideas about rights, obligations, and general imperatives, and may take many other forms too - is to get along with one another, to derive benefits and minimize harms. Let me be more explicit about this. It is in people's interests that they have as many as possible of a whole series of things, the first three members of which might be called cooperation, exploitation, and immunity.

Cooperation: there are many things people cannot do alone and many situations in which the outcome is better for everyone if people act in concerted ways. (And there are many different forms of concerted action, some of which require a cooperation that goes far beyond simple coordination.)

Exploitation: given that others have a motive to act cooperatively, it will often be in an individual's interest to present the appearance rather than the substance of cooperation.

Immunity: given that others have a motive to act exploitatively, it will often be in an individual's interest to be able to discriminate the appearance from the substance of cooperation.

The series can be continued indefinitely. (The fourth element would be: given that others have a motive to discriminate true from fake cooperation, it will often be in an individual's interest to outwit that discrimination.)

The aim is to get as much as possible of all these benefits. But that is a tremendously hard aim. In fact, I think it is simply beyond human capacities. Given that we are finite creatures, with limited powers to store and manipulate information, the task of thinking out optimum compromises to all these competing aims is just too hard for us. Footnote 3. (Are they competing aims? Yes: that is one thing the prisoner's dilemma shows.) There are deep facts about the human condition here, the way we are essentially intelligent but not quite intelligent enough. Stupider social animals would get by on built-in social routines; smarter social animals would be able to think everything out from first principles. We are fated to be in the middle, having to think but not able to think everything through. Footnote 4.

Being thus in the middle, inherently social but equipped neither with fixed social routines nor transcendent intellect, we are doomed but muddle through conceptually. Human beings in human societies need shared beliefs about psychology and morals. They need those twin creations, folk psychology and moral norms. (I believe they really are twins, symbiotically dependent on each other, but that topic is too large for this paper. Footnote 5.) Needing them, we create them. We form beliefs about actions and motives, in particular about the patterns of action and the types of motives which will maximize mutual benefit and minimize exploitation. These form the core of moral codes. They enable people to get the benefits of cooperation without falling prey to the dangers of exploitation. (Most people's desires are deeply social, defined with respect to the well being and preferences of others. That is quite consistent with the conception of morality defined here. Psychological atomism is neither required nor denied.)

Moral codes will vary, reflecting differing material conditions, different psychological assumptions, and different conceptions of the individual and communal good. But they will have the potentiality of being true. That is, they describe patterns of interaction which are in fact of mutual benefit, so when they say that it is in people's interest to act in these ways very often it is in people's interest to act in these ways, these are desirable patterns of action. This point needs to be put with some care: it is in each person's interest that, given that many people adhere to the prevailing code, most people do. It may on occasion also be in some person's interest that given that most people adhere she not

adhere. And it may often be in many people's interest that most people adhere to some variant code. The consequence of all of these is that each person in a situation in which many others adhere to a moral code usually has reason to want that more of them adhere to it. Sometimes, not always, this will give the person reason to want that she herself adhere to it. And this is an objective fact about human animals in nature. So we see a hint, at any rate, of how thinking naturalistically can explain why people are nearly always subject to moral norms, how their motivation is deeply shaped by them, and how they have good reason to take these norms seriously.

Under suitable conditions such a moral norm will not only give definite reasons for individuals to adhere to it and societies to support it, but also give meta-evaluations, reasons for thinking that it is 'right' from a more impersonal standpoint. (It will give "discriminations of right and wrong, better or worse, higher or lower, which are not rendered valid by our own desires, inclinations, or choices but rather stand independent of these and offer standards by which they can be judged.") The most suitable such conditions are rather idealised, but I think the best way to see the issues clearly is to state them and then to see how they are an ideal.

Imagine a society with a set of shared beliefs and patterns of action, and a number of norms which they expect each other generally to follow. The norms constitute a unique and obvious pattern of co-operative action which it is in everyone's interest to see generally followed. For if there were such a unique pattern then acting in accordance with it would be to act in a way that benefits everyone else while at the same time supporting something that is in ones own interests. It would be the sort of thing that people might agree to as a contract designed to obtain mutual benefit. But of course there is no need to think of this as anything like an explicit contract. I shall refrain from calling it a contract. I shall call it an 'understanding' instead.

There are many ambiguities here. In fact, they are in a way the intellectual substance of this paper. So let me state more carefully a set of conditions under which such an implicit morality- defining understanding can exist. To state them we have to consider a number of people, think of them as all the people in a society, who have open to them a number of options. Each option specifies basic facts about the person's attitude to the norms in question. For example option 2 might be, in part 'keep all your promises and be kind to everyone'; option 365 might be 'be nice to other people in your clan and go out of your way to do harm to everyone else'; option 35 might be 'always try to get as much money for yourself as you can'. The important things are not so much these options as combinations of them. One combination might be: person One chooses option 365, and person Two chooses option 35, ..., and person Three billion and six chooses option 2. The combinations are really just squares in a gigantic pay-off matrix, and so for each such combination there is a pay-off for each person. And - to specify the relevant conditions - there is one very special combination of options which has the following three features. Footnote 6.

(a) *Rough equilibrium* If many of the actual people deviate from it then everyone suffers. That is, if many people fail to act in accordance with it everyone's wants are satisfied less well than they would have if most people had acted in accordance with it. (Note the many, most, and the actual, wants.)

(b) *Perfection* In a situation in which everyone acted in accordance with it everyone would be maximally well off; that is, more of the desires they would have in that situation would be satisfied than would be satisfied by conformity to any other norm.

(c) *Intelligibility* There is a description of it which everyone can understand. Footnote 7.

Call any such pattern of action an *ideal moral understanding*. ((b) makes it ideal, (a) makes it moral, and (c) makes it an understanding.) Note that (a) gives the ideal understanding a kind of self-enforcing quality. If numbers of people deviate from it then they will suffer by doing so. Nevertheless, there can be a free riding (and thus Prisoner's dilemma) -like quality too, in that individuals or even small numbers of people breaking the contract may gain by doing so. Thus people may have a motive for setting up a system of incentives and punishments to keep deviation at a minimum.

There are a number of reasons why this ideal situation is not, and probably never can be, an actual one. But before considering them see one very important point. That is, the way such an ideal moral understanding can support meta-evaluations. The link with meta-evaluation comes from clause (b), perfection. That clause entails that in the ideal situation in which everyone acts in complete accordance with the norms, everyone would be maximally well off. This could be believed even if it were also believed that such ideal total compliance were impossible for the actual people

as they are: it is not a claim about actual people and their actual desires but a claim about ideally compliant people and what their desires would be. It is easy to think that ones own society and its norms satisfy (a) and (c) but not (b): that they are in an intelligible rough equilibrium but not perfect. All one has to do is to imagine a situation in which people live according to some other norm and satisfy more of their desires than you satisfy yours. (There are big and interesting problems here about comparing different degrees of desire satisfaction. I'm sure one does not just count desires and degrees of preference and add them up. This topic needs a full discussion that I shall not give it here. In fact I think it is the central underlying issue here. Let me just say that you should not be taken in by the prejudice of our culture that there is no way of comparing different people's different preferences.)

This is, of course, something that people do frequently. They imagine a different social and moral reality and say "they would be happier/better off/living better lives" than us. When people actually do this they almost never imagine completely different people with completely different values and norms. They usually imagine ways in which the desires and values of their own society would be better realised by a rather different set of norms and institutions. And the interesting fact is that there has never been a human society that was not vulnerable to such an "internal" criticism. No society satisfies its own values: rough equilibrium always falls a long way short of perfection. And as a result there is always room for meta-evaluation: comparing the actual norms to more or less distant possible or imaginary situations, to see which actual and possible desiderata might be satisfied in them. And so we see how to meet Taylor's requirement that morality should not be "rendered valid" simply by "our own desires, inclinations, or choices". Meta-valuation builds on possible preferences, and works by comparing them to the consequences of actual ones. As long as what is natural includes what is (actually, really) possible the natural can be judged naturalistically without begging the question. Footnote 8.

This contrast between rough equilibrium and perfection is crucial to my argument. A complex of moral ideas at work in a society defines what I shall call a *moral landscape*. That is, it defines a conception of the adequacy of the degree of compliance found in the actual situation and of how things might be if only people were more capable of good behavior. On a traditional moral landscape the actual situation is on a ridge leading from an abhorrent state of nature downhill to an unattainable peak of perfection uphill. The slope dips very steeply down from actuality towards chaos, so the actual situation with its imperfect compliance to norms does represent a kind of stability, but the road up to perfection while not so steep is very long. There are many variations on this image. Different ethos, different landscape.

3. Self-respect

If individual people are to act in harmony with others they have to know what actions to expect of others and what actions others will expect of them. They need this information in fairly simple and manageable form: they need to know the general shape of the landscape. One way in which they can know this is in terms of the combination of the idea of a value and that of a second order desire. Let me explain each in turn.

To value some situation or some general quality is more than to want it. One can want coffee, or like folk music, without attaching any value to it. For one can want things without thinking that they are valuable. To value something, as I shall understand the concept, is to want it to be wanted by you and others, and to want that it be wanted even if you or others ceased to want it. So if you want coffee but do not value it and anticipate your love of coffee declining the prospect will not bother you. And it will not bother you if some other person does not want coffee. But if you were to value coffee, the way some people value justice, love, or beauty, you would want to resist any danger that your want for it would decline, and you would try to criticise or convert any misguided person who did not like it.

Second order desires are desires about ones desires. One has a second order desire when one wants to want something. The standard example is of an addict who wants something (coffee, for example) which she would rather not want. It has been persuasively argued by Frankfurt and others that the possession of second order desires is a necessary condition of being a person, an agent capable of taking part morality: being responsible for her actions, being a party to a social contract, and the rest. Footnote 9. There are certainly practical advantages to having second order desires. A person who has second order desires which have some influence over her first order desires can make plans which turn on her possession of the right desire at the right moment; she can reflect on the desires which obstruct or compromise

the satisfaction of her other desires, and change or contain them.

It is not that easy for relatively primitive mammals like ourselves to have second order desires. In order to get the advantages of second order desires one has to have the grammatical sophistication to manage embeddings of one proposition in another, and one has to be able to connect these complex propositions with the plans of action that one carries out. And, very fundamentally, one has to have a conception of ones self, an I-concept, that is at once a central element of ones thoughts and desires for ones own life and as simple demonstrative me. Footnote 10. The content of ones desire is for example "that I not want so much coffee." And the I concept is the focus of many primitive and complex emotions, which do not respond readily to rational influences. One solution to the problem of having and managing second order desires - in fact the normal solution in human life - is to tie them to values. One absorbs from ones culture norms of character, motivation, and behavior, which are valued - most people want most people to conform to them - and one sets conformity to them as the easily grasped image of the way one wants to be. In terms of a set of values one defines ones identity, ones grasp of oneself and of what one is worth.

As a result, the values one holds constitute not only ones navigational equipment in ones moral landscape, but also ones grasp of the emotions such as respect, self-respect, contempt, admiration, guilt, and shame. Footnote 11. (Moreover, though this is not a theme of this paper, ones values are normally at the heart of ones attempts to understand oneself; without some sense of what goods ones efforts are tracking one is left with a sense of the inexplicability as well as the futility of ones actions.)

This is the emotional grip of metaevaluation. The second order thought that ones desires are the right ones is what allows one to consider oneself loveable. A deeper form of self respect, a real sense of self worth, is the sense that one is a suitable object of affection, a deeper form of moral condemnation is to think of someone as utterly unlovable. These connect with the primitive landscape described above. An admirable person is someone who is the way "we" all would be in Perfection, an adequate person is one who functions in the way needed to keep the rough equilibrium going. If a person assumes that everyone else values perfection and the maintaining of the rough equilibrium, and uses those values to shape her own second order desires, then she will want to be admirable, and the thought that she is adequate is necessary for her basic emotional functioning.

These are fairly limited claims. I am arguing that given normal human psychology under fairly ordinary social conditions there will be a link between the values that are widespread in the society, serving to maintain generally beneficial equilibrium, and the ways in which individual people maintain their self-respect. I mean this to entail a platitude, that for most people many evil actions are impossible: we could not live with ourselves if we thought of ourselves as murderers, cheaters, or people whose life-strategy was deception. It is an important platitude, though. It reminds us that meta-evaluation, the conviction that there is an answer to the question "why value this?", can be provided for an individual by an awareness of the way they want people to be, and for a whole society by the knowledge that the rough equilibrium they actually occupy gives some ideals a special role. The ideals sustain the equilibrium and also give a perspective from which to judge its imperfections.

But the claims are definitely limited. They do not show an inevitable psychological price for violating a moral norm. Monsters, psychopaths, and everyday banal wrongdoing are still perfectly possible. But if the picture I am presenting is right some categories of immorality should be more frequent than others. For example we should expect petty norm violation that does not threaten the rough equilibrium to be an everyday fact, and we should expect people to deceive themselves about the extent to which they violate the values on which they base their self-respect. Similarly, we should expect major violations to be committed either by pathetic loners whose lives are too simple to need a complex concept of self and self-worth, or by powerful megalomaniacs with the capacity to manufacture for themselves a role excusing them from the usual values and responsibilities. But most of all we should expect moral inadequacy to stem from the complacency of a rough equilibrium that pays only intermittent and hypocritical attention to the ideals that are needed to sustain it.

4. The real landscape

The ideal situation of rough equilibrium, perfection, and intelligibility was very ideal, in being an idealisation, based on not very realistic assumptions. Here are four fundamental reasons why the assumptions are not realistic, why the

landscape is not as it would be if a simple picture of morality were right..

Plurality of rough equilibria Condition (a), rough equilibrium, supposes that there is only one set of norms rough compliance with which will be generally good for everyone, given their actual preferences. But that is not true generally. Even given actual preferences, there are usually many other norms such that if people were to act in accordance with them they would be generally well off. Typically some would be better off and some worse. And difficulties in comparing different people's preferences and the pervasive incomparability among even one person's preferences make it hard to compare such alternative rough equilibria to the actual one. What is much more nearly true is that there is only one rough equilibrium which people in a given society will expect one another to adhere to. Given a folk psychology and a shared ethos, people will consider some of one another's actions much more probable than others. And if we only consider rough equilibria between probable actions then a single equilibrium is much more likely. But note that it is in part the shared adherence to a set of norms that is giving these norms their special properties. They are in a way self-justifying.

(Related to this plurality of rough equilibria is a scale-related indeterminacy within many rough equilibria. Construing the rules in one way may be more in the interest of those entering into interactions involving large numbers of people, while construing them in another way may be more in the interest of those entering into small scale interactions. Macroethics may differ from microethics. I will not say more about this.)

Plurality of better equilibria The perfection clause (b), supposes that there is a single best pattern of behavior which would result from everyone acting in accordance with the norms. But this is definitely false for any set of norms humans have yet discovered. Comparing the ideal situation of perfect compliance with actual norms with alternatives in which hypothetical people follow alternative hypothetical norms, it is never plausible that perfect compliance with actual norms wins. There are usually better alternatives, better even in that many actual desires and actual values would be better satisfied. And there are usually many of these possible ideals, each hard to compare to one another. In effect, the point is that the search for reflective equilibrium bites deeper than the naive perspective realises. It not shows much more than the imperfection of the actual situation: it suggests that greater perfection will involve not greater compliance with actual norms but compliance with rather different norms.

Non-existence of perfection Not only is there in general a plurality of situations which are improvements on general compliance with actual norms, but there is no reason to believe that there is in general a best one. For as we move from one set of norms to another our perspective changes, we see new and attractive preferences, values, and norms which we could not see before. Society shapes desire, as Rousseau taught us, and society shapes imagination, as Hegel taught us. our societies shape our capacity to imagine the options that are open to us. In a hunter-gatherer society you cannot imagine capitalism. So instead of a perhaps unattainable single peak of perfection we have a series of peaks scattered about the landscape, whose heights and climabilities are hard to estimate until one is on the slope leading to them.

Downhill detours This is the most subversive point of all. There are improvements on actuality the route to which involves a detour through definite un-improvements. That is, there are patterns of action such that if large numbers of people acted in accordance with them the result would be in almost everyone's interest, but such that if only a few people did, the result would be catastrophic. Consider a well functioning feudal society. Suppose that a few people began to act as they would in an advanced capitalist society. The result would be disastrous for those few people. Suppose that a few more began so to act. The result would be disaster for everyone. Suppose that almost everyone began so to act. The result would be a general improvement. Almost everyone would gain. No one wants to go back to feudalism. Is not the same surely true of our way of life? There are better ways to live, but were only a few to adopt them the consequences would be bad. So what is from one point of view visionary is from another point of view perverse.

To sum this up, the moral landscape as imagined by a traditional morality consists of a stable point on the shoulder of a slope up from chaos, with a single peak of perfection rising behind. But the truth is more complex. The low-lying swampy chaos is there, and there are shoulders and peaks. But there are many stable points safely above the swamp, and there are many peaks, whose heights are hard to compare from a distance. Moreover sometimes the route from one stable shoulder to a higher one involves a delicate route which at times slopes downwards and passes by precipices

above the swamp.

That's the way it really is. And we all at some level know it. Traditional cultures have many devices for keeping us from a too panoramic view. Religion is one such. The role of the gods, I think, is not as often supposed to provide a system of incentives and punishments to hold people to the rough equilibrium. Simple self-interest will do that, and if it will not the norms will not survive. Rather, the divine element is there first to reinforce a conception of the normal, of what patterns of behavior are overwhelmingly probable. Given that, the uniqueness of the rough equilibrium is greatly reinforced. And, more profoundly, it is to focus people's attentions on one among the many perfections looming in the distance. If you can sustain the myth of unique perfection you can imagine that there is only one set of norms that can be found o the road to it, and so you will fail to see any probability that a reflective and morally sensitive person will act in a way that does not sustain the norm. If you look only in the direction of Olympus you will not see disturbingly many other heights and passes.

5. Real values, problematic values

What remains once we realize the true landscape? All the components of the naive picture of morality remain, if we are careful how to define them, but they are related to one another in importantly different ways. Rough equilibria remain. We have to define them in terms of what is mutually advantageous among the patterns of action that are probable, given the people in question, with their awareness of a set of norms and their subscription to a folk psychology. But given all that, it is a pretty ubiquitous fact that people usually act in the context of norms whose general observance is in the general interest. Ideals remain, if not perfections. The rough equilibria are only locally optimal, given either a little imagination or serious moral questioning one can find many plausible candidates for ways we would all be better off. And in this way meta-evaluations are still possible: nothing in the true landscape undermines our capacities for standing aside from our current values and seeing how they approximate to and fall short of the possibilities for human life. And moral emotions remain. One can still be satisfied with the way one sustains others in the common attempt to cobble together a coordination of our aims. One can still condemn others for messing it up. One can still admire the moral courage of those who insist on something better.

These are real values. They are quantities that people in a society have good reason to want to be wanted. Moreover, they come with metaevaluation attached. Several kinds of it, in fact. For given a rough equilibrium there is a variety of ideals, a variety of better situations which different variations on the norms behind the equilibrium could produce, at a variety of "distances". And these ideals are objectively better, in that any bicultural person who had been socialised into both such an ideal and the rough equilibrium would choose the ideal. (There usually are no such bicultural people, and it is usually very hard to know what they would choose if they did exist. That is a problem about what we can know, not about what is real.) From the point of view of such an ideal the effect of the actual values can be evaluated. They will usually emerge as falling far short of the ideal. Contrast this with an attempt to validate a society's values by reference to a symbiotic ideal whose function is to preserve those very values. The values will usually be given high marks by such a validation, and a supernatural gloss on the process will underline the importance of adhering to them. But the whole process is not very comforting. The accounts have been cooked; the result is pre- ordained. So, paradoxically, in this time of intellectual whining about the fragility of our values, we have better ways of describing the objectivity of values, and better and higher standards of metaevaluation.

Paradoxical as it sounds, I think this conclusion is solid. We have better grounds now for believing that human values can be objectively judged and grounded than people have ever had. That is the factual and logical situation. But the emotional and practical situation is different. For the easy rationale for acting in accordance with our present imperfect values is undermined. To act in accordance with the norms of ones culture is no longer automatically to be on the road to something better. Sometimes the way to something better involves coordinated violations of the present norms. The coordination is a difficult numbers problem, as I described above, and usually involves a detour through something worse. Moreover the plurality of ideals makes the coordination even more difficult. One not only has to recruit enough people to make deviation reasonable, and one not only has to face the possibility that the ideal may be illusory, but one has to be sure that it is the same ideal that enough people are aiming at. The ideal is no longer defined as a perfect

version of the actual norm. Moral insight requires essentially political organization. Footnote 12.

Morally relevant emotions have to change too. Our simple conceptions of self-respect and moral admiration are tied to a simple relation between decent behavior in the immediate situation and attainment of an ideal. They cannot survive a more realistic realization that better situations may come about from people violating some of the norms of present good behavior. (The right number of people, in the right circumstances.) This is not to say that a more appropriate folk psychology, imbedded in a suitable way of life, cannot give us a way to emotions that are at the same time profound and motivational and based on the true landscape. But nothing remotely appropriate is in sight.

The most general conclusion I would draw is that moral thinking is really three rather different activitiesFootnote 13. . In the first place there is the enterprise of getting clear about the values and norms that we actually hold, seeing their structure and how they function in our lives, given our actual preferences. This is an enterprise requiring care, precision, and patience. It is a good job for analytical philosophers. Then in the second place there is the enterprise of seeing *better* values, norms, and preferences. This requires imagination, insight into human nature, and an ability to communicate a vision. This too is a task in which philosophers can play a part, though it requires a very different intellectual ethos. And then in the third place there is the enterprise of coordinating our attempts to move towards different ways in which we could live. This is politics, including intellectual politics, attempting to affect the ideologies and ethos of people around one. And although all of these are related, and each must be carried out with an eye on each of the others, they may well not be neighbouring activities. For each of them the considerations that most matter may come from disciplines and projects quite remote from the other two. The first may be closer to epistemology, the second to psychology, and the third to economics, than to either of the others.

Which of these kinds of thinking is most likely to reinforce our sense that our values are real? The first, the analytical task, can assure us that many of our values serve functions in allowing us to manage our joint activities in a reasonably successful way. That is a kind of reassurance, but it is undermined by the realization that these values are just a selection from the plurality of possible values which could serve this function more or less well. The second, the imaginative task, is more promising. It promises an insight into the range of ways in which people could live and the range of values they could have. (The really possible range for real human people.) It is an insight that could be unsettling. Parochial, complacent, or opportunity-missing systems of values will not do well in comparison with real human possibilities. And there surely are many ways in which our values and norms are parochial, complacent, and opportunity-missing. But if you want a real grounding of our values you must face that danger. Footnote 14.

Footnotes

1. See Charles Taylor *Sources of the self* (Cambridge. U.P. 1989.) When he talks of naturalism at different points Taylor has in mind (a) Scientism: the dogma that the only things we should take seriously are the results of The Scientific Method, whatever that is, especially when they take the form of the Hard Physical Sciences (b) Reductionism: the dogma that everything that exists can only be understood in terms of its physical components (c) Naturalism proper: the conviction that human beings are part of the natural world and can be discussed in the same terms as anything else in it. All three of these are independent of one another. (a) and (b) are false. (c) is true, I think. Does anyone deny (c)? Does Taylor deny (c)? Sometimes his rhetoric of the creation of selves seems to.

2. Alasdair MacIntyre *After Virtue* second edition, (Duckworth 1984), and *Three rival versions of moral enquiry* (Duckworth 1990.), Richard Rorty *Contingency, Irony, and Solidarity* (Cambridge University Press 1989).

3. See Binmore, Ken and Shin, Hyon Song 'Algorithmic knowledge and game theory', Wlodzimierz Rabinowicz, 'Tortuous labyrinth: noncooperative normal-form games between hyperrational players', Brian Skyrms 'Equilibrium and the dynamics of rational deliberation', all in Biccieri and Dalla Chiara, (eds.), *Knowledge, belief, and strategic interaction* (Cambridge University Press 1992). Also Hyon Song Shin and Timothy Williamson 'Representing the knowledge of Turing Machines' *Theory and Decision* 37, no 1, 1994.

4. This is a theme of modern studies of rationality, for example in Herbert Simon's *Models of bounded rationality* (vol

1, M.I.T. Press 1982), but the idea goes back at least to Aquinas and surfaces also for example in Heidegger.

5. See my 'Folk Psychology is not a predictive device', *Mind* 105, Jan 1996, 1-19. That paper is part of work towards a book whose working title is *Folk psychology is ethics*, as is this one.

6 For an exposition of the basic game theoretic concepts used in this section see Sean Hargreaves Heap, M. Hollis, B. Lyons, R. Sugden, R. and A. Weale *The theory of choice.* (Blackwell 1992).

7. See David Lewis *Convention* (Harvard U.P. 1968), and Robert Sugden *The economics of rights, co-operation, and welfare* (Blackwell 1986). The intelligibility condition is influenced by Christopher Bertram's 'Political justification, theoretical complexity, and democratic community', *Ethics* 1997.

8. There are obvious connections with Rawls' idea of a reflective equilibrium here. See John Rawls *A theory of justice*, (Harvard University Press 1971), ch. 1, and Norman Danielson 'Wide reflective equilibrium and theory acceptance in ethics' *Journal of Philosophy* 76, 1976, 256-82. There are equally basic links with recent work by philosophers and economists on comparing the preferences and needs of different people in different situations. See the essays in J Elster and J Roemer, eds. *Interpersonal Comparisons of Well-Being* (Cambridge University Press 1991), and in A Sen and M Nussbaum, eds. *The Quality of Life* (Oxford University Press 1993).

9. See Harry Frankfurt 'Freedom of the will and the concept of a person' *Journal of Philosophy* 68,1971, Michael Tooley *Abortion and Infanticide* (Oxford University Press 1983.)

10. See my 'Why there is no concept of a person' in C. Gill, ed. *The Person and the Human Mind: Issues in Ancient and Modern Philosophy* (Oxford U.P. 1989.)

11. See Patricia Greenspan *Emotions and Reasons*, (Routledge, 1989), and Gabrielle Taylor *Pride, shame, and guilt* (Oxford U.P. 1985).

12. Compare Richard Rorty " Modern, literate, secular societies depend on the existence of reasonably concrete, optimistic, and plausibly *political* scenarios, as opposed to scenarios about redemption beyond the grave. To retain social hope, members of such a society need to be able to tell themselves a story about how things might get better, and to see no insuperable obstacles to this story's coming true." (*Contingency, Irony, and Solidarity*, p 86.) But Rorty's emphasis on single ideals ("a story".."*this* story") and his disparagement of ideals that can not be readily obtained lead to very different conclusions from mine.

13. A similar conclusion is found in Amelie Oksenberg Rorty 'The many faces of morality', *Midwest Studies in Philosophy* to appear.

14. Very early versions of this paper were read at Armidale, Queensland, and Maryland. A middle version read at the Utrecht value realism conference prompted very helpful comments from Menno Lievers and members of the audience: I have learned from some of these and no doubt missed the point of others. In preparing the final version I have been influenced by conversations with Georges Rey