

Australian Journal of Professional and Applied Ethics

Volume 7, Number 1, June 2005

Contents

Editorial Board	ii
Papers	
• <i>The Tyranny of Romantic Love: Continuing the Conversation with Cephalus</i> Justin Jeffrey	1
• <i>Being Queasy about Reconstructing Animals</i> Michael Hauskeller	11
Book Symposium	
• <i>Introduction to Book Symposium on Bernard Gert's Common Morality: Deciding What to Do</i> Dean Cocking	22
• <i>Common Morality and 'Institutionalising' Ethics</i> Andrew Alexandra and Seumas Miller	24
• <i>Reply to Andrew Alexandra and Seumas Miller: 'Common Morality and 'Institutionalising' Ethics'</i> Bernard Gert	34
• <i>Applying our Common Morality: the Case of Privacy</i> Jeroen van der Hoven	38
• <i>Reply to Jeroen van der Hoven: 'Applying our Common Morality: the Case of Privacy'</i> Bernard Gert	44
• <i>Moral Arrogance and Moral Disagreement</i> Dean Cocking	48
• <i>Reply to Dean Cocking: 'Moral Arrogance and Moral Disagreement'</i> Bernard Gert	55
• <i>Timeless Wisdom or Moral Arrogance?</i> Thomas Pogge	60
• <i>Reply to Thomas Pogge: 'Timeless Wisdom or Moral Arrogance?'</i> Bernard Gert	65
Discussion	
• <i>Government and Faith-based Organisations in a Pluralist Society</i> Bruce Langtry	72
Critical Notice	
• <i>Neil Levy's What Makes Us Moral: Crossing the Boundaries of Biology</i> Richard Joyce	78
Notes on Contributors	86

Editorial Board

Editor

Andrew Alexandra *Charles Sturt University, University of Melbourne*

Associate Editors

C.A.J. Coady *University of Melbourne*

Stephen Cohen *University of New South Wales*

Seumas Miller *Charles Sturt University*

John Weckert *Charles Sturt University*

On-line Editor: Michael Collingridge *Charles Sturt University*

Editorial Board

John Alderson

Warwick Blood *University of Canberra*

John Boatright *Loyola University of Chicago*

Wendy Bowles *Charles Sturt University*

Terrell Ward Bynum *Southern Connecticut State University*

Margaret Coady *University of Melbourne*

Fred D'Agostino *University of New England*

Michael Davis *Illinois Institute of Technology*

James Griffin *University of Oxford*

Ravindra Jain *Jawaharlal Nehru University*

John Kleinig *John Jay College, Charles Sturt University*

Colonel Charles R. Myers *US Air Force Academy*

Denise Myerson *University of Cape Town*

Noel Preston *Queensland University of Technology*

Tom Regan *University of North Carolina*

Sam Ricketson *Monash University*

Alper Riza *Queen's Counsel, London*

Simon Rogerson *De Montfort University*

Charles Sampford *Griffith University*

Henry Shue *Cornell University*

Zhang Xiaoming *Chinese Academy of the Social Sciences*

The Journal

The Australian Journal of Professional and Applied Ethics is the official journal of the Australian Association for Professional and Applied Ethics.

It is published twice a year by the Centre for Applied Philosophy and Public Ethics, Charles Sturt University. Copyright for all material published in the journal is retained by Charles Sturt University.

The Tyranny of Romantic Love: Continuing the Conversation with Cephalus

Justin Jeffrey

University of North Carolina

'I cannot live without my life! I cannot live without my soul!'

He dashed his head against the knotted trunk; and, lifting up his eyes, howled, not like a man, but like a savage beast getting goaded to death with knives and spears.

(Emily Bronte, *Wuthering Heights*)

O that you and I escape from the rest, and go utterly off, free and lawless,

Two hawks in the air, two fishes swimming in the sea, not more lawless than we...

(Walt Whitman, 'From Pent-up Aching Rivers')

1. Introduction

At the beginning of book I of Plato's *Republic*, Socrates asks Cephalus, the elderly father of Polemarchus, whether he finds it difficult to have reached 'the threshold of old age'. Cephalus notes that many of his friends grumble about the lost pleasures of youth, including the pleasures of sex, but he finds that old age has brought him a welcome freedom and peacefulness that the young typically lack. Cephalus relates that he was once present when someone asked the poet Sophocles: 'How are you as far as sex goes, Sophocles? Can you still make love with a woman?' The poet supposedly replied, 'Quiet, man. I am very glad to have escaped from all that, like a slave who has escaped from a savage and tyrannical master' (Plato, 1992 (329c)).

The attitude Cephalus (and supposedly Sophocles) takes toward sexual, romantic, and other strong appetites is an early peek at the view of justice and the good life that Plato will develop and defend in later books of the *Republic*. Familiarly, on this view the appetites are supposed to harmonise with reason (and the spirit) such that each serves in its proper role; this means that the appetites ought to serve in a subordinate position with respect to reason and spirit. Plato, of course, is suspicious *generally* of appetites and their objects, and argues that the pleasures derived from the satisfaction of appetites are not only fleeting, but also not "real". We ought instead, he argues, to pursue the real and enduring pleasure derived from appreciation of the Forms.

For those of us who can't muster a belief in the Forms, and who don't find the broader subordination of the appetites to reason an appealing vision of the good life, there is still reason to be suspicious of *certain* appetites, and to wish, if not for their extinction, then at least their for their moderation or subordination to reason. For certain appetites - the kind Socrates would categorise as 'tyrannical' - are uncontroversially destructive to human lives and at odds with human flourishing. The desire for certain highly addictive drugs is one such example; another might be the unregulated desire to have power over others. Continuing from where Cephalus left off, in this discussion I will examine not specifically *sexual* appetites, but the related appetites surrounding romantic love as they are experienced by certain psychologically 'normal' human beings who nevertheless, in the common way of speaking, go a bit 'crazy' when they are in love.

Heathcliff from Emily Bronte's *Wuthering Heights* is an extreme and vivid example of a person in the sort of state I have in mind, though one need not bash one's head against a

tree trunk to qualify. The characterisation of the state might be less vivid than one might like, but I am hoping that most will recognize and recall in my description a state that they have been in themselves at one time or another.¹ I am thinking of a state in which the lover claims to need to be with their love-object, where the lover is preoccupied by thoughts of their love-object, where the lover feels willing to do just about anything for their love object. The lover often feels unhappy or unsatisfied unless in the presence of her love object, and is distracted and unable to focus in the usual way on tasks and projects unrelated to the love-object. Many people are in such a state during, for instance, the early stages of some of their more passionate romantic relationships. Many people - perhaps even most - also find themselves in this state when a love-object is threatening to leave them, or where one feels seriously threatened by a perceived romantic rival. For simplicity's sake, and in honor of Emily Bronte's afflicted character, I will refer to people in such a state as HCs, and to their romantic experience as HClove. Despite what seems to me to be the glamorous reputation of HClove, I will argue - appealing to a wide variety of evidence - that the psyches or souls of HCs approximate the 'tyrannical' in Plato's sense. Consequently, I will argue, HClove is a bad thing.

Some of the evidence I will appeal to in making my argument will be from our common, informal experience; other evidence I present comes from psychology and neurobiology. The empirical evidence offered from these disciplines is utterly relevant to the philosophical question of the goodness of HClove. It is evidence, I think, to which Plato himself would be delighted to have access.

Now, the goodness or badness of various forms of romantic love might seem like a fine thing to ponder, but also something that has little *practical* importance. For famously, falling in love seems to just *happen* to people. Falling out of love seems similarly out of one's hands, no matter how much one might desire to be free of his or her romantic feelings. This is perhaps particularly true of HCs. And if we accept the axiom that ought implies can, the question of what one ought to do in regards to being in HClove may not seem to be one that sensibly arises.

However, recent empirical evidence suggests that there are in fact ways to control one's susceptibility to HClove, and that a kind of romantic 'inoculation' can be delivered in the form of a pill. According to some researchers, elevated levels of serotonin in the brain are antagonistic to the neurological processes responsible for romantic love and sexual attraction (Helen Fisher 2004, 197). The level of serotonin in our brain is something that we can control now - at least to some extent - with certain modern anti-depressants, which function by elevating brain levels of serotonin. It follows that we now can control, at least to some extent, our susceptibility to being in romantic love, including, presumably, HClove.

If one believes that every form of romantic love is *always* a good thing among those who are psychologically healthy (according to conventional criteria), then the possibility that certain drugs could interfere with being in love could *only* be worrisome. In this spirit, one team of researchers warns that 'artificially elevating serotonin activity can endanger your ability to fall in love' (ibid). A different view, and one that I will defend in this paper, is that being in (romantic) love is *very often* a good thing, but that HClove is almost always a bad thing. If I am right about this, then medicine that can influence our romantic experience may be something to *welcome*, as a modern tool that can help in achieving the ancient goal of human flourishing.²

But whether we should do something about HCllove, and if so, what, I will leave for another time. In this vein it is enough, for now, to show that the question of the desirability of HCllove has a certain urgency and interest that it otherwise would lack. And luckily, just as technology has increased the urgency and interest of the question of the desirability of HCllove, it has also provided us with new resources to find an answer to the question.

2. The Tyrannical Soul

In this section I will explicate Socrates' characterization of lawless desires, and the related notion of a tyrannically souled individual. I will also sympathetically explicate his argument that the tyrannically souled individual is both miserable and prone to immoral actions in a conventional sense.

A. Tyrannical Souls and Lawless Desires

A soul is tyrannical, in Socrates' sense, when it is controlled by a certain small set of what he calls 'lawless' desires. Desires are lawless when they satisfy the following conditions. They must spring from the animalistic, non-rational part of the soul that generates desires for food, drink, sex, and other primitive desires. Their satisfaction must not be something that is beneficial to their possessor. And, what is most important for my purposes, these desires must be *insatiable*.

Importantly, in the souls of those who are psychically healthy, lawless desires almost always exist but are held in check by reason.³ When these desires are in *control* of a person's soul, however their possessor recklessly and single-mindedly pursues their object without any regard to the demands of reason or the satisfaction of their other desires. This must be why Socrates calls such desires 'lawless': when sufficiently strong they cause their possessor to trample over every obstacle in their way. This is due to a combination of their insatiability and strength of these desires in the tyrannical soul. Note that insatiability alone does not cause a desire to be recklessly and single-mindedly pursued. For I might have many such unsatisfiable desires whose expression I am able to hold in check by keeping the strength of these desires sufficiently low. And strength of desire alone does not cause its possessor to be single-minded about its satisfaction in any kind of sustained way. For, no matter how strong a desire might be, once it is satisfied, that desire no longer has the kind of motivational force that would cause one who possessed the desire to continue to pursue its satisfaction and wholly neglect the satisfaction of one's other desires. A desire that is *both lawless and sufficiently powerful* to dominate a soul I will call a tyrannical desire (though this is not a term Socrates himself uses). And the soul that contains tyrannical desires is what Plato would call a tyrannical soul.

Note that a tyrannical desire is not, of *logical* necessity, a malicious or otherwise morally objectionable desire. They often in fact have immoral objects, but this is a *by-product* of their insatiability and strength. They tend to settle upon inappropriate objects for two reasons. First, the tendency for strong, insatiable desires to shift restlessly from one object to another, which increases the chance that one will eventually end up desiring inappropriate objects. For instance, suppose one has a powerful and insatiable sexual desire. Morally acceptable means of satisfying this desire will not satisfy it, simply because there is *no* means of satisfying the desire. And so those who have been thus 'implanted with the sting of longing', are prone to desiring inappropriate sexual objects. This does not mean that those with powerful and insatiable desires will end up desiring just *anything*. For example, an insatiable and powerful sexual desire might be limited to objects—that is, persons—who are of a certain morally appropriate age. Though a person in the grip of such a desire might

desire, say, sexual encounters outside of a relationship such as marriage that prohibited these encounters, the same person might have no desire for such encounters with children. Desires can be more or less narrow in their scope, and powerful insatiable desires that are nevertheless sufficiently narrow in their scope, where the scope is ‘narrowed’ so as to exclude certain morally inappropriate objects, will be less likely to settle upon morally inappropriate objects.

The second reason that powerful, insatiable desires can end up settling upon inappropriate objects is that one naturally forms a *derivative* desire to remove obstacles that come between the person in the grip of such a desire and the object of their desire. One, then, can come to have a powerful derivative desire to remove any number of obstacles, and the removal of these obstacles may well not be morally permissible. If I have a powerful, insatiable desire for a drug, but I do not have enough money to purchase the drug, I might form a desire to threaten my drug supplier with bodily harm unless he or she gives me the drug. Or I might form a desire to steal money from my younger sibling. And so on.

B. Tyranny and Happiness

The tyrannically souled person will pursue the satisfaction of her lawless desires vigorously, and, because they can never be satisfied, pursue their satisfaction with an undiminished intensity. This naturally leads to the neglect of other, more beneficial desires, and the demands of reason that would otherwise receive appropriate attention. A lawless desire in a tyrannical soul can be likened to a child in a crowded day care centre who not only cries the loudest, but who can never be pacified, leading to the neglect of her less clamorous (and better behaved!) companions. Leaving a discussion of the Forms aside, we can now see why the tyrannically souled person is supposed to be the most miserable of all persons. The happiest person, according to Socrates, has a soul in which reason governs the entire soul, and so allows each part of the soul to be satisfied by the objects appropriate to it. But in the tyrannically souled person, the best parts of her—including her reason and her beneficial desires—are utterly unsatisfied. This, according to Socrates, means that the individual is unhappy.

C. Tyranny and Conventional Justice

The discussion of lawless desires and tyrannical souls comes in the context of the broader project in the Republic, of course, a large part of which is a defense of a particular account of the nature of justice. The view of justice Socrates defends is rather counterintuitive. He claims that justice and injustice are direct functions of an individual’s *internal* state—the state of the soul—rather than of external actions. One is just, then, simply when one’s soul is properly ordered, which means, very roughly, that each part of the soul serves its proper part, with reason governing the spirit and the appetites. Plato has famously been criticized for not making a convincing case for the connection between *conventional or external* morality and immorality, and justice and injustice in the soul (David Sachs 1963). It is worth noting that Plato himself was not primarily concerned to show the connection between his internal notion of justice and conventional, external justice, since, to his mind, he had given a compelling argument for thinking that internal justice was the only sort of justice worthy of the name. For those of us who remain skeptical on this point, we will want to know more about the connections between Platonic justice and conventional justice. If Plato has given me a reason to have a properly ordered soul, in what way will this give me a reason, say, to be kind and fair to my neighbour?

The criticism is fair enough, I think, when it comes to the connection between a Platonically *just* soul and conventionally *just* actions. We have good reason to wonder

whether a perfectly Platonically just person might not also commit conventionally unjust actions, even to the point where, by conventional standards, our Platonically just person would be considered horribly unjust. But the connection between the internal structure of the soul and external behaviour is more obvious in the case of internal *injustice*. This has already been hinted at in IIA: since lawless desires pay no heed to anything besides their own satisfaction, pursuing the satisfaction of these desires will cause the tyrannically souled to trample on conventional justice whenever doing so will serve in the cause of their (the lawless desires') satisfaction. This could be because the tyrannically souled individual has an inappropriate object of primary, underived desire, or because the tyrannically souled has an inappropriate object of derived desire (that is, a derived desire to remove some obstacle whose removal is morally inappropriate). The lack of coincidence of the actions of the tyrannically souled and the requirements of conventional morality is, then, to be expected.

3. The Tyranny of HCllove

I have explicated Socrates' characterization of lawless desires and the tyrannical soul, and I have also sympathetically explicated his argument that the tyrannically souled individual is both unhappy and prone to acts of conventional immorality. I will now argue that the desires of those in HCllove approximate my description of tyrannical desires; that HCs approximate Socrates' description of the tyrannically souled; and that HCs are, in the Platonic sense, unhappy.

A. Methodology

I will take it that establishing that an individual has a tyrannical soul also establishes that they are possessed by tyrannical desires, and that entailment runs the other way as well. I will also take the establishment that an individual has a tyrannical soul *or* is possessed by tyrannical desires to establish the further claim that the individual is unhappy, in the Platonic sense. These are conceptual claims. For a person who has a tyrannical soul *by definition* fails to satisfy their non-tyrannical and higher desires. A tyrannically souled person might claim to be happy, for instance, and sincerely believe that they are happy. But if they possess a tyrannical soul then, regardless of their happiness in some non-Platonic sense, it is logically not possible that they are happy in the Platonic sense. Whether one thinks that her state is undesirable will depend upon whether one accepts something like Plato's intuitively appealing account of happiness, a full defence of which, unfortunately, is beyond the scope of this discussion. These conceptual relationships I will take as understood, and will not be pointed out in the discussion that follows.

Importantly, though, conceptual analysis can only do so much here. Empirical data will be necessary in establishing the truth that 1) HCs are tyrannically souled or 2) are possessed by tyrannical desires. *Only* once at least *one* of these claims is established can the conceptual argument just outlined be of any help. And the best way to establish one of these claims, in the absence of knowing its logical corollary, is to use empirical evidence. So I will appeal to what I think are common, non-clinical observations about those in love, and also some more formal research of the behaviour and neurobiology of those in love.

Some of the arguments I will offer below will attempt to show that those in HCllove exhibit behaviour consistent with the tyrannically souled, thus suggesting that those in HCllove are tyrannically souled. But to *prove* that HCs have a tyrannical soul it is not quite enough to show that HCs evince behaviour consistent with the behaviour of the tyrannically souled. To do so would be to 'affirm the consequent'. For Socrates has only established that the tyrannically souled evince certain characteristic behaviours, *not* that *all* such behaviour

comes from the tyrannically souled. Still, showing that an individual evinces behaviour consistent with that evinced by the tyrannically souled individual is highly suggestive, since the best explanation of the phenomenon - in the absence of other plausible candidates - looks to be the presence of a tyrannical soul.

With that said, there will still inevitably be a desire to avoid the uncertainty of behavioural research on the psyche or soul, and to probe the psyche more *directly*, through neurological research. Such research on the brain of those in love is terribly exciting, and *potentially* the best way to confirm the tyrannical nature of HCllove, since it holds out the promise of letting us look *directly*, even if incompletely,⁴ into Plato's 'soul'. This enthusiasm must be tempered, though, with a certain humility about our limited ability to *interpret* neurological data.⁵ Our position is not unlike that of an anthropologist who, after years of unearthing and dissecting and analysing artefacts in order to figure out the laws of a particular society, has discovered a perfectly preserved tablet that elaborately enumerates these very laws. Unfortunately, she has only a rudimentary understanding of the language in which the tablet is written. Her problem, like ours in large part, is one of interpretation, rather than of access to good data. Of course, our anthropologist would be foolish to ignore her recent discovery, and we would be equally foolish to ignore the neurological evidence. We must simply maintain a proper modesty about our abilities of interpreting it.

B. Evidence from Common Experience

HCs are famously single-minded in their thoughts and actions. In focusing their thoughts and actions on the object of their love, HCs often act in ways that are contrary to what is reasonable or moral, even when they can *acknowledge* that this is so. Their reason can still see, even if only dimly, the unreasonableness of their behaviour; but their passion is evidently lawless and also terribly powerful, so that those under the sway of this passion will ignore the demands of reason and the satisfaction of desires that would be more beneficial. For instance, many of those in HCllove grant that it would be *reasonable* to spend less time with their beloved and more time with their friends or family, or more time on their studies and other commitments and projects. Or they might admit that they are spending unreasonable amounts of money on their beloved, or taking unreasonable risks to be with their beloved, or hurting others in ways they acknowledge they ought not. They might acknowledge that being with their beloved is bad for *themselves*—perhaps their lover neglects them or mistreats them in certain ways. Or they might acknowledge that remaining with their partner is bad for their *partner*—perhaps, because of circumstances, they can see that their partner would be better off without them, or better off with someone else. But they persist in their relationship against their better judgment, or against the promptings of their reason.

When things go badly, such as when they are rejected, HCs are not infrequently engulfed with anger and take steps to harm their love-object. This need not take the form of physical violence, and, for HCs, as I have characterized them, it usually does not. After all, I have claimed that, among other things, most of us have been in HCllove at one point or another, but most of us have not committed physical violence against our lovers. More common, probably, are attempts at emotional violence. Frequently a jilted lover is verbally malicious to their love object, or tries to induce psychological pain through jealousy, by damaging their lover's reputation, or some other means.

The best explanation of all of this seems to be that HCllove is both insatiable and terribly strong, and so tyrannical in nature.

C. Psychological/Neurological Evidence

Researcher Helen Fisher has done some remarkable work in which she took fMRI scans of those who were evidently passionately in love. These scans allowed her to record the level of blood flow to different parts of her subject's brains when exposed to pictures, gifts, and other reminders of their love object. She screened her subjects by selecting only those who seemed, to her, most in love. The most important factor in determining the degree to which her candidate subjects were in love was the amount of time her subjects reported thinking about their beloved. All but one of her twelve subjects was evidently psychologically normal according to accepted psychological standards, though each claimed to have obsessive thoughts about his or her partner, often saying that they thought about their partner over ninety percent of the time. She also selected only those candidates who appeared to be yearning for or craving their beloved (ibid., 63).

When Fisher scanned the brains of those thinking of their lover, she found that the most central and most primitive area in the brain became active. Interestingly, it was more active the more in love her subjects reported themselves to be. This core area of the nervous system is part of what is called the *reptilian brain* since this brain region evolved long before the mammals proliferated about 65 million years ago. The reptilian brain, among other things, is part of the brain's 'reward system', the network for arousal, sensations of pleasure, and the motivation to acquire rewards (ibid., 69). This part of the nervous system, then, is evidently *not* a part of the brain that is responsible for humans' unique ability to reason. Thus, romantic love seemingly meets one of Plato's criteria for being lawless, namely that it springs from a non-rational, appetitive part of the soul or psyche.⁶ (With some trepidation I venture that if we were to try to classify the brain into a three-part system in the way that Plato attempted to classify the soul, this most primitive part of the brain would roughly match Plato's description of the appetitive part of the soul.⁷)

In her research, Fisher found that the *Ventral Tegmental* area of the brain also became highly active in her subjects when they were focusing on their love object. This area is what she describes as the "mother lode" for dopamine-making cells. She explains that dopamine produces focused attention, fierce energy, and concentrated motivation to attain a reward - all of which is characteristic of tyrannical desires.

She additionally notes that the elevated levels of central dopamine characteristic of romantic love are also characteristic of the brains of those in the grip of some other basic drive, such as the drive for hunger or thirst (ibid., 75). She suggests, partly based upon these neurological findings, that romantic love can be thought of primarily as such a basic drive - rather than, say, an emotion. This means that romantic love shares the characteristic features of basic drives, of course, which includes being focused on a specific reward and being difficult to control. Again, these are characteristic of tyrannical desires.

Further evidence for the focused and near-irresistible nature of romantic love comes from a study that suggests a neurological similarity of passionate love to the psychological disorder known as obsessive-compulsive disorder (OCD). This is a disorder that, familiarly, is characterized by irresistible thoughts and actions, usually highly focused on some particular subject matter. As mentioned, Fisher's subjects reported obsessive thinking, which is by itself suggestive of this similarity. Additionally, though, the study found that those who suffered from unmedicated OCD and those who were in love had significantly lower levels of serotonin than did those who 'suffered' from neither (D. Marazziti, 1999).⁸

So far we have gone over psychological evidence that suggests that HCllove springs from the non-rational part of the psyche, and that it is focused and extremely difficult to control. But we have not yet presented psychological evidence that it is *insatiable*. An important piece of psychological evidence for its insatiability comes from its similarity to chemical addiction, a disorder characterized by, among other things, a desire (for a substance) that either cannot ever be satisfied, or whose satisfaction is typically fleeting. Fisher writes that nearly all drugs of abuse affect a single pathway in the brain, the mesolimbic reward system, which is activated by dopamine. Romantic love, she found, stimulates parts of the same pathway with the same chemical. Another team of neuroscientists backed up this finding. They found that the brain scans of those who reported being passionately in love were active in many of the same areas as those who had taken cocaine or opioids (Helen Fisher 2004, 182).

Besides these *neurological* similarities between lovers and drug addicts, according to Fisher passionate lovers show the classic *behavioural* symptoms of drug addiction as well. First, lovers grow tolerant to romantic love. Though at first lovers need to be with one another only occasionally, as time passes they need to see more and more of each other, just as the addict, up to a point, needs more and more of their drug to feel satisfied.⁹ Second, if the relationship is interrupted or broken off, lovers show the common signs of drug withdrawal: depression, crying spells, anxiety, insomnia, loss of appetite, irritability, and chronic loneliness. And, like drug addicts who in withdrawal will sometimes take extraordinary means to acquire their drug, lovers will sometimes go to unhealthy, humiliating, immoral, and even physically dangerous lengths to be with their beloved (ibid., 183).

D. Summarizing the evidence

A quick summary and analysis of the evidence will be helpful.

First, we have provided evidence that HCllove is lawless: 1) neurological evidence on HCllove suggests that the faculty of reason *does not generate its ends*; 2) evidence from common experience suggests that the satisfaction of the HCllove is *not beneficial* in the way that, say, the satisfaction of the desire for (satiabile) hunger is; and 3) common experience and psychological evidence suggests that HCllove is also *insatiable*.

Showing that HCllove is lawless is part of the argument that it is tyrannical, since a desire is tyrannical if (and only if) it is both lawless and sufficiently powerful. The second part - showing that HCllove is sufficiently powerful to count as tyrannical - has come from common experience and psychological research. The evidence here strongly suggests that HCs are extraordinarily motivated to pursue the objects of their desire, such that they find their desires nearly irresistible.

If the argument has been successful, then it has been established that HCllove is tyrannical. The conceptual argument given in IIIA would entail, then, that HCs are tyrannically souled, and also unhappy in the Platonic sense.

4. Objections

A. One might object that love, passionate or otherwise, could not possibly qualify as a tyrannical desire, since tyrannical desires are malicious or perverted, while love is often or usually altruistic, and so morally admirable.

I should first note that I believe that there are surely many forms of love, including romantic love, which are not tyrannical, and are instead wonderful and beneficial and healthy in all of the ways that matter.

That said, I maintain that even altruistic love could be tyrannical. Part of the reasoning here is contained in 2A, where I argued that a desire is not tyrannical in virtue of the *moral appropriateness* of its object, but rather in virtue of being both lawless and sufficiently strong. If some form of romantic love meets both of these conditions, then, regardless of the appropriateness of its object, it is tyrannical.

One might respond by claiming that, if altruistic love, of all things, can be tyrannical, then this just shows there must not be anything bad about tyrannical desires! Now, I have conceded that there is no *logical* connection between tyrannical desires and immoral desires and actions. However, as a matter of practice, a desire that is both insatiable and sufficiently strong will have a tendency to generate immoral desires and actions that serve to satisfy those desires. So in the case of altruism, I might want the good of my lover so strongly and insatiably that I shirk even my most stringent duties to others whenever their fulfillment conflicts with the promotion of my lover's welfare. Or I might make unreasonable personal sacrifices, like giving up on some life project or commitment, in order to promote her welfare. And so on.

B. One might object that I have not given good enough evidence that HCllove meets each of the three conditions for being a lawless desire, and so I have not shown that HCllove is tyrannical. Specifically, the argument that romantic love springs from a non-rational part of the brain might be thought highly conjectural. One might even think that the claim about its non-rational origins is falsified by the large element of reasoning that evidently goes into most forms of love, including HCllove.

One response is that it is not logically inconsistent that one might devote great cognitive resources to an end that is itself just a basic desire generated by a non-rational part of the psyche. I might, for instance, have the basic desire to eat, a desire generated by a non-rational part of my psyche. But I am also, unfortunately, lost in a foreign country whose language I do not speak. In order to satisfy my non-rational desire to eat I will need to use my reason extensively. An end, then, could be set by a non-rational part of the psyche, while the means required to achieve that end could be calculated by a rational part.

Additionally, though, even if I have not shown that HCllove is generated by a non-rational part of the psyche, and so not, strictly speaking, lawless or tyrannical, I have given compelling evidence to think that HCllove is a bad thing. Regardless of its psychic origins, if I have shown that HCllove disposes its possessor to be imprudent, obsessive, and immoral—the arguments for which do not turn on any premises regarding HCllove's origins—then I have shown enough for my purposes.¹⁰

Notes

¹ One might hope that this state would have, like many emotions are alleged to have, some cognitive content that distinguishes it from other states. However, I will also not attempt to articulate any distinctive cognitive content it might have, and I am additionally going to remain neutral as to whether the species of romantic love I am concerned with is properly characterized as an emotion in the first place. I leave these questions open because it is not clear to me how to answer them, and also because the settling of these questions is not clearly needed for the purposes of this discussion. The most obvious features of the state I have in mind, as I am conceiving of it,

are its motivational and affective aspects, and so it is these that I will focus on.

² Of course, even if we can show that a certain kind of romantic love is bad, it does not automatically follow that one ought to take steps to prevent oneself from falling in (this kind of) love, much less that one ought to take steps to prevent others from falling in (this kind of) love! In a Humean spirit, from the mere goodness or badness of any states of affairs no practical conclusion *automatically* follows (though a teleological account of ethics will happily insert the needed premise(s) to generate a practical conclusion). Further, even if one successfully argues that *some* such steps ought to be taken, it does not, of course, automatically follow that pharmaceutical intervention is an appropriate way of ‘immunising’ oneself against passionate romantic love.

³ And reason’s helpmate, spirit. I will not give any kind of detailed explanation of the role of spirit in the governance of Plato’s soul, though, as it is not needed or helpful for the purposes of this discussion.

⁴ A believer in wide mental content might think that even a direct probing of the human brain will be unable to reveal, on its own, certain aspects of a person’s ‘soul’. For certain mental content, the thought goes, is determined not simply by what is ‘in the head’, but also by one’s environment.

⁵ As someone trained in philosophy rather than psychology, I cannot help feeling especially humble about my abilities of interpretation.

⁶ Of course, the fact that a desire emanates from an irrational part of the soul does not mean that it *could* not be controlled by the rational part. It has been pointed out to me that the neurological evidence suggests that desires do not lead in an unmediated way to action, but rather are ‘routed through’ the prefrontal cortex, and so are not isolated from ratiocination. Even so, as long as the *source* of a desire is an irrational part of the psyche, then it satisfies one of Plato’s criteria for being lawless. See also section 4, part B of this paper.

⁷ See Paul MacLean (1990) for an account of what he calls the ‘triune brain’, the parts of which, in a very rough way, intriguingly correspond to Plato’s three parts of the soul.

⁸ It is important to note that a limitation of this study is that brain levels of serotonin were not directly measured, but only the level of serotonin in blood in the rest of the body.

⁹ This tolerance is usually not unlimited in either case: passionate lovers seem to eventually reach a threshold where they are satisfied with the amount of time spent together, even if this is a great amount of time. And drug addicts will similarly usually reach a dosage beyond which they do not desire to go.

¹⁰ I am grateful to an anonymous referee of this journal for many helpful comments on an earlier version of this paper.

References

Fisher, Helen. (2004). *Why We Love*. Henry Holt and Company, New York.

MacLean, Paul. (1990). *The Triune Brain in Evolution*. Plenum Press, New York.

Marazziti, D., H.S. Akiskal, A. Rossi, and G.B. Cassano. (1999). ‘Alteration of the platelet serotonin transporter in romantic love’, *Psychological Medicine* 29: pp. 741-45.

Plato. (1992). *Republic*. Translated by G.M.A. Grube, (revised by C.D.C Reeve). Hackett Publishing company, Inc., Cambridge.

Sachs, David. (1963). ‘A Fallacy in Plato’s Republic’, *Philosophical Review*, vol. 72, pp. 141-58.

Being Queasy about Reconstructing Animals

Michael Hauskeller

University of Exeter

1. Introduction

Although a considerable number of people are concerned about biotechnological manipulations of living beings they usually find it hard to explain what exactly they think is wrong with it. When being asked on which moral grounds such a manipulation (or what forms of manipulation) can be reasonably opposed they are at a loss for an answer. They just feel that it is wrong but can't say why. Hence it is easy to dismiss their concerns as unfounded. Scientists and philosophers tend to believe that biotechnological manipulations of living beings can only be wrong if there is some suffering involved. Some would even say that it is only the suffering of human beings that is morally relevant. The suffering of animals, then, simply does not count. Yet there are things that can be done and are done to living beings, including humans, which, even if they do not involve or produce any suffering, are still considered to be morally wrong by a large proportion of the public. Among them are, for instance, changing the nature of living beings by means of genetic engineering in order to enhance their health, or, more likely with animals and plants, their utility, or impairing their ability to live autonomously, or unduly instrumentalising them.

Often when such concerns are expressed, they are rejected on the grounds that there is no rational basis for them. They are claimed to be irrational, to 'make no sense' (Gibbard 1990). Accordingly, they are not regarded as proper moral concerns at all. But if they are not moral concerns, what kind of concerns are they? One answer often heard is that these concerns are at best aesthetic ones which therefore, being a matter of mere taste, can be justly disregarded. It is this two-fold claim which interests me in this paper. Is the distinction between proper moral concerns and 'merely' aesthetic ones really sound *in the context in which this distinction is applied*? And can we really make no (moral) sense of the uneasiness which many people feel about biotechnological manipulations of living beings, regardless of whether, as a consequence, those beings suffer or not? While dealing with these questions, I will not attempt to give a precise definition of the aesthetic as opposed to the moral. Rather, I will use the term 'aesthetic' in the same wide and admittedly vague sense in which it is used in the arguments that I will discuss. Since nothing seems to hinge on it, a precise definition is not necessary. What is relevant here is the rhetorical function the term 'aesthetic' acquires in claims that certain common concerns are not proper *moral* concerns – in spite of the fact that people who share those concerns think of them as moral. The question is not whether those concerns are aesthetic but whether they are *merely* aesthetic in the sense of being not morally considerable.

2. Eliminating suffering

Human beings have always manipulated and exploited nature, which, as such, is unobjectionable and does not differ much from what other living beings do to survive (Rowlands 2000). After all, it is simply a matter of survival to learn about the many ways in which the environment can be put to use, and to act accordingly. With the rise of genetic engineering it has apparently become even easier to create the exact world we want and need for our survival and wellbeing. Thanks to this comparatively new technique we can now do more than merely exploit animals as we happen to find them, or at best adapt them slowly

and patiently to our needs by conventional breeding. Instead we can literally redesign them in a very short time in such a way that they serve our various needs more effectively. Thus genetic engineering is applied - or its application is urged - in order to enhance the productivity of farm animals to a hitherto unknown extent (e.g. GM super salmon), but also to deliberately create animals that produce medically important proteins for use as 'pharmaceutical factories'. Others serve as experimental 'models' of human diseases, or produce organs for xenotransplantation (so called 'spare parts factories') (Bowring 2003, 119-122).

It is clear that in all these cases we are not primarily concerned with the good of the animals we are using. We do not think about what is good for them but rather what they are, or might become, good for. On the other hand, we *could*, of course, use genetic engineering also to actually help animals, for instance to cure them from genetic diseases or to enhance their disease resistance. We might even help them to suffer less or, ideally, not at all from the conditions human society forces upon them. Thus genetic engineering is not intrinsically objectionable in terms of animal welfare considerations. This position has been prominently adopted by Bernard Rollin (1995). Considering it unlikely that we will simply stop using animals as much as we possibly can, Rollin argues that from an animal welfare point of view it is justified, or even morally required, to effect genetic changes in the animals such that we can reduce the pain and suffering they would otherwise endure. If, for instance, we could, by means of genetic engineering, create cattle without horns (so we would not have to dehorn them later, which is quite painful to them) or eliminate the drive to nest in chickens kept in battery cages, then why should we not do it? What could be wrong with removing a source of suffering? Hence Rollin is convinced that, if 'changing the animals by genetic engineering is the only way to assure that they do not suffer (...), people will surely accept that strategy, though doubtless with some reluctance.' (1995, 175)

However, presenting the alternatives like this seems to anticipate the outcome. If you already take it for granted that there are no other alternatives than either to let the animals suffer or to modify them, you will feel that it is hardly justifiable to oppose the modification. But I don't think we can or should take it for granted that changing the animals' environment is not possible or realistically to be expected. It is like asking, Would you rather let the Iraqis attack the US with weapons of mass destruction than invade Iraq by military force and take these weapons from them before it is too late? The consequence of putting the question this way is that no other alternatives are seriously considered.

Moreover, why should people feel reluctant to use genetic engineering if it is only in the interest of the animal? The main reason, says Rollin, is a certain 'queasiness' which is at root aesthetic, i.e. not moral: 'The chicken sitting in a nest is a powerful aesthetic image, analogous to cows grazing in green fields. A chicken without that urge jars us.' (1995, 175) It is odd, however, that Rollin, while on the one hand asserting that there is nothing intrinsically wrong with changing an animal's nature (or *telos*, as Rollin prefers to call it), on the other hand says that it is 'certainly a poor alternative' to alter the animal instead of its environment (1995, 171). He even calls it 'the lesser of two evils' (1995, 192) - the other evil being the pain the animal would have to endure if we did not alter it. But if there is nothing *morally* wrong with changing an animal's genome and through it the animal's physical appearance and structure of behaviour, then why should this be classified as 'a poor alternative', let alone an 'evil'? If the revulsion we (or some of us) experience here is, as Rollin believes, purely 'aesthetic', then from a moral point of view this option is not poorer or less desirable than the other (and most certainly not 'the lesser of the two evils').

Yet how exactly do we distinguish here between purely aesthetic concerns and purely moral concerns? What Rollin apparently means by 'aesthetic' is that our revulsion does not rest on any beliefs concerning the situation the animal is objectively in, but rather on the fact that this situation or the animal itself is somehow *unpleasant to our senses*. Our revulsion is a matter of taste and does not involve any kind of moral judgement. If, however, it is just not as pleasant to see a hen in a cage as to see it sitting in a nest, can and should we then not ask *why* it is not so pleasant? Does our revulsion really have nothing to do with the situation the hen is objectively in? Are we, as Rollin suggests, just being 'queasy'? How can we be so sure? After all, to many of us, it at least *seems* as if we had a genuine moral concern here, as if changing an animal's nature so that it no longer suffers from the bad conditions it is forced to live in, were actually something wrong, i.e., something we simply ought not to be doing. How can we, without begging the question, say with confidence that our concerns, though they seem to us to be *moral* concerns, are not moral concerns at all, but 'merely' aesthetic?

To distinguish so sharply between moral and merely aesthetic concerns gets even less plausible when we consider another of Rollin's suggestions. Rollin believes that the common practice of using animals as models for human diseases is, at least potentially, so beneficial for humans that it is extremely unlikely that we will ever give it up. Since, however, this practice involves a lot of suffering for the animals being used which, in order to study the disease properly, have to be kept alive for quite a long time while anaesthesia can only be applied for a comparatively short time, it is morally not sufficient to merely change the animals' subjective experience. If we really care for the animals' welfare we ought rather to try to eliminate their consciousness completely by, for instance, removing or destroying their cerebral cortex (1995, 205). We might even some day be able to genetically engineer them in such a way that they are already born decerebrate and with the diseases we want to study. Assuming that 'the capacity for genetically engineering models for all manner of genetic and other diseases is imminent, and that the research community will forge ahead in creating such models', decerebration is, according to Rollin, not only 'the only viable way to control suffering' but also in perfect harmony with the moral principle of conservation of welfare (1995, 205). Thus the moral principle Rollin applies and which seems to be the only one he accepts as a *valid* moral principle ends up as something like this: Avoid making sentient beings suffer but if you can't (or don't find it convenient to) avoid it by any other means then you should kill it, either physically if you no longer need it (1995, 197), or mentally, if you still need its living body for your purposes. This, however, is surely a curious moral principle and certainly not in accordance with our moral intuitions, or at least with what we take to be our moral intuitions.

With Bentham, Rollin obviously assumes not only that what morally matters is whether animals can suffer, but also that this is the *only* thing that matters. According to this view, genuine moral concerns are concerns about good and bad subjective experiences. All other concerns are, whatever else they might be, definitely *not* moral concerns. That is, the norm that is accepted here does not permit any other concerns to be included in the moral domain (cf. Gibbard 1990, 7). But why should we accept this norm? If we ask ourselves how we know that inflicting pain and suffering on other living beings ought to be avoided or at least requires moral justification, we can hardly say more than that is how we feel about it. Thus, the grounds for our belief that wilfully inflicting pain on others is bad does not seem to be much different from the grounds of our 'queasiness' about reconstructing animals according to our needs and in such a way that they become something quite different from what they, as it may seem to us, by their very nature are meant to be.

A claim like this does, of course, need some backing which I would like to provide

through a discussion of H.G. Wells' novel *The Island of Dr Moreau*. My reasons for choosing this slightly unusual method of *logon didonai* by analysing the contents of this particular novel are not only its subject matter (the intentional remodelling of living beings) but even more so the implicit deconstruction of a normative system of moral rationality which is believed by its proponents to be entirely autonomous and free from any aspects that might be denounced as merely 'aesthetic'. What a good novel like Wells' can do – and sometimes better than a purely conceptual analysis – is to show the inadequacy of conceptual distinctions. In this case, it is the distinction between moral concerns and so-called aesthetic ones which are shown to blend seamlessly into one another (cf. Blackburn 1998, 12). Never mind that it is 'fiction', what counts is that it gets to the heart of the matter.

3. The ethics of the matter

The Island of Dr Moreau tells the story of the amateur biologist Edward Prendick who has the misfortune to become stranded on an island which is inhabited by the notorious vivisectionist Dr. Moreau, the alcoholic physician Montgomery who assists Moreau, many animals captured in order to be researched upon, and finally the strangest human beings Prendick has ever encountered and which he is at first not able to make head or tail of. He just knows that there is something wrong with them and that he can hardly bear the sight of them, a feeling that he shares with the crew of the ship on which he travels to the island together with Montgomery and one of those disturbing creatures. On being asked by Montgomery how this creature, which he mistakenly takes to be human, strikes him, he answers: 'He's unnatural. ... There's something about him. ... Don't think me fanciful, but it gives me a nasty little sensation, a tightening of my muscles, when he comes near me. It's a touch ... of the diabolical, in fact' (1923, 42). Now this revulsion which Prendick describes here is surely a good example of what Rollin would classify as an aesthetic reaction, for Prendick does not seem to think that the creature is suffering under its condition, or at any rate this is not what gives him this 'nasty little sensation'. In fact, there is no belief whatsoever underlying his aversion to the creature. He has no clear reason for being revolted: he just is. It is only later that he begins to suspect the truth: 'Suddenly, as I watched their grotesque and unaccountable gestures, I perceived clearly for the first time what it was that had offended me. (...) Each of these creatures, despite its human form, its rag of clothing, and the rough humanity of its bodily form, had woven into it, into its movements, into the expression of its countenance, into its whole presence, some now irresistible suggestion of a hog, a swinish taint, the unmistakable mark of the beast.' (1923, 50). Although at first he misinterprets what he sees, thinking that Moreau has somehow changed humans into animals, he soon finds out what is really happening, namely, that Moreau has in fact reconstructed the captured animals in such a way that they resemble, or have become, human beings. Learning, however, that the creatures which have struck him as so strange are in fact former animals turned halfway into men, does not lessen his abhorrence. On the contrary, it gives his initial, instinctive reaction a certain justification. He understands now *why* he has felt the creatures to be abhorrent, and incorporates this knowledge into his perception of them. Yet his feelings about them remain unchanged. In a discussion with Moreau about the justifiability of his experiments he still thinks of the beast-people, as he calls them, as 'abominations' (1923, 95). They are, in fact, regular monsters which perfectly fit the definition given by Brittnacher (1994) in his seminal study on the aesthetics of horror. For, according to Brittnacher, a common denominator of all beings designated as monsters is their 'excessive deviation from the norm of physical integrity. In the physical extremity of the monster, the human and the animal spheres overlap and the idea of an animal kingdom which is neatly organized in species is being revoked' (1994, 184).¹

Now it is interesting to note that, when Prendick first arrives on the island he wonders why Moreau and Montgomery are being so secretive about the kind of research they are undertaking there. This irritates him since he suspects Moreau to be practising just the usual vivisection of animals which Prendick considers, however painful it might be for the animals involved, 'especially to another scientific man, (...) nothing so horrible (...) as to account for this secrecy' (1923, 40). This statement, which expresses the common conviction of many scientists that for the sake of knowledge or human good we are justified in subjecting animals to a great deal of suffering and pain, echoes a remark made earlier by Moreau himself, when he tells Prendick that what he is doing on his island is 'nothing very dreadful really – to a sane man' (1923, 35). Moreau, however, *knows* what he is doing, and what he takes to be 'nothing very dreadful' is not only the infliction of pain on the animals during the process of reconstructing them but also the whole business of creating beings that are half human and half animal. Thus, while Prendick considers the infliction of pain justifiable – though also seeing the need to justify it – he is by no means prepared to accept the creation of hybrid creatures such as the beast-people inhabiting Moreau's island.² So if his reaction is really purely aesthetic, as Rollin presumably would have us believe, then it is more strongly grounded and much harder to overcome by considerations of utility than the 'proper' moral concerns he has about the pain the animals have to endure during the process of reconstructing them. Of course, Moreau himself believes that neither the infliction of pain nor the outcome of his experiments need to be justified by the prospect of any good resulting from it. When Prendick asks him for a reason and for a justification for inflicting all this pain Moreau replies that

it is just this question of pain that parts us. So long as visible or audible pain turns you sick, so long as your own pains drive you, so long as pain underlies your propositions of sin, so long, I tell you, you are an animal, thinking a little less obscurely what an animal feels. This pain (...) is such a little thing. A mind truly opened to what science has to teach must see that it is a little thing. It may be that, save in this little planet, this speck of cosmic dust, invisible long before the nearest star could be attained – it may be, I say, that nowhere else does this thing called pain occur. But the laws we feel our way towards ... (1923, 92-93).

And later he adds:

Pain! Pain and pleasure – they are for us, only so long as we wriggle in the dust ... You see, I went on with this research just the way it led me. That is the only way I ever heard of research going. I asked a question, devised some method of getting an answer, and got – a fresh question. You cannot imagine what this means to an investigator, what an intellectual passion grows upon him. You cannot imagine the strange and colourless delight of these intellectual desires. The thing before you is no longer an animal, a fellow-creature, but a problem.

To this Prendick can only reply, rather helplessly, that, whatever Moreau might say, 'the thing is an abomination' – a declaration which moves Moreau to the laconic comment: 'To this day I have never troubled about the ethics of the matter' (1923, 94).

This passage strikes me as important for two reasons. First because of the remarkable extent to which Moreau's more than a century old description of the scientific method, including its conceptual transformation of fellow-creatures into problems, is mirrored by the language and the practice applied by many present day scientists working in the field of biotechnology.³ And second – more relevant to the object of this paper – because of the emphasis that is put on the – again in a loose sense - *aesthetic* origin of our conviction that inflicting pain on another living being is morally wrong. 'So long', says Moreau, 'as *visible* or *audible* pain turns you *sick*, so long as *your own pains* drive you', your mind is not truly opened, not fit to realize that pain is really a little thing, not of much importance, not to be

considered as an obstacle to scientific research. Prendick himself acknowledges the truth of Moreau's remark when, earlier in the novel, he overhears Moreau vivisectioning a puma who cries horribly while being tortured. Prendick reflects: 'The crying sounded even louder out of doors. It was as if all the pain in the world had found a voice. Yet had I known such pain was in the next room, and had it been dumb, I believe – I have thought since – I could have stood it well enough. *It is when suffering finds a voice and sets our nerves quivering that this pity comes troubling us*' (1923, 44, my italics). Thus the abstract knowledge that some other creature suffers pain does not press itself on Prendick in such a way that he would object to it. It is only when he actually *perceives* the pain with his senses and gets bodily struck (nerves set quivering) by what he perceives that he cannot help but acknowledge the wrongness of the deed. This hints at how very much our moral evaluation of situations and actions depends on, first, to what extent we are capable of a *direct inspection* of these situations and actions,⁴ and second, whether what we thus inspect tends to make us sick with revulsion. 'We cannot really think that injustice is bad', writes Mary Midgley, 'if it does not at some point sicken us' (1981, 92).

But if this is true, then what makes us realize that an action or a situation is morally wrong is exactly what Rollin would classify as an *aesthetic* quality or reaction. In other words, the conviction that pain is an evil ultimately rests, as Stephen Clark has pointed out, 'only on our dislike of it' (1997, 65). In this respect it does not differ from the conviction that it is wrong to create animals with a reduced or even annihilated capacity for suffering, or hybrids of animals and human beings, or of different species in general. Hence the distinction between purely aesthetic concerns and purely moral concerns breaks down.⁵

4. The aesthetic dimension of ethics

To insist that, if our actions do not bring about any pain or suffering they cannot possibly be considered to be *morally* wrong, is simply question-begging. Even if we do not agree with someone who thinks that certain actions are wrong regardless of any suffering involved, this does not justify our assumption that their disapproval of these actions is not a moral one at all. There are quite a few people who feel that decerebrating a living being in order to experiment with it without violating the principle of conservation of welfare, is even worse than just hurting it and making it suffer. It is certainly arguable whether these people might not have the *wrong* kind of moral intuitions but that is not to say that their intuitions *are* not moral in the first place but rather 'aesthetic'. If there is a kind of disapproval that is really *purely* aesthetic it would be one where it is only the sight (or sound, or smell) of something which bothers us or is abhorrent to us. In the cases we have been considering, however, it is for most people the action itself, or the results of this action, to which they object. It is not that they don't want to *see* it done or don't want to think about it being done, but that they simply want it *not to be done*. Mark Packer (1996) has shown that many of our moral intuitions are aesthetically grounded in so far as their only justification is the vague but often strong feeling of shock or outrage we feel when we contemplate certain actions, and that those feelings are commonly regarded as being sufficient to justify a condemnation of these actions: 'We are inclined to disallow the actions in question because there is something, perhaps we cannot say, exactly what, that is simply offensive about them.' We need not be able to explain *why* it is offensive to us, but this does not make our judgement less valid. Consider the (fictional, but not entirely unlikely) case Packer describes:

Within the next few decades, genetic technology may develop to the point that it will be possible to manufacture sides of beef and chicken parts entirely from DNA, without the need to create or raise living animals. Potentially unlimited quantities of T-bone steaks and poultry wings might

emerge in just a few days or weeks out of dishes of genetic material that is extracted painlessly from individual cows and chickens. Since no animals would be killed and none would suffer any physical or psychological harm, this process may appear to provide moral license to even the most committed animal rights activist to eat meat. In the absence of suffering and death, there would seem to remain no sound ethical reason for vegetarianism.

But an animal rights activist might rejoin to this proposal with a rather arresting suggestion. If it were possible to manufacture animal meat this way, then human limbs and organs might also be brought to the deli counter by means of a similar process of DNA cultivation. Why should the food industry restrict its products just to non-human animals? Higher profits are possible with this expanded array of delicacies, as consumers might be willing to pay more in order to enjoy the naughty thrill of cannibalism without any pangs of conscience. Nobody would suffer any pain, and no one would be killed. In short, there would be no harm. (1996, 58)

I think Packer is quite right to conclude that most people would, in spite of the fact that there wouldn't be any harm done, vehemently resist such a proposal. If so, our reaction might be described as *aesthetic* insofar as it is caused 'by nothing more than characteristics that are entirely inherent in the behaviour or objects themselves' (1996, 61), but must also be described as *moral* insofar as it is an expression of the conviction that far more is at stake here than just a violation of individual taste. Instead we feel very strongly that this is something, whatever the consequences might be, that *ought not to be done*, period, not by ourselves and not by anybody else. This might appear irrational, since we cannot properly explain what's wrong with it, and a 'sane man' like Dr. Moreau would certainly feel quite differently about it, but then moral convictions are, after all, *always* non-rational, not in the sense of being contrary to reason but rather in the sense of being ultimately not based on or guided by any process of rational decision-making (Hauskeller 2001). This is *not* to say that our moral convictions are independent of our knowledge of the situation in question. *For us*, our moral convictions follow naturally from our beliefs concerning the facts of the situation, but others may have different moral convictions following just as naturally from the *same* factual beliefs, and to some – those who have a completely objective mind (if there is such a thing) – there does not follow anything at all in terms of moral value. Thus the entirely 'sane man' does not have any moral convictions at all. That is what those who demand a 'rational analysis of facts'⁶ fail to realize. The difficulty here is obviously how a 'rational analysis of facts' could ever decide the question of whether it is wrong or right to create transgenic animals. After all, there is no one way of 'ably presenting' the facts. Every presentation of facts is bound to be either coloured by the presenter's values, or to be neutral in respect to their moral evaluation. For this reason, if we do not *feel* strongly about the rightness or wrongness of an action or a state of affairs, we know at best what others *expect* us to do, and not to do, and hence what is generally *considered* to be right and wrong. In other words, it is not a moral, but a *social* knowledge we have, a knowledge of certain facts (namely what is valued by others) but not about what actually is valuable. David Hume thus had every reason to claim that all true morality depends on our sentiments (1739/40, III, II, V). If the sight or contemplation of an action or situation did not elicit any emotion in us, if we didn't feel any indignation or outrage whatsoever, then we wouldn't even *know* that there is something wrong with it. When those parts of our brain which support our emotions get damaged or functionally destroyed, we are still conscious of the moral norms considered to be valid in the society we live in but we no longer feel obligated by them. They have lost their binding force on us (Damasio 1994). To have an emotion always means to evaluate the situation one is in or is confronted with (Solomon 1980). Furthermore, having no emotion means *not* to evaluate the situation. It is tempting to argue that if someone says that she just can't *bear* the thought of humans eating human flesh even if it is especially created for that purpose and nobody is killed or hurt for it, or of animals being deliberately engineered without their cortex or

reduced in their capacity to suffer under the circumstances they are forced to live in, then she doesn't have a good reason for the claim that it is immoral to do so. If, however, we *can* bear the thought easily then what good reason could we possibly have to assert its immorality?

It all depends on what we mean when we say we cannot bear it. Maybe it is just that we cannot bear the sight of it in such a way that we tend to get sick when seeing it (or imagining it) without really minding the fact that it happens. Even though we felt nauseated we would then refuse to see more in it than a spontaneous, mechanical and entirely meaningless physical reaction to a visual stimulus. We would, instead of asking ourselves *why* we are nauseated and whether there might not be something in the object or situation we react to that justifies our nausea, just make sure that we are not exposed to the stimulus, and then don't bother ourselves with it anymore. Yet on the other hand, this might be a misunderstanding of what our revulsion really *means*,⁷ and we may discover at some stage that our revulsion in fact runs deeper, that what we are not able to bear is not really the *sight* (the sound, smell, or taste) or the *thought* of a certain state of affairs or course of action, but rather this very state or course of action *itself*, and as soon as we realize this we cannot any longer believe that our reaction is purely idiosyncratic, that we are just expressing a personal dislike which is of no relevance to others. Instead it will seem to us that the object in question *justifies* our negative reaction. We won't say then that it is *we* who cannot bear what's being done but rather that what's being done is in itself *unbearable*. Hence, we will expect or even *demand* from others to share our feelings of revulsion or abhorrence towards the action, just like Kant thought we would if making an *aesthetic* judgement. So the fact that there is something we simply cannot bear, may well be a perfectly good reason for saying that it is immoral to bring it about or to permit it to be.

Conclusion

In a well-known passage David Hume (1739/40, III, I, I) states the emotional basis of moral evaluation:

Take any action allow'd to be vicious: Wilful murder, for instance. Examine it in all lights, and see if you can find that matter of fact, or real existence, which you call vice. In which-ever way you take it, you find only certain passions, motives, volitions and thoughts. There is no other matter of fact in the case. The vice entirely escapes you, as long as you consider the object. You never can find it, till you turn your reflexion into your own breast, and find a sentiment of disapprobation, which arises in you, towards this action. Here is a matter of fact; but 'tis the object of feeling, not of reason. It lies in yourself, not in the object. So that when you pronounce any action or character to be vicious, you mean nothing, but that from the constitution of your nature you have a feeling or sentiment of blame from the contemplation of it.

Whereas I entirely agree with the first part of Hume's statement saying that our feelings alone reveal an action to be wrong, I disagree with the second part saying that in proclaiming something to be wrong we *mean* nothing but that we have a certain feeling of disapproval. Instead, what we mean is something quite different, namely that the action in question *is* wrong, and that we know it is wrong because we feel it to be wrong. And we do not think that our feeling it to be wrong *makes* it wrong. That is why we can in principle always do more than just state our feelings, that is: if we care enough to do so. We can reflect on our feelings and try to find out what it is that we spontaneously object to. We can, for instance, reflect on the fact that in treating animals along the lines of Rollin's argument we treat them as if they were mere tools or machines designed for human good and that it is this kind of *instrumentalisation* which is morally wrong (Holland 1990, 170).⁸ Or we can say that our attitude is not in agreement with the principle of *respect for nature* (Taylor 1986), that it

violates the animals' *telos* (Fox 1990, Hauskeller 2005), or their individual or species *integrity* (Rolston 2002, Verhoog 2002). These notions are not just feeble attempts to rationalize some vague 'aesthetic' dislike but, rather, serious efforts to *understand* our immediate emotional reactions, and to understand them in the light of the nature of, and our relation to, the world in which we live. They are expressions of the genuinely ethical refusal to let our moral world be governed by a convenient theory telling us what we are legitimately allowed to regard as wrong or bad and what not. Hence, when, for instance, Stephen Clark declares that there 'is something wrong, ugly, horrifying about the sight of a living creature made so wretched that it can no longer even care; there is something wrong, ugly, horrifying about mammals transformed into milk machines, or microcephalic lumps' (Clark 1997),⁹ then he is not just being queasy. For what he is really doing is articulating a common moral concern without pretending that such a concern has much to do with being what is thought to be rational or sane. And if a moral theory is not able to incorporate these concerns, this ought to be considered a weakness of the theory and not of the concerns many people actually have.¹⁰

Acknowledgements:

I am grateful to Nigel Pleasants and Christine Hauskeller for their helpful comments.

Notes

¹ The original is German and the translation is mine. Cf. Midgley 2000.

² Though the fact that the hybrids in this case are *human-animal* hybrids probably adds to Prendick's (and the reader's) horror, it is already the hybridization *as such* that is considered to be somehow wrong. When Moreau explains himself to Prendick he says: 'You begin to see that it is a possible thing to transplant tissue from one part of an animal to another, or from one animal to another, to alter its chemical reactions and methods of growth, to modify the articulations of its limbs, and indeed to change it in its most intimate structure? And yet this extraordinary branch of knowledge has never been sought as an end, and systematically, by modern investigators, until I took it up.' (1923, 90) Prendick asks him then 'why he had taken the human form as a model. There seemed to me then, and there still seems to me now, a strange wickedness in that choice. He confessed that he had chosen that form by chance. 'I might just as well have worked to form sheep into llamas, and llamas into sheep. (...)'' (91-92).

³ Cf. Clark (1994), p.13: 'Biotechnology is the art of manipulating living forms as though they were machines.' That they are in fact little more than machines is argued by, for instance, Leahy (1991) and Frey (1980). Thus it is claimed that if animals have needs then they have them in the same way as trees *or even machines* have needs (Leahy, pp. 44ff.; Frey, pp. 79 ff.), and the same holds for their suffering: 'Plants can suffer from too much sun or too little water, or a watch from rough handling.' (Leahy, p.223)

⁴ For an experimental proof of this see Milgram (1974). For a definition of the term 'aesthetic' cf. Eaton (2001), p.11: 'A is an aesthetic property of O (an object or event) if and only if A is an intrinsic property of O and A is culturally identified as a property worthy of attention (i.e. perception and reflection). (...) F is an intrinsic property of O if and only if direct inspection of O is a necessary condition for verifying the claim that O is F.'

⁵ As Eaton (2001) argues: 'A deep mistake has colored value theory. The mistake is believing in the general separability of the aesthetic and the moral.' (p.57). Eaton, however, addresses the issue in an (successful, I think) attempt to refute the dogma that the evaluation of art has to be free from all moral concerns. She does not say much about whether moral evaluation requires aesthetic considerations as well.

⁶ Murphy 1990, 13: 'Making animals that are a mixture of different species is not a new idea, but it is a new reality. (...) For many people, the transgenic manipulation of animals is a very frightening concept. (...) (However:) When we look at transgenic animals we must be aware that our reactions to them are unlikely to be based on the rational analysis of facts ably presented; the nightmare images from past fantasies are too likely to

escape from the Pandora's Box of our imagination and distort our vision.'

⁷ Cf. Gaita (2002), p.206 f.: 'Some people (...) become vegetarians because they find they cannot eat meat. At first they might not eat it for a practical reason (...) and then gradually they find the very thought of eating meat repulsive. It would be wrong to say of the vegetarians I am thinking of that they had just become squeamish, if that meant that their disgust was not a moral disgust. (...) But because they are likely to identify the moral basis of vegetarianism with decisions of principle, they might hesitate to say that theirs is a moral revulsion. And if, quite literally, they can't eat meat because the mere thought of it makes them nauseated, then they might be misled into believing that they cannot eat meat only because of the sheer strength of a 'merely' psychological revulsion against it. That would be a pity for (...) the impossibility they express is interdependent with a perception of what it means to eat an animal.'

⁸ '(...) genetic engineering essentially involves using living things as instruments. Moreover, it reduces them to instruments of a particular kind – mechanical ones. It does not simply view living things as instruments; it views them as *mechanisms*.'

⁹ Cf. Fox (1990), p. 41: 'Molecular farming – the incorporation of human genes into non-human beings for human benefit – is seen as a form of parasitism: genetic parasitism. It is as *abhorrent* to some people as the practice of grafting pig livers and chimpanzee hearts into humans.' (my italics).

¹⁰ Cf. Midgley (1985), p. 149: 'Today, this intellectualist bias is often expressed by calling the insights of common morality mere 'intuitions'. This is quite misleading, since it gives the impression that they have been reached without thought, and that there is, by contrast, a scientific solution somewhere else to which they ought to bow as there might be if we were contrasting commonsense 'intuitions' about the physical world with physics or astronomy. Even when they do not use that word, however, philosophers often manage to give the impression that whenever our moral views clash with any simple, convenient scheme, it is our *duty* to abandon them.' For instance, Grice (1967), pp. 146-7. Midgley comments: 'Grice demands that we withdraw our objections to harshness, in deference to theoretical consistency. But 'harsh' here does not mean just 'brisk and bracing' like cold baths and a plain diet. (...) It means unjust. (...) An ethical theory which, when consistently followed through, has iniquitous consequences is a bad theory and must be changed. Certainly we can ask whether these consequences really are iniquitous; but this question must be handled seriously. We cannot directly conclude that the consequences cease to stink the moment they are seen to follow from our theory.' (150)

References

- Blackburn, Simon (1998). *Ruling Passions. A Theory of Practical Reasoning*. Clarendon Press: Oxford.
- Bowring, Finn (2003). *Science, Seeds and Cyborgs. Biotechnology and the Appropriation of Life*. Verso: London.
- Brittnacher, Hans Richard (1994). *Ästhetik des Horrors*. Suhrkamp: Frankfurt.
- Clark, Stephen R.L. (1994). 'Modern Errors, Ancient Virtues', in Anthony Dyson and John Harris (ed.) *Ethics and Biotechnology*. Routledge: London, pp. 13-32.
- Clark, Stephen R.L. (1997). 'Natural Integrity and Biotechnology', in David S. Oderberg and Jacqueline A. Laing (ed.) *Human Lives. Critical Essays on Consequentialist Bioethics*. Macmillan Press: London, pp. 58-76.
- Damasio, A.R. (1994). *Descartes' Error. Emotion, Reason, and the Human Brain*. G.P. Putnam's Son's: New York.
- Eaton, Marcia Muelder (2001). *Merit, Aesthetic and Ethical*. Oxford University Press: Oxford.
- Fox, Michael W. (1990). 'Transgenic Animals. Ethical and Animal Welfare Concerns', in Wheale/McNally (eds.), *The Bio-Revolution. Cornucopia or Pandora's Box*. Pluto Press: London, pp. 31-54.
- Frey, R.G. (1980). *Interests and Rights: The Case Against Animals*. Clarendon Press, Oxford. Grice, G.R. (1967). *Grounds of Moral Judgment*. Cambridge University Press: Cambridge.
- Hauskeller, Michael (2001). *Versuch über die Grundlagen der Moral*. C.H. Beck: Munich.

Hauskeller, Michael (2005). 'Telos. The Revival of an Aristotelian Concept in Present Day Ethics', in *Inquiry* Vol. 48, No. 1, pp. 62-75.

Hume, David (1739/40). *A Treatise on Human Nature*. John Noon: London.

Gaita, Raimond (2002). *The Philosopher's Dog*. Routledge: London.

Gibbard, Alan (1990). *Wise Choices, Apt Feelings. A Theory of Normative Judgment*. Oxford University Press: Oxford.

Holland, Alan (1990). 'The Biotic Community: A Philosophical Critique of Genetic Engineering', in Peter Wheale/ Ruth McNally (eds.), *The Bio-Revolution. Cornucopia or Pandora's Box*. Pluto Press: London, pp. 166-188.

Leahy, Michael P.T. (1991). *Against Liberation. Putting Animals in Perspective*. Routledge: London/New York.

Midgley, Mary (1981). *Heart and Mind. The Varieties of Moral Experience*. Harvester Press: Brighton.

Midgley, Mary (1985). *Evolution as a Religion. Strange Hopes and Stranger Fears*. Methuen: London.

Midgley, Mary (2000). 'Biotechnology and Monstrosity. Why we should pay attention to the "yuk factor"', in *Hastings Center Report* 30, No. 5, pp. 7-15.

Milgram, Stanley (1974). *Obedience to Authority. An Experimental View*. Harper and Row: New York.

Murphy, C. (1990). 'Genetically Engineered Animals', in Peter Wheale/Ruth McNally (eds.) *The Bio-Revolution. Cornucopia or Pandora's Box*. Pluto Press: London, pp. 13-17.

Packer, Mark (1996). 'The Aesthetic Dimension of Ethics and Law: Some Reflections on Harmless Offence', *American Philosophical Quarterly*, Vol. 33, No. 1, pp. 57-74.

Rollin, Bernard (1995). *The Frankenstein Syndrome. Ethical and Social Issues in the Genetic Engineering of Animal*. Cambridge University Press: Cambridge, Mass.

Rolston, Holmes III. (2002). 'What do we mean by the intrinsic value and integrity of plants and animals?', in David Heaf and Johannes Wirz (eds.) *Genetic Engineering and the Intrinsic Value and Integrity of Animals and Plants*. Proceedings of a Workshop at the Royal Botanic Garden, Edinburgh, UK (Hafan, Llanystumdwy), pp. 5-10.

Rowlands, Mark (2000). *The Environmental Crisis. Understanding the Value of Nature*. Macmillan: London.

Solomon, Robert C. (1980). 'Emotions and Choice', in Amelie Oksenberg-Rorty (ed.) *Explaining Emotions*. University of California Press, Berkeley/ Los Angeles, pp. 251-281. Taylor, Paul W. (1986). *Respect for Nature. A Theory of Environmental Ethics*. Princeton University Press: Princeton.

Verhoog, Henk (2002). 'Does genetic engineering impact the intrinsic value and integrity of animals?', in David Heaf and Johannes Wirz (eds.) *Genetic Engineering and the Intrinsic Value and Integrity of Animals and Plants*. Proceedings of a Workshop at the Royal Botanic Garden, Edinburgh, UK (Hafan, Llanystumdwy), pp. 35-39.

Wells, H.G. (1921). *The Island of Dr. Moreau*. Heinemann: London.

Introduction to Book Symposium on Bernard Gert's *Common Morality: Deciding What To Do* (New York: Oxford University Press, 2004.)

Dean Cocking
Charles Sturt University

In 2004, the Centre for Applied Philosophy and Public Ethics was fortunate to have Professor Bernard Gert, from Dartmouth College (Hanover, New Hampshire, U.S.) as a visiting fellow. During his visit Professor Gert was an especially lively and thought provoking contributor to the intellectual life of the Centre and to our seminar program at CAPPE in particular.¹ Many of us at CAPPE benefited significantly from his interest and attention to our own work, and in my case, I would especially like to thank Bernie for this, and for his (very) generous claims regarding contributions to some of his recent work. In September 2004, CAPPE conducted a workshop on Professor Gert's most recent book *Common Morality: Deciding What To Do*. The papers, and replies from Gert, presented at that workshop make up the symposium here.

Professor Gert's grand moral theory, first published in 1970,² is strikingly original - in particular, his accounts of the concepts of rationality, impartiality and morality, and of relations between them - and has received significant attention.³ *Common Morality* is developed in light of Gert's work in the area over the past forty years, and aims to build a theoretical account of morality based on our commonly shared moral practises and related substantive values. Since the theoretical account in *Common Morality* is concerned to identify, describe and classify the phenomena of our common moral practices and related substantive values, one important question to ask will be how well it does so.

Gert provides, for instance, a list of ten moral rules, which, he says, taken together, and given appropriate interpretation,⁴ exhaust the moral territory of what is morally prohibited and required. These rules are: Do not kill; Do not cause pain; Do not disable; Do not deprive of freedom; Do not deprive of pleasure; Do not deceive; Keep your promises; Do not cheat; Obey the Law, and; Do your duty. Do these ten moral rules then, really exhaust our moral prohibitions and requirements? Are some of these rules redundant?

In their paper, 'Common Morality and 'Institutionalising' Ethics', Andrew Alexandra and Seumas Miller argue that Gert's ten rules do not exhaust our moral prohibitions and requirements, since moral rules, such as against theft, damage to property, fraud and defamation are not accounted for by these ten rules. On the other hand, they claim some of Gert's rules, specifically 'Obey the law' and 'Do your duty', are redundant. And in 'Applying our Common Morality: The Case of Privacy', Jeroen van den Hoven suggests that, while Gert's systemisation of the moral rules seems to work well for our inter-personal morality, in particular the case of privacy, it might be insufficient for application to ethical issues arising from our emerging technologies, in particular our computer-related technologies.

Another key feature of Gert's account is that he understands his project of identifying, describing and classifying the key concepts that make up our 'common morality', in particular the concepts of impartiality, rationality and morality, as the concern to explain, rather than revise, "their central and coherent use by thoughtful people in everyday life".⁵ It

is, according to Gert, the coherent use of these concepts by ordinary thinking people, rather than the understanding of these concepts provided by philosophers, that should be our guide.

An important question to ask here then, will concern the extent, or ways in which, we should agree with Gert regarding his views about the accounts from philosophers of such concepts, and the primacy he gives to moral practices over the moral theorising of philosophers. Has he got the philosophical accounts right? Should we reject the revisionist approaches from philosophers regarding our moral practices to the extent, or in the ways, Gert claims? (Conversely, should we revise our theoretical accounts in light of our moral practises to the extent, or in the ways, Gert claims?)

One central point Gert wants to make against the moral theories of other philosophers, such as Immanuel Kant and John Rawls, is that these theorists understand morality as providing one unique, right answer to all moral questions. Gert thinks this false, i.e., he thinks there is often not just one 'right' answer to a given moral question. Moreover, he thinks the claim to the contrary presents a view of morality that should be regarded as a form of moral arrogance.⁶ In his paper, 'Timeless Wisdom or Moral Arrogance?', Thomas Pogge argues Gert has got both Kant and Rawls wrong, i.e., neither are committed to the claim that morality provides a uniquely correct answer to all moral questions. Indeed, Pogge argues that the relevant theories of both Kant and Rawls clearly allow that morality does not prescribe that there is only one correct answer to all moral questions, i.e., that there may be conflicting but equally morally acceptable answers to moral questions. And in my paper, 'Moral Arrogance and Moral Disagreements', I argue, first, that Gert has misdiagnosed the problem of moral arrogance, and second, that he has not given us the right account of how we should understand the sources of our moral disagreements.

In response to the problems put by the contributors to this symposium, Professor Gert has provided careful, detailed and comprehensive replies, and in doing so has contributed various fresh developments to his own views. As mentioned, his earlier work has by now received significant attention. However, 'Common Morality' is very recent and this symposium represents the first large discussion of the work. We are therefore, fortunate and grateful for the attention, time and effort Gert has provided us. Moreover, as all the contributors recognise, 'Common Morality' is a significant, original and comprehensive contribution to the understanding of moral theory and practice, and of relations between the two.

Notes

¹ His paper 'Moral Arrogance and Moral Theories' developed and read at CAPPE is now forthcoming in *Nous*, Supplementary Volume 15 on Applied Ethics, 2005.

² Most notably, see: Bernard Gert, *The Moral Rules: A New Rational Foundation for Morality*, (Harper and Row, 1970), and Bernard Gert, *Morality: Its Nature and Justification* (Oxford University Press, 1998).

³ Most notably here, see: *Rationality, Rules and Ideals*, Walter Sinnott-Armstrong and Robert Audi (eds) (Rowman and Littlefield, U.S., 2002), and; 'Book Symposium', *Philosophy and Phenomenological Research*, 62, pp. 421-81, 2001.

⁴ For Gert's account of how to interpret the ten moral rules, see *Common Morality: Deciding What To Do*, (Oxford University Press, New York, 2004) pp. 29-52.

⁵ Bernard Gert, 'Precis of *Morality: Its Nature and Justification*', in *Rationality, Rules and Ideals*, Ibid, p. 5.

⁶ See, 'Moral Arrogance and Moral Theories' Op cit.

Common Morality and ‘Institutionalising’ Ethics

Andrew Alexandra and Seumas Miller

University of Melbourne, Australian National University

1. Introduction: Moral theory and Moral Practice

In *Common Morality* Bernard Gert understands morality functionally: “The point of morality is to lessen the suffering of those harms that all rational persons want to avoid: death, pain, disability, loss of freedom and loss of pleasure” (4).

In fact, according to Gert we have a system in place which governs our actions and decisions in order to achieve that function – what he calls ‘common morality’. Specifically, “it is a system that rational persons put forward as a public guide for the behaviour of everyone who can understand it and guide their behaviour by it, that is, moral agents” (4)

As we understand his view, Gert is not claiming to offer a new understanding of the content of the moral system, though he clearly thinks that he is developing a new theoretical account of that system. Gert conceives of the relationship between moral practice and moral theory by analogy to that which holds between linguistic practice and grammar. Linguistic practice has a kind of priority over grammatical theory in at least two ways. Firstly, people can be competent speakers of a language, while having no explicit grasp of its grammar. Secondly, linguistic practice provides the standard against which grammar is measured – where there is a difference between what a grammatical theory says counts as correct use of language and what is generally accepted as correct, it is grammatical theory that should be revised. Similarly, Gert points out, it must be the case that just about everyone is morally competent, at least in the sense of knowing what the appropriate moral judgment is in most cases, and where there is a conflict between what moral theory claims is correct and what most morally competent persons believe, this shows an inadequacy in the moral theory, not in the beliefs of the persons.

This understanding of the relationship between moral theory and practice contrasts with what are perhaps the dominant philosophical approaches, since the latter have denied the priority of moral practice. On some accounts, the relationship between ordinary moral practice and moral theory is like that between what might be called the ‘common sense’ view of the material world and modern physics. Physics arose from that common sense view, but modern physics tells us that at various points the common sense view is deeply mistaken. However, those mistakes typically have little practical effect – we are adept at dealing with our physical environment despite the mistaken beliefs we have about it. Similarly, some ‘two-tier’ accounts of morality take it that although most people have false beliefs about morality - holding, for example, to absolute principles - the practical effects of those mistaken beliefs are not bad – indeed may be good – since they lead to our acting in ways that are, by and large, correct, or at least the best that can be hoped for from imperfect beings like us. Other revisionist accounts agree that most people have false beliefs about morality, but are not so sanguine about the effects of those beliefs. (Perhaps a relevant analogy here would be the relationship between certain traditional beliefs about disease, which see it as caused by magic, and modern medical theory. The traditional beliefs are not only false by the lights of modern theory, but also positively harmful, since they stand in the way of understanding how to avoid or treat disease) For example, if the possession of sentience is taken as the qualifying condition for moral standing as some utilitarians hold, then clearly we are unwittingly

committing many moral wrongs in our treatment of animals. We need to radically revise our moral judgments and practices in the light of moral theory.

Even if the priority of moral practice over theory is accepted, moral theory remains of practical, and not merely theoretical, importance. Just as a grasp of grammar can help us to become more self-conscious about the structure of language, and hence allow us to use it more effectively to communicate with each other, so attention to moral theory can inform our moral judgments, both individually and collectively. Understanding moral theory can lead a person to see, for example, that they have been acting in a way that is incompatible with a moral precept which they accept. It can help people who disagree in some matter of moral significance to locate the source of that disagreement and thus, perhaps, make it more amenable to resolution. It can help us understand why in some cases we should tolerate a range of different and conflicting views about a matter of moral significance. And since moral judgments do not wear their status on their face, an adequate moral theory may help us distinguish between claims that have moral weight and those which do not.

One test for any Gertsterian-style moral theory, then, will be the extent to which it is adequate to the moral phenomena which it is its task to systematise and explain. We think that Gert's own theory comes out of this test rather well by comparison with others. However, perhaps inevitably, we also believe that there are points at which it needs to be clarified, supplemented, or altered, in order to fit with common moral practice.

2. The Moral Rules and the Moral Ideals

The fundamental categorical distinction which Gert discerns within the system of 'common morality' is between what he calls 'the moral rules', and the 'moral ideals'. The common feature of the rules is that they prohibit the causing of harm.

According to Gert there are ten moral rules, which fall into two groups (15). The first five simply reflect the five basic harms rational persons want to avoid. The first five moral rules are: Do not kill; Do not cause pain; Do not disable; Do not deprive of freedom; Do not deprive of pleasure.

The second five rules prohibit those kinds of actions that *indirectly* cause the same five harms (17). They are: Do not deceive; Keep your promises; Do not cheat; Obey the law; Do your duty.

While we follow the moral rules simply by not harming, we have to engage in positive action to conform to the moral ideals. 'The moral ideals encourage people to prevent or relieve the harms that the moral rules prohibit them from causing.' (15)

The difference between the rules and the ideals include:

1. *Deontic standing*: The ten moral rules 'account for all of the kinds of actions that are morally prohibited and required' (15). It is morally obligatory to take (or refrain from taking) an action if and only if it falls under one of the moral rules. Actions in accordance with the moral ideals, on the other hand, are not morally obligatory, though they are of course permissible, and indeed are to be encouraged.
2. *Justification*: Violations, or attempted violations, of the moral rules may be justified, but such violations require excuse or justification. On the other hand 'there is usually no need to be concerned with justifying failing to follow a moral ideal' (17).

3. *Punishment/Reward:* Consequently, a person who violates, or attempts to violate, a moral rule is rendered liable for punishment, while someone who does not live up to a moral ideal is not rendered liable for punishment (even if they do not live up to *any* moral ideal); on the other hand, that a person conforms to the moral rules is not a reason to reward them (since they are just acting as they are required to act), while the fact that they have acted in accordance with a moral ideal does provide such a reason.

These differences themselves appear to stem from the fact that the moral rules are more basic, in the sense of more necessary to human well-being, than the moral ideals: “it is considered more important for people to obey the moral rules than to follow the moral ideals. . .” (23)

But in addition Gert thinks that these moral rules can only be concerned with not harming, as opposed to assisting. Gert quotes Mill ‘. . . a person may possibly not need the benefit of others, but he always needs that they do not do him hurt.’ (17). This is an empirical claim, and an unproven one at that. Recall in this context that “The point of morality is to lessen the suffering of those harms that all rational persons want to avoid: death, pain, disability, loss of freedom and loss of pleasure.” Reduction in the suffering from these harms is very often achieved by assisting others, as well as by not inflicting the harms on them. At any rate, our general point here – and one to which we will return later on – is that the distinction between the moral rules, i.e. the obligatory, and the moral ideals, i.e. the merely desirable, does not track the harm/assist distinction. Specifically, there can be moral obligations to assist others. Indeed, Gert himself acknowledges at least one such case, namely, your duty to assist a needy person, given that you are the only person able to do so and doing so is not unduly costly to you. (‘Do your duty’ – Rule 10).

3. The Second Five Moral Rules

To recapitulate, according to Gert there are ten moral rules, which fall into two groups. The first five moral rules are: Do not kill; Do not cause pain; Do not disable; Do not deprive of freedom; Do not deprive of pleasure. The second five rules are: Do not deceive; Keep your promises; Do not cheat; Obey the law; Do your duty.

According to Gert, “Although widespread violation of the second five rules always results in an increase in the amount of harm suffered, a particular violation of the second five rules does not always result in anyone suffering some harm.” (17) Here it is important to note that by Gert’s lights deceiving someone, breaking your promise etc. are not in themselves harms – or at least they do not harm persons. So causing a false belief (deception) or a false expectation (insincere promise) are not harms as such. Further, breaking one’s promise, disobeying the law on some occasion or failing to do one’s duty on some occasion do not necessarily harm anyone. On the other hand, such actions presumably constitute *institutional* harms. Thus, in so far as the institution of promising consists of many promises being made and kept, then even a one-off, causally self-contained breach of a promise constitutes an institutional harm, notwithstanding that it has no harmful effect on any other acts of promise-making or promise-keeping. Moreover, arguably truth, or at least moral truth, is a moral value and falsity a disvalue. If so, then to cause someone to have a false belief may well be to harm them. Further, asserting what one knows to be false, breaking one’s promise and cheating are all actions that involve a breach of trust, and therefore – it might be argued – damage or harm the relationship of trust between individuals.

At least three of the second set of five moral rules, namely ‘Keep your promises’, ‘Do not cheat’ and ‘Do your duty’, are conditional in nature, and the antecedent is a voluntary action which is in itself not morally obligatory. One does not have to make a promise, but if one does, this rule tells us, one is obliged to keep it. One does not have to take part in a competitive game, but if one does, one should play according to its rules. And so on.

4. Three Additional Moral Rules - ‘Do not steal or damage other people’s property’, ‘Do not defraud’ and ‘Do not defame’

Let us turn to those moral rules that prohibit actions that would otherwise indirectly cause harms. Perhaps the two most obvious omissions from the list are ‘Do not steal or damage other people’s property’ and ‘Do not defraud’. Theft and fraud certainly indirectly harm people, notably by depriving them of resources by means to live and by means to exercise their freedom.

A third apparent omission is ‘Do not defame’. Defamation is harming someone’s reputation. Certainly, harming someone’s reputation indirectly causes harm, e.g. loss of economic opportunities and therefore resources to freely act. Arguably, harming a reputation is not in itself a bad thing, although undeserved destruction of a reputation appears to be bad in itself. So perhaps harming reputations ought to be regarded as a sixth basic harm alongside do not kill, disable, deprive of freedom or pleasure, and do not cause pain.

It might be thought that theft and/or fraud and/or defamation are reducible to one or more of the other moral rules, specifically the rules not to deceive, to keep one’s promises and/or not to cheat. But if I openly walk into your house and remove your household appliances, you will know all that has happened and by whom; so there is no deception. In addition, I might not have promised not to take your goods; so there is no breach of promise. Finally, it is difficult to see in what sense, if any, I have cheated you. Certainly, theft of this kind is not akin to (say) insider trading. In the latter case, but not the former, there is a competitive market situation governed by rules, and insider trading is breaking one of those rules to give someone an unfair competitive advantage. So theft does not necessarily involve deceit, the breaking of any promise or cheating.

Naturally, theft typically involves breaking the law. But what is called for is an explanation for the common existence of laws against theft. There is at least one such explanation available, namely that theft causes indirect harm, notably to freedom. However, the question remains as to why theft is not an additional moral rule within the second set of rules.

It might be argued that theft is necessarily a legal offence; there is no such thing as theft outside a legal framework of property rights. If so, then the moral rule against theft is redundant within Gert’s scheme of moral rules, since there is a moral rule to obey the law. This argument is invalid if there is in fact no general rule to obey the law, or if theft is not necessarily a legal concept.

Arguably, people have a moral right to resources, living areas and artefacts, i.e. a right to property, and this right is not logically dependent on law. For example, it seems coherent to claim that a community has a moral right to exclusive use of a territorial area which they have long occupied or that an individual has a moral right to his personal property, e.g. his clothes. Again, absent any prior agreement, an artist has ownership rights in relation to the painting that he has laboured to create, or a philosopher in the distinctive moral theory that he

has working on for many years. Accordingly, other things being equal, it is an act of theft when land is forcibly taken from a people who have been the sole occupants of it for thousands of years, or when one's shirt is ripped from one's back, or when the artist's masterpiece is taken without his consent, and hung in someone else's house. Moreover, these would be instances of theft, even if there were no applicable law (as there well might not be, at least in the case of a people's territory or in the case of (say) unauthorized reproduction of the image in the artist's painting, or unattributed renditions of the philosopher's theory). Furthermore, the destruction or damaging of someone else's property is a further violation of a moral right. Suppose a people's territory is not only stolen, but also made uninhabitable by (say) dumping nuclear waste on it. Or suppose a thief not only take the artist's painting, but also destroys or defaces it.

Now consider fraud. Fraud necessarily involves deception. But fraud also involves taking a benefit from someone that one is not entitled to. Accordingly, fraud is not reducible to deception. Moreover, while some frauds involve breaking promises, many do not. Consider a fraudster who pretends to be a blind, rich, old woman's long lost son. The old lady gives him an expensive present – a Rolls Royce car. No promise has been made or broken. Nor is the case one in which the fraudster has broken some rule in a competitive context, and thereby gained an unfair advantage; it is not a case of cheating. So fraud does necessarily involve deceit, but it involves more than deceit. Moreover, fraud does not necessarily involve breaking a promise or cheating.

Naturally, as is the case with theft, fraud typically involves breaking the law, and there is a need to explain why this is so. There is at least one explanation available; as is the case with theft, fraud causes indirect harm, notably to one's freedom. However, the question remains as to why fraud is not an additional moral rule within the second set of rules.

It might be argued that fraud is reducible to deception and theft; it is theft by deception. However, fraud is not necessarily theft. Consider the above-described example again. The old lady has freely given the Rolls Royce to the fraudster; he did not steal it. Moreover, fraud does not necessarily involve relieving someone of their property. Suppose the twin sister of a famous pop star pretended to be her pop star sister in order to receive adulation from fans at a function celebrating the pop star's achievements. Or suppose the slightly drunk pop star's boyfriend engaged in sexual intercourse with the sister, believing it to be the pop star herself. This is fraud, but not theft of property, nor is it necessarily illegal.

It has been objected by Gert (in conversation) that since fraud necessarily involves deception, and there is already a moral rule against deception, the putative moral rule against fraud is not required and is therefore redundant.

There are at least two problems with this objection. First, we note that Gert admits the rule against cheating among the second set of moral rules. But surely cheating necessarily involves deception. Gert himself argues that it does not. He provides the example of a boss who always wins at golf because he breaks the rules; however, the boss breaks the rules quite openly, and no one protests since he is the boss. Contra Gert, this appears not to be a case of cheating at golf.

Here two possibilities are salient. Perhaps this is a special case of playing golf with an unofficial handicap. This would be so, if there is a tacit agreement between the boss and the other players that he is allowed to break *some* rules to assist his prospects of success, but the players still try to win within an overall framework of compliance with most of the rules; so

there is a genuine competition and it is quite possible for the boss to lose, notwithstanding the advantage he has. Accordingly, the boss is not cheating, rather he is allowed to break some rules in accordance with the informal provision of a handicap. This is common enough in, for example, family games of cricket where younger members of the family will not be given out, no matter what, until they have scored a run. The second possibility is that the so-called round of golf has ceased to be an instance of a competitive game governed by rules. This would be so if the boss was allowed to break pretty much any rule he liked to ensure that he won. But if so, the boss is not cheating, since there is no actual competition governed by rules, but at best the pretence of one.

Let us now consider the second problem with Gert's claim that a moral rule against fraud would be redundant. Why should the fact that breaking one putative moral rule necessarily involves breaking another moral rule disqualify the first moral rule from being a moral rule? Let us assume that psychological pain is involved in all infringements of the rule against depriving someone of freedom; loss of freedom always brings with it psychological pain. Would that render the rule against deprivation of freedom redundant? Surely not, since the psychological pain accompanying the loss of freedom does not go to the heart of what is morally unacceptable about being deprived of one's freedom. Again, if we were to accept this argument, then the moral rule against killing would be redundant, since all instances of killing are instances of disabling, deprivation of pleasure and deprivation of freedom. But once again, what is wrong with killing is not simply that it brings with it an inability to walk or a loss of freedom or a deprivation of pleasure. We conclude that if fraud is not reducible to the other acts prohibited by the moral rules then there is a case for having a separate moral rule against fraud.

Finally, let us turn to defamation. It might be thought that defamation necessarily involves deceit, but this seems to be false. If I make known some actual unsavoury but essentially irrelevant episode in the distant past of a public figure this may destroy his or her reputation, but undeservedly so. This is not deceit but it is defamation. Moreover, as we saw with the case of fraud, the fact that defamation may involve deceit is not necessarily a reason for there not being a moral rule proscribing it; since deception is not all or even the most important moral element in defamation. The most morally important element in defamation is reputation – defamation being the undeserved destruction or diminution of reputation. Not does defamation necessarily involve the causing of pain or embarrassment, since the defamed person may never discover that he or she has been defamed. Indeed, it might be impossible for them to find out – they might, for example, be dead.

5. Two Rules Too Many – ‘Obey the Law’ and ‘Do your duty’

By contrast with Gert's conception of the moral rule to obey law, we offer the following view. There is a moral obligation to obey *specific* laws and *specific* legal systems, but only because those laws/legal systems embody the moral rules and/or achieve collective goods not otherwise obtainable. On this account legal systems or laws as such do not generate moral obligations, even presumptive moral obligations that can be overridden. So the obligation to obey the law is entirely unlike to obligation to keep one's promises. Other things being equal, making a promise creates a moral obligation. Naturally, some promises – such as to kill innocent people – do not create obligations, and some promises that do create moral obligations can be overridden. However, other things being equal, the fact that there is an extant legal system prescribing a particular set of acts and omissions does not entail that there is an obligation to obey those laws; rather it all depends on the laws in question. Consider a totalitarian state of the kind described in Orwell's *1984*. Surely there is no

obligation to obey the laws in such a set up. We conclude that there is no fundamental moral rule to the effect that one ought to obey the law.

What of Gert's moral rule, 'Do your duty'? The duties Gert has in mind are those associated with specific institutional roles, such as that of doctor or police officer. Now there is an ambiguity in relation to the term "duty". It could mean "institutional duty" or "moral duty". Surely the duties in question are institutional duties, as opposed to moral ones. For if they were moral duties then this would be a meta-moral rule to the effect that one ought to perform those actions prescribed by the moral rules, whatever those actions might be. We conclude that the moral rule ought to be understood as the prescription 'Do your *institutional* duty', though with one important caveat to be discussed shortly.

However - as was the case with legal duties - institutional duties are only morally obligatory in so far as they embody moral rules; there is no moral requirement to, for example, participate in the institution of slavery, whether as a slave or slave owner. Quite the contrary; slaves and slave-owners have a moral 'duty' or obligation to abandon or subvert their roles, if it is possible to do so. Whether or not the institutional duties definitive of some institution or institutional role morally ought to be performed depends on what that institution or institutional role is; so it appears that there is no fundamental moral obligation to do one's institutional duties. So we seem forced to the conclusion that there is no moral rule to the effect that one ought to do one's duty. We will return to this issue below.

6. The Moral Duty to Aid

Amongst Gert's duties there is one that is not an institutional role duty, namely a duty to assist those in need when one is the only person able to help and can do so at little cost. Gert is surely right to hold that there is such a moral duty. Moreover, it is presumably a duty to assist in relation all the basic five harms.

Of course Gert postulates a further and much larger set of institutional duties attached to a wide variety of institutional roles; the institutional duties in question are also moral duties, indeed positive moral duties, i.e. duties to assist. An important question that now arises for Gert concerns the moral basis for the existence of such moral (institutional) duties. Does a fire crew have a positive moral duty to put out an incipient grass fire in the bush which – given the very high temperature and strong winds – will, if not extinguished, develop into a full-scale bush fire that will destroy property, and perhaps even lead to loss of life? By the lights of Gert's theory, the members of the fire crew, and indeed anyone else, would be doing a good thing if they put out the grass fire; they would be acting in pursuance of a moral ideal. (Here we are presuming that the fire has not been deliberately lit, but was caused by some natural event.) However, unlike ordinary members of the public, the members of the fire crew are morally *required* to extinguish the fire; it is their morally required institutional positive duty to extinguish the fire. So now the question to be asked of Gert is, What is the moral basis for this institutional duty?

It might be thought that the institutional duty is based on a promise. Each member of the fire crew has made a promise to put out fires (a promise he has presumably made on the condition that he will be suitably remunerated). However, this cannot be Gert's view, since this would reduce most moral duties to promises, and on his account keeping one's promises and doing one's moral duties are fundamental and separate moral rules.

Gert's own view seems to be that institutional duties are consequent on the occupancy of an institutional role.¹ Once a person actually takes on a role to (say) assist those who are ill, reasonable expectations are generated on the part of the ill that their illnesses will be treated. As a consequence a doctor comes to have a duty to treat those who are ill. This suggestion is akin to the view that institutional duties are based on promises: in both cases the duties are generated by an action that is not itself obligatory. However, the notion of intentionally creating expectations that one will assist does not seem sufficient to generate a moral obligation to assist, given there is no prior moral obligation to assist. Consider a winemaker. The winemaker has no moral obligation to make and sell wine to a community, but in fact the winemaker does make and sell wine to the community. Accordingly, the members of the community reasonably assume that they will continue to be able to procure wine from the winemaker, albeit at a price. Now assume that the winemaker decides that while he will continue to make wine, but in smaller quantities and only for selected customers – customers who understand and fully appreciate wine, as distinct from ordinary drinkers. Accordingly, the reasonable expectations of many former customers - the non-connoisseurs - are dashed. Moreover, since there is no other winemaker to which they have access, these customers suffer various harms, notably deprivation of pleasure. Notwithstanding the deprivation of pleasure caused by his actions, surely the winemaker does not have a *duty* to make and sell wine to ordinary drinkers. Certainly, given his limited stocks of wine, he is discriminating against the ordinary drinkers in favour of the connoisseurs, but he this discrimination is not arbitrary; indeed it has a certain rational justifiability. We conclude that the mere creation by a role occupant of reasonable expectations that he will provide benefits does not generate a moral duty to so provide those benefits. In the case of our finicky winemaker surely no-one, not even his disappointed former customers, would claim that he was liable for punishment on account of his actions. But if he is not liable for punishment, he has not, on Gert's account, broken the moral rule to do his duty. We conclude that moral institutional duties to assist cannot be grounded on reasonable expectations generated by role occupancy.

There is a further important point to be made in relation to such attempts to ground moral institutional duties to assist on the basis of the individual promises or the reasonable expectations generated by role occupancy. Even if some such account might be adequate as an explanation as to why some individual role occupant, such as a fireman, has a duty to (say) put out fires, it does not provide a moral basis for the institutional role itself; so it does not provide a moral basis for fire brigades and the associated institutional role of fireman. Specifically, it does not acknowledge that there is a moral requirement to see to it that there are fire brigades and firemen to extinguish fires in order to avert the very serious harms consequent upon uncontained fires. Indeed it seems that it cannot do so. On such accounts, if no-one chooses to assume the role of fireman, or in some other way generate reasonable expectations that they can be relied to put out fires, and no-one else chooses to put out fires, then no one has done a moral wrong and, in particular, no-one has failed to act in accordance with their moral duty. We believe, on the contrary that there is a moral duty to see that fires are put out, a moral duty which was not being discharged in this case.

The upshot of our discussion is that there appears to be a serious lacuna in Gert's theory, namely the absence of the provision of a moral basis for the creation of institutions and the consequent institutional moral duties to assist of members of those institutions. In what remains of this paper we will attempt to fill this lacuna.

We would argue that in addition to the duty of an individual to assist where he or she is the only one able to do so and there is minimal cost involved, there are also duties or

responsibilities of the members of groups to so assist; in short there are collective responsibilities. Suppose in our above bushfire scenario there are no firemen in the vicinity. Indeed, the only people in the areas are a group of bushwalkers from a well-known philosophy department in New Hampshire. No single member of the group could put out the small grass fire, but if they act jointly they can do so without risk to themselves, or indeed any great inconvenience. Moreover, they know that there is no way for them to warn the local community of the impending danger, should they refrain from extinguishing the fire. Presumably – given their awareness of the near certainty of a conflagration if they do not intervene - the members of the group have a collective moral responsibility to do so. To see this, let us assume that the group, under the sway of a slightly faulty, if easily fixed, moral theory, decides that they have no such obligation, and in the spirit of savouring the raw power of nature in its elemental outback state, seeks a suitable vantage point from which to view the inevitable spectacular firestorm with its attendant large-scale destruction of life and property.

Notwithstanding our sympathy with the desire of our erstwhile American friends to observe, understand and analyse the unique phenomenon that is the Australian bushfire, we *would* be inclined to blame these philosophers for failing to do what we believe they were morally obliged to do. We note that in some jurisdictions in the world the failure of individuals or groups to assist in these sorts of circumstance renders those individuals and groups liable for punishment. In the case of our philosophical bushwalkers, we strongly recommend (at the very least) immediate cancellation of their planned seminars and tourist travel itineraries.

So the need of groups for various forms of assistance that can only be adequately rendered by other groups generates collective responsibilities on the part of groups to so assist. Moreover, where such collective responsibilities can most effectively be discharged by establishing institutions and institutional roles whose institutional duties consist of providing such aid, e.g. fire brigades and firemen, then members of the group who have the collective responsibility have a derivative responsibility to establish and support such institutions.

Further, members of a given group may have collective moral responsibilities towards the membership of that very group, i.e. the group of which they are members. Assume that there is a high probability of bushfires destroying the property and taking the life of some very small percentage of Australia's population, but that with respect to any individual Australian the risk is close to zero. In this situation all or most Australians have a collective moral responsibility to prevent bushfires, and that responsibility can be discharged at minimal cost and inconvenience to any individual Australian. However, from the perspective of narrow individual self-interest, each Australian would not contribute. In the first place, given the almost zero possibility of harm to his or her self-interest, each does not have an incentive to contribute to a cooperative fire prevention scheme. And in the second place, even if such a cooperative scheme existed, self interest would dictate that each putative participant free-ride. We conclude that the fire prevention activity in question is undertaken in large part as a consequence of a perceived collective moral responsibility, and might initially take the form of local volunteer fire prevention associations. In due course, a division of labour tends to evolve, and the institution of professional fire brigades is established to discharge the collective moral responsibility to avoid the harms caused by fires.

A further point to be made here is that collective moral responsibilities can hold not only within groups, but also between groups. Consider cross-generational responsibilities. The current generation of Australians have a collective moral responsibility to address problems of climate change that if left unaddressed would only harm future generations.

Here, we'll assume, there are no benefits and only costs to the current generation of Australians. Moreover, the harms in question need not have been caused by the current generation who bear the costs, e.g. the harm that will arrive as a consequence of the large-scale land clearing in Australia by previous generations. Once again, this collective moral responsibility is most effectively discharged through the establishment of appropriate institutions and associated roles, e.g. the Murray Darling Commission.

On this conception the duty to assist is wider than Gert allows, though our widening of this duty takes as its starting point Gert's narrow notion of an individual duty to assist another individual. Furthermore, this broader duty to assist may, in certain cases, imply the duty to establish and support institutions to achieve the object of the duty. Once such institutions with their specialised role holders are in place it may be that we generally have no further duty to assist within the area of the institutions' operations. Indeed, it may be that generally we should not even *try* to assist, given our relative lack of expertise and the likelihood that we will get in the way of the role holders. Still, on occasions when no role holder is available to assist and we possess relevant capacities etc. the individual or collective duty becomes re-animated.

Since the collective moral responsibility or, if you like, duty to assist, is a duty in relation to the avoidance of the basic harms, arguably it should not be seen as a separate single duty with respect to multiple types of harm, but rather an extension to our understanding of the five basic harms. Accordingly, our duty in relation to life is not simply not to *individually* kill but also to *collectively* preserve life. Similarly for the deprivation of freedom and the other harms (including our suggested additional harms). These collective responsibilities are embodied in institutions and institutional roles, and justify the establishment and maintenance of such institutions and roles.

Notes

¹ Provided in conversation.

Reply to Andrew Alexandra and Seumas Miller: “*Common Morality* and ‘Institutionalising’ Ethics”

Bernard Gert
Dartmouth College

Andrew Alexandra and Seumas Miller present a clear and correct account of what I am trying to do in my moral theory. They appreciate that I am attempting to provide an account of our common morality, not to revise or improve it. They also appreciate the primacy of practice over moral theory, but they do not seem to fully appreciate that morality must be such that all moral agents, those about whom moral judgments are legitimately made, must know what kinds of actions morality prohibits, requires, discourages, encourages, and allows. We cannot make a moral judgment about an action of a person who is legitimately ignorant that what he did was immoral. This is one important way in which morality differs from law and religion, for both religion and law allow judgments to be made about the actions of a person who is legitimately ignorant that his action was prohibited by his religion or legal system. I mention this because in the first sentence of their paper, Andrew and Seumas say, “In *Common Morality* Bernard Gert understands morality functionally.” It is true that I do hold, “The point of morality is to lessen the suffering of those harms that all rational persons want to avoid: death, pain, disability, loss of freedom and loss of pleasure.” But this is also the point of medicine. Morality cannot be understood simply by providing its function; it is also necessary to provide its content and form.

Andrew and Seumas use some terms, e.g., “institutional harm” that I am not completely clear about. Insofar as there is something appropriately called “institutional harm,” it is morally relevant only insofar as it significantly increases the risk of individuals suffering one or more of the five harms I list as the basic harms. Andrew and Seumas also use the term “moral” in a somewhat wider way than I do, e.g., “Moreover, arguably truth, or at least moral truth, is a moral value and falsity a disvalue.” My colleague, Walter Sinnott-Armstrong, claims that false beliefs should count as a basic evil, and hence ‘Do not deceive’ should be listed with the first five rules rather than as one of the second five. He may be correct, but I consider false beliefs to be bad because they make one more likely to suffer one of the five harms and less likely to gain the goods, or to satisfy one’s rationally allowed desires. (In addition, knowledge is an ability, and false beliefs are incompatible with knowledge.) But in this context I do not understand exactly what is meant by “moral truth” and “moral value.” As they acknowledge, I do not use the term “moral duty.” For me, if one has a duty, based on either a social role, what they call an institutional duty, or on circumstances, then it is immoral not to perform that duty unless one has an adequate justification for not doing so. There are no moral duties, although there are moral requirements, namely, to obey the moral rules, including the moral rule requiring you to do your duty, unless you have an adequate justification for not obeying them.

In addition to clearly summarizing some features of my account of morality, Andrew and Seumas also have some criticisms of it. In particular, they claim that three additional rules are needed in order to account for all of our moral judgments about what is required or prohibited. “‘Do not steal or damage other people’s property,’ ‘Do not defraud,’ and ‘Do not defame.’” Their argument that these rules are not reducible to one of the original ten rules is correct. However, it is not necessary that each of these three rules be reducible to one of the original ten rules in order for them to be redundant. All that needs to be the case is that any violation of one of these three rules always involves a violation of one of the original ten

rules, though it need not be the same rule every time. The question is; can any of the three rules, 'Do not steal or damage other people's property,' 'Do not defraud,' and 'Do not defame,' be violated without any of the original ten being violated? This does not seem likely because causing a significant increase in the risk of suffering any of the harms counts as a violation of the first five rules. It is difficult to imagine a case of stealing or damaging another's property, defrauding, or defaming that does not involve violating at least one of the original ten rules. If it does not violate one of the second five, then unless it causes a significant increase in the risk of suffering one of the harms, I doubt that we would say that the action is immoral and needs justification. We probably wouldn't even call the act stealing, damaging, defrauding, or defaming.

However, Andrew and Seumas may be correct that there are such cases. But even if they are, this is primarily an aesthetic matter, i.e., ten rules divided neatly into two tablets of five each appeals to me aesthetically. I prefer the list of rules I have now, but that is because I think they are sufficiently general to rule out all of the kinds of acts that would be immoral if there were no adequate justification for violating them. However, if they are not, then I can add some less general rules. The ten rules are primarily a mnemonic; and I prefer to interpret them as ruling out as immoral all cases of stealing, damaging, defrauding, and defaming, but it is not an important philosophical matter if one prefers to interpret the rules narrowly and hence requires additional rules in order to account for all of our moral judgments about what is required or prohibited. I do not claim that the rules I list are the only moral rules, I claim only that these rules are both necessary and sufficient to account for all of our judgments about what is morally prohibited or required. They are what I call the basic general moral rules. There are many non-basic moral rules, including the three that Andrew and Seumas put forward, but I prefer to formulate the moral system in an aesthetically appealing way, consistent with all of the rules being clearly understood.

I have claimed that the rule against cheating is a basic general moral rule, not only because I think it necessary to account for some judgments not accounted for by the other nine rules, but also because cheating is an important, but neglected concept. It is not only important in academic context, but also accounts for the significance of fairness and for much that is misleadingly called social contract theory. Andrew and Seumas, possibly misled by Dean Cocking, seem to think that the example I give of the boss cheating at golf is not really an example of cheating. Their analogy between this case and family games of cricket is interesting but not persuasive. All those playing the family game of cricket have freely agreed to the change in rules. (Although a slightly older sibling who is not given this advantage might not agree, and then the change in rules is made explicit to him.) The employees have not agreed to change the rules of golf. This is shown both by their attitude toward their boss and how they report the game to their friends and family; complaining that their boss cheated again in order to win. Andrew and Seumas may claim that these employees are not competent speakers of the language, or they may hold, parallel to the Utilitarian view about morality, that philosophical analysis of concepts trumps ordinary use of terms. They may also hold that it is impossible to cheat at solitaire.

Andrew and Seumas also object to two rules that I do include; 'Obey the law' and 'Do your duty.' However, their objection to these rules seems to rest upon a view about the relationship between moral rules and moral obligation that I do not hold. Although it is a common view, I do not claim that there is a moral obligation or moral requirement, even a presumptive one, to obey any of the moral rules. I claim only that it is immoral to break any of these rules without an adequate justification. This difference in my formulation is not accidental. All of the examples that Andrew and Seumas give to show that 'Obey the law'

and ‘Do your duty’ are not moral rules depend upon holding that we have a moral obligation or moral requirement to obey a moral rule. But on the formulation that I put forward, we have no such moral obligation or moral requirement, not even a presumptive one.

A doctor does not have a moral obligation or requirement not to cause pain to a patient if causing pain is necessary to prevent death, greater pain or significant disability, and the patient gives his valid consent to the doctor to cause him pain. All of the particular examples that they offer where we have no moral obligation or moral requirement to obey the rules ‘Obey the law’ and ‘Do your duty’ are intended to be examples where we are strongly justified in breaking these rules. If we are strongly justified in breaking these rules, then we have no obligation to obey the law or do our duty. But if they are claiming that, without knowing anything more, it is morally irrelevant whether there is a law prohibiting our action, or whether we have duty to do it, then I think they are mistaken. Even under the Nazi regime, or a totalitarian state of the kind described in Orwell’s 1984, most laws, which means ordinary laws like traffic laws, parking regulations, and building codes, should be obeyed.

Andrew and Seumas consider my formulation of the duty to assist too limited. I say that a person has a duty to assist only when he is in a unique or close to unique position to do so, assisting is relatively cost free, and the person assisted will suffer serious harm if not assisted. Their criticism of my formulation is completely correct. Their example of a small group of bushwalkers being confronted with a small grass fire is completely persuasive. That group of people, given that they are in a unique or close to unique situation to put out the fire, that putting it out is relatively cost free, and failure to put it out will result in some people suffering serious harm, do have a duty to put out that grass fire. My concentration on individuals was excessive, and they have shown that I should have enlarged the situation in which circumstances impose duties to include being in a small group when all of these circumstances apply. If I have an opportunity to revise my formulation of that duty, I should change it to accommodate their criticism.

However, they move from this entirely correct criticism to a position with which I do not agree. Expanding on their example of the bushwalkers they present the following example, “Assume that there is a high probability of bushfires destroying the property and taking the life of some very small percentage of Australia’s population, but that with respect to any individual Australian the risk is close to zero.” I shall take their assumption to correctly describe the actual situation in Australia. Accepting this assumption, they conclude, “In this situation all or most Australians have a collective moral responsibility to prevent brushfires, and that responsibility can be discharged at minimal cost and inconvenience to any individual Australian. Thus they seem to conclude that all or most Australians have a moral (?) duty to take some action to prevent brushfires, either, setting up local fire prevention associations or instituting professional fire brigades. If they are using “duty” in the sense that I do, then if someone does not do what they have a duty to do, it is appropriate that they be liable to punishment. I wonder if they want to hold that those who do not help in setting up local fire prevention associations or instituting professional fire brigades, e.g., some classical musician or serious metaphysician, should be liable to punishment.

Once they have made this jump from being a member of a small group facing a small bushfire to all or most Australians acting to prevent future bushfires, it is not surprising that they make another jump. Now they want to claim, “our duty in relation to life is not simply not to *individually* kill, but also to *collectively* preserve life.” Again, the question arises, do they think it appropriate to make liable to punishment anyone who, without an adequate justification or excuse, does not take part in this collective responsibility. I hope that they do

not hold this, for it would expand the powers of government far more than I, and I hope also they, would like to see.

I think that what explains why they have put forward the view that they do is given in the final sentence of their paper, “These collective responsibilities are embodied in institutions and institutional roles, and justify the establishment and maintenance of such institutions and roles.” What they want to do is to “justify the establishment and maintenance of such institutions and roles.” They seem to hold that in order to justify them they must posit duties to establish and maintain these institutions and roles. However, that is not necessary. The moral ideals “justify the establishment and maintenance of such institutions and roles.” They not only allow the establishment and maintenance of such institutions and roles, they encourage doing so. Indeed, these moral ideals are sufficiently strong that they justify violating moral rules depriving people of their freedom not to participate.

Many philosophers find it surprising that the moral ideals, which are only encouraged, can justify breaking moral rules, which are required. But we are only morally required to obey the moral rules when we have no adequate justification for violating them. The moral ideals to establish and maintain institutions and roles that are necessary for preserving life, freedom, etc., do provide adequate justification for violating the rule against depriving of freedom not to participate. There is no need to invent individual moral duties to justify the establishment and maintenance of such institutions and roles; the moral ideals are sufficient to do that. I agree with their goal, which is to “justify the establishment and maintenance of such institutions and roles” that may deprive some people of some freedom, but I do not agree that my account of the moral system needs to be modified to achieve that goal.

Applying our Common Morality: The Case of Privacy

Jeroen van der Hoven

Australian National University

Bernard Gert has provided a clear, concise, powerful and accessible articulation of our ‘Common Morality’. Common Morality (CM) comprises ten moral rules, a set of related moral ideals, a two-tier procedure to arrive at moral judgments, which among other things facilitates the description of situations in terms of morally relevant properties, and the impartial application of the moral rules. CM is very practical in its focus. Its main purpose is “to attempt to describe, explain and justify our common morality” and its basic moral rules in order to provide “...a procedure for properly describing actions so that people can make better moral decisions and judgments” (7).

1. Moral Non-Decidability

The notion of a *procedure* needs to be qualified. Gert criticizes much of modern ethics for assuming that a moral theory or a moral framework should be complete and comprehensive, in the sense that it provides a unique solution for each moral problem, a correct and unique answer to all moral questions. Gert’s reconstructed CM does provide a procedure for moral decision making, but it does not provide an *effective* procedure; that is, it does not provide an algorithm or recipe that unfailingly leads via a number of recursive steps to a unique correct result. Gert points out that CM cannot be mechanically applied in this sense. Although he admits he can imagine different computer programs correctly implementing CM, they could lead to different outcomes. There will no doubt always be agreement on a vast majority of cases, but there will be irresolvable disagreement as well. Gert provides a list of the sources of reasonable disagreement, which is similar to Rawls’ list of sources of reasonable disagreement, discussed under the heading of ‘the burdens of judgment’ (Rawls 1993, p. 56ff): different scope of moral standing, different rankings of harms and benefits, different interpretations of the rules

2. CM and the Privacy Issue

I will look at Bernard Gert’s treatment of the right to privacy and will provide an analysis which corroborates the idea that CM provides an interesting conceptual environment for applied ethics and more in particular has something interesting to offer in the field of privacy.

The illustration of the privacy debate may help to illustrate two points: 1. the idea that there may be different solutions, without one of them being clearly wrong, and 2. the uncovering of a fine-grained structure of moral considerations which the list of moral rules may provide.

Gert does not provide a fully-fledged account of privacy but discusses privacy briefly at several places. I am sure he has more to say on the subject than he has done in CM. According to Gert in CM loss of privacy should be construed as a loss of freedom, since loss of freedom “includes loss of opportunity and resources and also loss of bodily integrity and privacy” (40). The right to privacy is not one of the ten principles but is an interpretation of the moral rules, and it does not collapse into one of the moral rules, like some of the basic

rights, which can be construed as straightforward counterparts of the moral rules dressed up in rights vocabulary.

Intentionally touching somebody, whom one has good reason to believe does not want to be touched, counts as violation of her freedom, and so is intentionally bumping into someone in the subway. Unintentional bumping is not a violation of privacy. In some circumstances, listening or looking at another person counts as a violation, since the freedom or option to be unobserved is thereby taken away. Another privacy case which is discussed is that of the celebrity who is followed by a photographer all day. She may be hurt (trespass, stalking, harassment) or may have her freedom diminished.

Gert admits that it is very difficult to determine when the rule “do not deprive of freedom” is exactly violated, or when an action leads to loss of freedom in a morally relevant sense. Typically violations take the form of making a person unable to do something by altering his environment or situation, rather than by altering the person himself.

2.1 Reasonable Disagreement

In the case of privacy there are two fundamentally different starting positions, which represent different fundamental conceptions of the relation of individual to society, or to the group of which he or she is part. First, one could argue - as the staunch defender of individual privacy rights does - that the moral rule “do not harm” or “do not deprive of freedom” is violated by surveillance, monitoring, and the unconstrained collection of personal data since these diminish the freedom of the individuals involved. The starting position however could also be the opposite, i.e. that there should in principle be no, or almost no, limits to the freedom to gather information. According to this idea privacy constraints themselves are violations of the rule “do not diminish freedom” (since they constrain other people’s freedom to collect information and acquire beliefs), or they themselves are violations of the rule “do not harm” (since the lack of information about individuals (anonymity) allows for a lot of free-riding and harm done to innocent individuals).

With respect to most privacy issues there is this fundamental ambiguity of starting position of the debate. Whose freedom should be taken into account and how should it be weighed, and how should it be weighed against violations of other rules? It is not evident that one of the basic positions is wrong.

In the Netherlands a court ruled in 2004 that the Prince Royal could not prevent photographers to stand and wait on the public street near his house. He had asked for a ‘journalist free zone’. The freedom of information for others overruled his privacy concerns. On the other hand, a month earlier the court ruled that damages should be paid to a Dutch female singer who was caught topless on camera, while lying on a public beach and found herself portrayed in a gossip magazine.

The privacy issue lies at the heart of an ongoing debate in nearly all Western democracies between liberalists and communitarians over the question how to balance individual rights and collective goods, individual rights and community interests. The privacy issue is concerned more specifically with the question how to balance the claims of those who want to limit the availability of personal information in order to protect individuals and the claims of those who want to make information about individuals available in order to benefit the community. This essential tension emerges in many privacy discussions, e.g. undercover actions by the police on the internet, use of Closed Circuit Television in public

places, making medical files available for health insurance purposes or epidemiological research, linking and matching of databases to detect fraud in social security, and soliciting information about on-line behaviour of internet users from access providers in criminal justice cases.

Communitarians typically argue that the community benefits significantly from having knowledge about its members available. According to communitarians modern Western democracies are in a deplorable condition and our unquenchable thirst for privacy serves as its epitome. Who could object to having his or her data accessed if honourable community causes are served? Communitarians also point out that modern societies exhibit high degrees of mobility, complexity and anonymity, and crime, free riding, and the erosion of trust are rampant under these conditions. Political philosopher Michael Walzer observes that “Liberalism is plagued by free-rider problems, by people who continue to enjoy the benefits of membership and identity while no longer participating in the activities that produce these benefits. Communitarianism, by contrast, is the dream of a perfect free-riderlessness”.

Recent analyses of privacy have rendered privacy as a cluster of moral considerations. Too many different things have been lumped together under the privacy heading. And it seems time to distinguish the moral considerations that have been conflated under the name of ‘privacy’. I have argued elsewhere that all we would like to be able to express about privacy can be expressed in terms of moral justifications for data-protection or in terms of moral justifications for constraints on access to personal information. Or to put it in terms of Gert’s CM framework: the privacy debate is about the moral justification of violations of a specification of the rule “do not deprive of freedom”, i.e., the freedom that personal information may be collected, processed and disseminated without constraint.

I think the following fine-grained type of moral justification along the lines of Gert’s CM for data-protection and limits to access can be distinguished: 1) do not harm in various ways identified by rules 1,2, 3 and 5, and; 2) do not deprive of freedom (rule 4), where a further threefold distinction seems to be appropriate.

3. Moral principles involved

3.1 Do not cause pain and do not disable, do not harm

The first type of moral reason for data-protection and for justified violations of the freedom of information is concerned with the prevention of harm, more specifically harm which is done to persons by making use of personal information about them.

Cyber criminals and hackers are known to have used computerised databases and the Internet to get information on their victims in order to prepare and stage their crimes. The most important moral problem with ‘identity theft’, for example, is the risk of financial and physical damages. One’s bank account may get plundered and one’s credit reports may be irreversibly tainted so as to exclude one from future financial benefits and services. Stalkers and rapists have used the Net and on-line databases to track down their victims and they could not have done what they did without tapping into these resources. In an information society there is a new vulnerability to information-based harm. The prevention of information-based harm provides governments with the strongest possible justification for limiting the freedom of individual citizens to access personal information. Policies that encourage rigorous security measures must be put in place to protect citizens against information-based harm. This seems a matter of security though and not of privacy. No other

moral principle than John Stuart Mill's Harm Principle is needed to justify limitations of the freedom of persons who cause, threaten to cause, or are likely to cause, information-based harms to people. Protecting personal information, instead of leaving it in the open, diminishes the likelihood that people will come to harm, analogous to the way in which restricting the access to fire arms diminishes the likelihood that people will get shot in the street. We know that if we do not establish a legal regime that somehow constrains citizens' access to weapons, the likelihood that innocent people will get shot increases. In information societies, personal information is comparable to guns and ammunition.

3.2 *Do not deprive of freedom*

Next to the rules against harming in various ways, three types of limiting freedom of individuals may be distinguished. These may be invoked to justify violations to the rule that people have the freedom to acquire information about the world (including fellow human beings) and to access personal data and use personal information.

3.2.1 *Freedom of Choice*

When someone reads your letters unbeknownst to you, snoops and spies, then your free choice to not be spied upon, or not have your stuff searched, is taken from you, and *ceteris paribus*, you are left with one less option. Moreover, your picture of the world is not adequate, you think you are unobserved, or that you are the first to read the letter, but that is wrong, someone else has read it before you, or is looking at you through binoculars when you are in your apartment. Everything you do is based on this mistaken assumption. Even if you know about the observation, the observer may cause you to feel awkward and self-conscious, and impose his view upon you, and thereby alter your condition.

3.2.2 *Economic opportunity*

If a consumer buys coffee at the shopping mall, information about that transaction can be generated and stored. Many consumers realize that every time they come to the counter to buy something, they can also sell something, namely, information about their purchase or transaction, the so-called transactional data. Likewise, sharing information about ourselves on the Net with web sites, browsers, autonomous agents may pay off in terms of more and more adequate information (or discounts and convenience) later. Many privacy concerns have been and will be resolved in *quid pro quo* practices and private contracts about the use and secondary use of personal data. But although a market mechanism for trading personal data seems to be kicking in on a global scale, not all individual consumers are aware of this economic opportunity, and if they are, they are not always trading their data in a transparent and fair market environment. Moreover they do not always know what the implications are of what they are consenting to when they sign a contract. We simply cannot assume that the conditions of the developing market for personal data guarantee fair transactions by independent standards. Data protection laws should be put in place that constrain the use of data in order to guarantee equality and a fair market for personal data. Data-protection laws in these types of cases typically protect individual citizens by constraining others in the use of those data and requiring openness, transparency, participation and notification on the part of business firms and direct marketers to secure fair contracts.

3.2.3 *Discrimination*

Many people do not object to the use of their personal medical data for *medical* purposes, whether these are directly related to their own personal health affairs, to those of their family, or perhaps even to their community or the world population at large, so long as they can be absolutely certain that the only use that is made of this data is to cure people from diseases, or the service of some other clinical, therapeutic, or medical goal. They do object,

however, to their medical data being used to disadvantage them socio-economically, to discriminate against them in the workplace, refuse them commercial services, deny them social benefits, or turn them down for mortgages or political office. Similarly, people do not usually mind if their library search data are used to provide them with better *library* services, but they do mind if these data are used to criticize their tastes, and character. They would also object to these informational cross-contaminations when they would benefit from them, as when the librarian would advise them a book on low-fat meals on the basis of knowledge of their medical record and cholesterol values, or a doctor poses questions about AIDS, on the basis of the information that one has borrowed a book from the public library.

We may thus distinguish another form of informational wrongdoing that privacy constraints seek to avoid and prevent, namely, disrespect for the boundaries of what we may refer to, following Michael Walzer, as “spheres of justice” or “spheres of access”. I think that what is often seen as a violation of privacy is often more adequately construed as the morally inappropriate transfer of data across the boundaries of what we intuitively think of as separate “spheres of justice” or “spheres of access.”

The meaning and value of personal and other information is local, and allocation schemes and local practices that distribute access to information should accommodate these local meanings. It is *pro tanto* morally wrong to use personal information without informed consent in another sphere than it was provided for. The information is relevant for some spheres and not for others. If the information is used to characterize the person, in one sphere on the basis of information from another sphere, that characterization may be morally irrelevant. Moreover, if there are negative consequences associated with that characterization, this amounts to unjustified discrimination.

3.2.4. Moral Freedom

There is a third and final freedom at stake when use of personal data would be unconstrained. That is *moral* freedom, the freedom to shape our own moral biographies, to reflect on our moral careers, to evaluate and identify with our own moral choices, without the critical gaze and interference of others and a pressure to conform to the “normal” or socially desired identities. Privacy, conceived along these lines, would only provide protection to the individual in his quality of a *moral* person engaged in self-definition and self-improvement against the normative pressures which public opinion and moral judgement exert on the person to conform to a socially desired identity. At this point Gert’s device of the identification of morally relevant features of the situation is needed. Especially a version of question 6, which prompts us to ask whether those protected by the moral rule, do not use the protection offered to violate this or other moral rules themselves.

Information about Bill, whether fully accurate or not, facilitates the formation of beliefs and judgements about Bill. Judgements about Bill, when he learns about them, or suspects that they are made, fears that they are made, may bring about a change in his view of himself, and may induce him to behave differently than he would have done without.

To modern individuals, who have cast aside the ideas of historical and religious necessity, living in a highly volatile socio-economic environment, and a great diversity of audiences and social settings before which they make their appearance, the *fixation* of one’s moral identity by means of the judgements, on the basis of very limited information, by others is felt as an obstacle to ‘experiments in living’, as Mill called them. The modern individual wants to be able to determine himself morally, or to undo his previous determinations, on the basis of more profuse experiences in life, or additional factual

information. Constraint on the freedom of others to access information about him can provide the leeway to do just that.

This conception of the person as being morally free, as being the author of his or her own moral career, provides a justification for protecting his personal data and limiting the freedom of others in this respect. Data-protection laws thus provide protection against the fixation of one's moral identity by others than one's self and have the symbolic utility of conveying to citizens that they are morally free.

I think this justification in terms of moral freedom and moral identity is essential to defending a right to privacy and limiting the informational and epistemic freedom of others to gain and use knowledge of persons. As far as I can see, this is not part of Gert's reading of his fourth moral rule, but maybe it can be read into it.

4. Applying Common Morality to moral problems of science and technology.

Although CM is successful as a scheme for dealing with practical inter personal moral problems, which involve the common moral wrongdoings, it may have trouble dealing with the moral problems of high technology, institutional design, socio economic systems, global infrastructures, in short the moral phenomena of so-called risk societies, characterized by teamwork, tight coupling, interdependence, manufactured risks, moral luck, constant innovation, and problematic expertise. Analysis of so-called "normal accidents" such as the Millenium Bug, the BSE crises, Tsjernobyl, or Bophal show that they are unlike the paradigm case of moral wrongdoing: Tom punched Dick on the nose, took his toy away and Harry didn't do anything. In normal accidents, agency is distributed in teams or groups, the causal chains are unclear, actions and consequences are not contiguous in time and space.

The vocabulary of the moral rules needs several layers of interpretations in order to be adequate for these types of cases. It may be slightly contrived to discuss the moral wrongs of software engineers who did not actively put the Millenium Bug on the public agenda in 1997 in terms of CM. What kind of moral obligations do government policymakers have regarding open source software, and how should the ICANN be restructured to guarantee fair allocation of new top level domain names, or how should the choice be made for cryptographic standards for Privacy Enhancing Technologies in Identity Management Systems for E-Government systems? No doubt there is a story to be told here which would eventually bring us back to the rule "Do not deprive of freedom", but I just cannot imagine how a real debate by policy makers and technical experts in these areas would invoke the moral rules in their canonical form in the setting of discussing the appropriate policies in these areas.

Gert observes: "We often hear the complaint that scientific advances are outstripping moral advances, as if we need to make new moral discoveries to deal with the new scientific discoveries and technology. We do need to understand how common morality applies to new situations, but there is no need for moral advances. Common Morality, together with an understanding of the new situations created by scientific discoveries and technology, is sufficient to deal with any problem with which we are confronted".

In a sense that may be true, but we may need to interpret the rules in such ways that to claim that they are still the same moral rules may be a bit implausible.

Reply to Jeroen van der Hoven: “Applying our Common Morality: The Case of Privacy”

Bernard Gert
Dartmouth College

I have learned much from Jeroen van den Hoven’s paper ‘Applying our Common Morality: The Case of Privacy.’ One of the advantages of providing a systematic account of morality is that it has many implications that are not obvious on a first, or even a second, reading. When these implications are made explicit, this provides an opportunity to test the theory. Insofar as the implications are acceptable or desirable, this tends to confirm the theory, but insofar as the implications are unacceptable or undesirable, this tends to disconfirm the theory. If the latter occurs, this does not mean the theory has to be abandoned, but it does mean that it has to be revised. In the more than four decades that I have been developing my moral theory, unacceptable or undesirable implications have resulted in many revisions. In recent years these unacceptable or undesirable implications have become less significant and so the required revisions have been less extensive. At the same time, some acceptable or desirable implications also have been discovered, and this has tended to confirm that the theory had some valuable features and makes it worthwhile to try to revise it to eliminate the unacceptable or undesirable implications.

Jeroen van den Hoven has discovered some implications of my theory with regard to privacy that are new to me. These are genuine implications of my theory, for Jeroen has presented a correct account of the relevant parts of my theory. Luckily, most of these implications are acceptable or desirable, so that they do not call for further revisions, but rather tend to confirm the theory. However, for this reason, any unacceptable or undesirable implications that he has discovered must be taken with utmost seriousness.

Jeroen correctly points out that the rule “Do not deprive of freedom” can be interpreted to prohibit “surveillance, monitoring, and collecting personal” information about persons and it can also be interpreted to prohibit any restrictions on gathering information. So the rule prohibiting the deprivation of freedom can be interpreted so that it protects privacy, or it can be interpreted so that it prohibits privacy restraints. These are what he calls the two basic positions and as he correctly says, “it is not evident that one of the basic positions is wrong.” Jeroen is correct that the rule can be interpreted in two conflicting ways, and thus the rule might be in conflict with itself. It this sounds a little too paradoxical, we can say that there are two acceptable interpretations of the rule and they may lead to conflicting moral judgments about what policy to adopt.

How is one to decide which interpretation of the rule “do not deprive of freedom” should be adopted? I shall adopt Jeroen’s labels and call the two interpretations, the liberalist and the communitarian, the former interprets the rule as prohibiting “surveillance, monitoring, and collecting personal information” and the latter as prohibiting any restrictions on gathering information. It seems as if one could, and perhaps should, adopt both interpretations, that is, one can say that, without an adequate justification, it is immoral to engage in surveillance, monitoring, and collecting personal information” and also that without adequate justification it is immoral to place any restriction on gathering information. Indeed, after reading Jeroen’s paper, it now seems to me, that it is correct to adopt both interpretations and to say that a justification is needed for violating the rule on either the liberalist or the communitarian interpretation. This is an implication of my theory, and so the

question arises about whether it is an acceptable or desirable implication. I see nothing unacceptable about it, so I regard it as an acceptable implication of my theory that it requires justification for violating the rule on either interpretation. Indeed, it seems to me correct to say that a person needs a justification for engaging in surveillance, monitoring, and collecting personal information, and also for restricting the gathering of information.

Notice that it is possible not to violate the rule on either interpretation. It is possible not to place any restriction on gathering information and also not to engage in surveillance, monitoring, and collecting personal information. These two interpretations are not contradictories, that is, it is not the case that if you obey one you must violate the other. They may not even be contraries, that is, a person might violate the rule on both of its interpretations, that is, both place restrictions on gathering information by others while at the same time engaging in surveillance, monitoring, and collecting personal information. The two interpretations of the rule can be treated as though they were two distinct rules, and so the question arises; what counts as an adequate justification for violating the rule on either interpretation? As with all of the other moral rules, it is possible that violations of the rule on either interpretation be strongly justified, weakly justified, or unjustified.

It may seem surprising that the rule “Do not deprive of freedom” has, at least, two interpretations that seem to conflict with each other. But, as pointed out above, these are not conflicting interpretations in any important sense, rather the rule can be interpreted as prohibiting a wide variety of actions, both actions in which private information is collected, and actions which prohibit collecting private information. To put it in an even more paradoxical way, we can say that the rule “Do not deprive of freedom” prohibits actions that deprive people of the freedom to deprive people of freedom. However, although this sounds paradoxical, all that is being said is that any deprivation of freedom, even the deprivation of freedom to deprive someone of freedom, needs to be justified. Depending on the freedom one is depriving a person of, it may be strongly justified, weakly justified, or unjustified to deprive him of the freedom to do that. It is, except in unusual circumstances, strongly justified to deprive a person of the freedom to kill another person, but unjustified to deprive him of the freedom to dye his hair green.

Determining whether an action is a violation of a moral rule is not, usually, the most difficult or important factor in making a moral decision or judgment. All that is decided by determining whether an action is a violation of a moral rule is deciding that the action needs to be justified. Sometimes, this is important, as in deciding that sexual activity between unmarried consenting competent adults, of whatever sex, does not violate any moral rules and hence does not need justification. However, generally, there is agreement about whether an action violates a moral rule and hence needs justification. Most moral disagreement is about whether that particular violation of a moral rule has the morally relevant features, that make it justified, strongly, weakly, or not at all. On the liberalist interpretation, it is clear that surveillance, monitoring, and collecting personal information is strongly justified when there is valid consent from the persons who are subject to that surveillance, monitoring, and collecting of personal information. On the communitarian interpretation, it is clear that placing restrictions on gathering information is strongly justified, when the restrictions prohibit people from gathering information in order to benefit those gathering it by harming those about whom the information is being gathered.

However, these are the easy cases, and are not the ones that are the topic of discussion. Of course, some of the difficult cases are not so much about gathering information, but about disseminating it. Although surveillance, monitoring, and collecting

personal information is strongly justified when there is valid consent from the persons who are subject to that surveillance, monitoring, and collecting of personal information, these persons may have given consent only for a limited number of people to have that information. So the problem now becomes disseminating information without consent, that is, what Jeroen, following Walzer, calls “the morally inappropriate transfer of data across the boundaries of what we intuitively think of as separate ‘spheres of justice’ or ‘spheres of access.’” One possible way to deal with this problem is to ask for consent to disseminate that information across these spheres. Since, if valid consent is obtained, there is no problem, the question arises, why not attempt to gain valid consent for disseminating that information? One answer is that the person about whom the information is being gathered may not want that information disseminated. People have the right to control some information about themselves, that is, disseminating some information about a person without his valid consent may be interpreted as violating the moral rule prohibiting depriving of freedom. If that interpretation is accepted then any dissemination of that information without an adequate justification is immoral. The difficult cases, as always, are deciding whether, in the absence of valid consent, there is an adequate justification for violating the rule.

Jeroen, using my terminology, frames the privacy debate quite well; “[It] is about moral justifications of a specification of the rule, ‘do not deprive of freedom,’ i.e., the freedom that personal information may be collected, processed, and disseminated without constraint.” His argument that it is justified to restrict people from gathering information in order to benefit those gathering it by harming those about whom the information is being gathered is one with which I completely agree. His arguments for restricting people from gathering or disseminating information that might be used to discriminate against them, seems to me to be an example of this kind of justified restriction. I hold that simply being deprived of your freedom to be unobserved is sufficient for requiring a justification for spying without bringing in further arguments. I am less clear about his argument for restricting information about “transactional data.” Insofar as collecting that data significantly increases the risk of the person suffering harm, then no further argument is needed to claim that it is immoral unless there is an adequate justification for collecting that data. However, it is not clear that, if dissemination of that information is controlled, that collecting that data does significantly increase the risk of the person suffering harm. All of these issues, especially the latter, seem to me to require information about the risks and benefits of collecting this “transactional data.” Indeed, Jeroen has shown that many moral problems cannot be settled without the appropriate empirical data. Further, that empirical data is constantly changing, so that one must continue to keep up to date to decide whether current policies should be modified with regard to the collecting and disseminating of private information.

I am much less clear about the concept of “moral freedom” that Jeroen puts forward as an additional justification for restricting access to private information. As I indicated in my reply to Andrew and Seumas, I use the term “moral” in a very restricted way. Every use is related in an intimate way to the moral system. I have no problem in talking about moral rules, moral ideals, morally relevant features, moral virtues, etc., but when the term “moral” is used in a way that is not related to the moral system, I do not know what it means. So, I cannot accept Jeroen’s claim, “moral freedom and moral identity is essential to defending a right to privacy and limiting the informational and epistemic freedom of others to gain and use knowledge of persons.” However, I do not see that as a significant loss, for I think that the other arguments that Jeroen presents for limiting access to information about persons are sufficient.

Jeroen is rightly skeptical of my claim that common morality is sufficient for dealing with the new situations created by scientific discoveries and technology. If he takes me to be claiming that an application of the moral rules is sufficient for doing this, he is correct. His remark, “we may need to interpret the rules in such a way that to claim that they are still the same moral rules may be a bit implausible.” is characteristically generous. I do not claim that the moral rules can be interpreted so as to answer questions like; “how should ICANN be restructured to guarantee fair allocation of new top level domain names, or how should the choice be made for cryptographic standards for Privacy Enhancing Technologies in Identity Management Systems for E-Government systems.” What is needed in these situations is not an interpretation of moral rules, but using the moral system to develop policies that can be taken as incorporating the point of morality, lessening the amount of harm suffered, more than competing policies.

I claim that the moral system, with its concepts of impartiality and rationality, and its goal of the lessening of harm suffered by all those protected by the moral system, does provide the framework for deciding what policy to adopt. Particularly important is recognizing that although the moral framework rules out all morally unacceptable solutions, it does not provide a unique correct answer to many controversial moral questions, including questions about what policy should be adopted, e.g., how should ICANN be restructured? There are many rationally acceptable rankings of the harms involved in adopting cryptographic standards, etc. Recognizing this fact should itself be of value, for it allows those who are seeking a solution to compromise without giving up their moral integrity. It is the view that there is a unique correct answer to all of these questions that causes much unnecessary conflict. Requiring agreement only on the range of morally acceptable answers may allow for all concerned to let the final decisions be made by those who are most likely to know what the effects of the different policies will be. It also allows for a change in policy without requiring holding that the previous policy was mistaken. As long as the moral system is used to limit possible solutions to morally acceptable ones, recognition that many of these decisions are political decisions can be of significant value.

Moral Arrogance and Moral Disagreement

Dean Cocking

Charles Sturt University

In *Common Morality*, Professor Bernard Gert develops a theoretical account of morality concerned to identify, describe and systematise the substantive rules and values of the moral practices common to most of us, i.e., what Gert calls our common morality. Such an approach will seem refreshingly useful, especially to many working in practical and applied ethics, since the more direct focus on providing an account of the nature and application of our common practices and related values promises a theoretical account with which we may more usefully engage and apply.

Central to any assessment of the ethical theory presented by *Common Morality* will accordingly be questions concerning how well the theory does identify, describe and classify the phenomena of our common moral practices and related substantive values.

In this paper I focus on two related aspects of Gert's theory: his account of moral arrogance in relation to whether or not there are uniquely right answers to all moral questions and his account of fundamental moral disagreements.

For Gert the problem of moral arrogance comes about on account of philosophers who put forward their favoured moral theory as providing uniquely right answers in all cases. In the first instance one might think that the problem of moral arrogance imagined here is that there are not uniquely right answers to all moral questions. Thus, those philosophers who claim that their favoured moral theory delivers uniquely right answers to all moral questions are mistaken about the moral territory and arrogant not to admit, in some circumstances, that their answer is not uniquely right.

But while this may be part of what Gert has in mind it is not the core part of his worry about moral arrogance. The core part of his worry is that philosophers who preach that there are uniquely right answers in all cases will develop or contribute to moral arrogance in those who listen to them – such as their students. For the effect on the hapless listeners will be to imagine that there are uniquely right answers to all moral questions and so to reject the conflicting answers of others as mistaken in circumstances when these conflicting answers are equally morally acceptable.

But whether one imagines the problem of moral arrogance with respect to moral philosophers and/or those who listen to them it is not clear that the proper diagnosis of the problem is that it is sourced by taking the view that there are uniquely right answers to all moral questions or even the view that one has these answers. In the first section of the paper I argue that the problem of moral arrogance is not sourced in these ways.

In the second section I query Gert's account of fundamental moral disagreements. Here I suggest Gert has not given us the right account of what sources our moral disagreements. On the alternative account I present, the sorts of disagreements Gert claims are fundamental turn out either to reflect plural and conflicting values about which we have, or should have, some agreement; uncertainties with respect to the circumstances within which we pursue these values, and/or; implausible views, which, while in fundamental conflict, since they are implausible, should be rejected.

1. Moral Arrogance

Gert thinks there need not always be a uniquely right answer and that it is morally arrogant to hold that there is a uniquely right answer in cases where there is not, i.e., in circumstances where there are conflicting but equally morally acceptable or reasonable answers. And since Gert makes a strong connection between thinking one's answer uniquely right, its attendant arrogance and our capacity to engage in democratic deliberative procedures and to have fruitful discussion with those who hold conflicting views (145-8), his analysis of moral arrogance promises some broad significance. In querying his diagnosis of arrogance I'm thereby querying whether correcting the failure to recognise one's own view as not the uniquely right view is the solution either to moral arrogance or to the threat one may rightly be concerned moral arrogance might present to our capacity to have fruitful discussion with those with whom we disagree and to a culture of deliberative democracy.

Gert claims that often enough there is not one uniquely acceptable or reasonable answer concerning what it is right to do. There may be conflicting, incompatible but nevertheless equally acceptable or reasonable answers. Thus, for Gert, a moral theory that holds there is a unique answer to all moral questions is mistaken. Many moral theorists nowadays hold a conflicting and plural value view of good and right action. And some such theorists are happy to admit that in some cases where irreducibly different values are in conflict that there is no uniquely correct answer as to what is good and/or right. In some cases, either of the conflicting judgements or actions in question may be morally acceptable or reasonable. Thus, e.g., I may pursue values of family and/or friendship at the expense of values concerned with aiding others, fulfilling other social/institutional roles or defending my country and sometimes the conflict in values may be unresolvable and it may be morally acceptable or reasonable for me to go either way.

On this view, the justificatory force of either of the conflicting options need not be over-riding. Moral theorists who allow this typically think of this situation as a moral dilemma. It is a rather striking feature of Gert's account of how there can be incompatible but equally morally acceptable answers that he does not talk in terms of plural and conflicting values, nor talk in terms of moral dilemmas. In fact, (from conversation) it turns out he doesn't think they are dilemmas. Rather, he thinks one person may simply judge or act in a certain way and another do so in an incompatible way, but neither need think they face any kind of dilemma. Each simply judges or acts according to her own incompatible but morally acceptable lights.

Nevertheless while, e.g., I may pursue obligations to family at the expense of obligations to country or the wider community and *I* may do so without recognition of a dilemma, this may well *in fact* represent a dilemma. Of course, Gert can agree with this. He will say however, that the force of this is just to say what he's saying, i.e., that there need not be a uniquely right answer. For Gert (again, from conversation) the moral dilemma theorist invariably fails to take the extra step and say what he (Gert) is saying. They recognise that neither of the conflicting values in play need be over-riding but fail to make out or focus attention on the implication of this, i.e., that there need be no uniquely right answer.

I agree that this is a correct implication of the dilemma situation but it seems that Gert is failing to take a necessary step of his own. For if two different people have two incompatible but equally acceptable judgements in a given case and both are – as Gert insists – equally and fully informed of the salient facts, then it seems they ought to accept or recognise the others' judgement to be as morally acceptable as their own. After all, ex-

hypothesi, this is precisely what we are imagining is true of the situation. And if one ought to recognise the conflicting path to be as morally acceptable as one's own path, then one ought to recognise neither option – including one's own - provides over-riding justificatory force. Again, this just is the situation. So each ought not simply pursue her own path indifferent to the force of the path of the other. If each recognises the equivalent moral acceptability of the other's incompatible path, then they ought not be indifferent to the force of the other's path. They ought to feel the force of being in something of a dilemma.

In fact, if, as Gert says, we must recognise our own view need not be the only correct view, wouldn't it be arrogant to not take on board the equal force of the conflicting view and so regard oneself as in something of a dilemma? Otherwise it just looks like one is merely patronising the other by paying only 'lip service' to the idea that the other's view is equally justified. It would be like the person who says to their detractors: "I hear what you are saying", but whose judgement or choices continue on their merry path utterly unshaken by what his detractors are saying.

If this is right, then holding the view that there is a unique answer to all moral questions is not necessary for moral arrogance. And while I agree that the view that there is a unique answer to all moral questions is likely mistaken, I think one can be mistaken without being arrogant –and I think one can be so with respect to thinking there is only one answer to a certain moral question. So I also think holding the mistaken view that one's answer is uniquely correct also need not suffice for arrogance.

The case upon which Gert bases his analysis of moral arrogance is where one thinks one's answer uniquely right but one is mistaken. One's answer is not uniquely right.

The morally arrogant person Gert describes takes the attitude that there is only one correct and so morally acceptable answer and that their answer is it in circumstances when *as a matter of fact* it is false that there is only one correct answer. Gert holds that without the recognition that there are conflicting morally acceptable answers you will lack moral humility and possess arrogance.

But being humble and not being arrogant seem compatible with holding that there is only one right answer. Suppose I think there is only one right answer. So I think it is not true that there are conflicting morally acceptable answers. Am I thereby arrogant or lacking in humility? It does not seem so. After all, I might not even *yet* think *I* have the right answer. I just think there *is* only one right answer. This seems to suggest that the problem regarding arrogance is not the thinking that there is only one unique right answer. At best, it is something about the attitude one takes to thinking there is only one answer, such as being dogmatic and resistant to listening to the un-cooperative reasons to the contrary put by others.

Thus, a humble attitude, without arrogance, could hold that there is only one right answer but be open to the possibility that a conflicting view is the correct one. My thought then, is that you are arrogant if you possess a lack of open-ness to the possibility that you may be wrong, but you may have this sort of receptivity and also think there is only one right answer. Indeed, it seems compatible with this receptivity that you might not only think there is a unique morally acceptable answer; you may also think that your answer is it. So one could think there is only one right answer, think one has this answer, but also think it possible one might be mistaken on both counts. Why would this count as being arrogant?

If this is right, then it is not the case that on account of the dis-value attached to moral arrogance, one must give up the view that there is only one correct answer, nor even the view that one has this answer. Thinking one's view uniquely right is neither necessary nor sufficient for moral arrogance. One just has to give up being arrogant about it, i.e., refusing to consider the reasons put to you that would suggest you're mistaken.

2. Fundamental Moral Disagreements

My second query concerns whether Gert has given us the right account of what sources our moral disagreements. Here my claim is that Gert's account of the conflicting concerns that are fundamental either 'dummy' for more substantive values that a plural and conflicting value or dilemma theorist would posit (and so values about which we have some agreement) or they do not really present fundamental conflicts (139-142).

I've suggested that in fact, *contra* Gert, it looks more like arrogance to not take seriously the force of a conflicting judgement one recognises as equally morally acceptable and to not regard oneself as posed with something of a dilemma in such circumstances. But this suggests some fundamental type of agreement: namely, we each recognise the force of the conflicting values at stake.

For Gert some key sources¹ of fundamental conflict are:

1. Dispute over the scope of morality –who is protected and who is not.
2. Dispute over the rankings we give to benefits and harms.
3. Differences in Interpretations of the Rules (14-5).

Take the third of these sources of fundamental conflict - different interpretations of the rules. As his key example, Gert presents us with the case of the dispute over whether removing the respirator (with autonomous consent) counts as killing. Now he must mean 'killing' here in the morally loaded sense, i.e., killing as a harm or evil. Otherwise, what moral conflict or disagreement could there be? If one did not think the killing a bad thing, as something harmful or evil, presumably one would not take up the option of being against it. If there is a dispute, where one party claims it would be wrong to turn off the respirator because to do so would be to kill the patient, then this party to the dispute must presumably think the killing is a bad thing.

For Gert, conflicts over whether taking the consenting patient off the respirator counts as a killing, are sourced by his other sources of fundamental conflict, i.e., regarding the scope of who is protected and who is not and the rankings of benefits and harms.

Thus, the conflict would be sourced by either (or both) disputes over whether the killing counted as a benefit or harm and disputes over the extent to which the patient counts as protected by morality. I agree that either or both these considerations may be party to the dispute. Certainly, the dispute seems primarily about whether the killing is a harm or a benefit to the person. In any case, neither source of moral disagreement seems to show any fundamental dispute with regard to the respirator case.

For *ex-hypothesi*, assuming autonomous consent in the standard circumstances of the relief of incurable misery, the killing is a benefit to the patient, not a harm. If someone thinks otherwise then either they are confused about this, e.g., they think the killing in-itself necessarily bad even though it may be an outstanding relief and great benefit to the person killed, or they may be uncertain. Even if they are not uncertain about its being of great

benefit, they may be uncertain whether autonomous consent really applied. They may, e.g., think the person can't really consent or they may (rightly) be very uncertain about the consent involved and therefore uncertain as to whether the killing is justified.

Thus, on the strength of the example offered, I don't see how disagreement about the interpretation of the rules is shown to suggest a fundamental moral disagreement.

The disagreement in the respirator case seems based on confusion, say, that the killing must necessarily be bad, and/or based on uncertainty – will it be of more benefit than harm? Does the person really give their autonomous consent? If we suppose autonomous consent and that we are certain they have given it, and if we suppose that the killing would in fact be singularly beneficial to the patient and that we are not confused about this, then, it seems, we take away the disagreement. Certainly, it is hard to see how there is fundamental disagreement that is sourced by fundamental disputes regarding the scope of morality and disputes over the rankings of benefits and harms.

Now consider Gert's second source of fundamental disagreements - dispute over the rankings of benefits and harms. The key sort of dispute Gert has in mind here is between "whether reducing the risk of being killed or injured by a specified amount outweighs a specified loss of freedom, for example, whether it is morally justified to have a law requiring everyone to wear seat belts" (14-5). Similarly, Gert has in mind here such examples as the dispute over the precise degree of probability required that someone will self-harm or harm others before it is permissible to commit them involuntarily. How fundamental are these disagreements?

Again, it seems the source of the disagreement is more plausibly supposed as the result of: (i) our *uncertainty* about whether, e.g., committing the person will be a harm or benefit to them, and; (ii) the conflict between the values that under-pin the benefits and harms, with respect to which, we are uncertain we will bring about.

On the uncertainty, we first want to know whether they would act autonomously if they did self-harm. Perhaps this is fairly easy for us to satisfy – only in rare cases, e.g., the political hunger striker – do we see a reason worth respecting as an expression of their autonomy. Typically, the self-harming person has some serious problems that have diminished their autonomy. Suppose then we are clear that if they were to self-harm they would not be acting autonomously. How do we assess the probability *they* will self-harm? i.e., how do we establish with any certainty that, say, 10% probability applies to *this person* rather than, say, 30% probability? In fact, given, for instance, the highly individualised nature of such problems, it is hard to see how we could have conclusive evidence either way. Thus, this case of divergent yet fully-informed views regarding the probability of harms balanced against a loss of certain freedoms seems more about how we may well disagree on the basis of reasonable guesses that track limited and conflicting evidence.

Perhaps though, this line of thinking misses the point. The point is not that there is conflicting and limited evidence upon which one may reasonably guess either way – say, whether or not this particular person is a 10% or 30% risk to self-harm. Rather, the point is that people may reasonably diverge about whether 10% or 30% risk is required to justify restricting a person's freedoms by committing them to an institution. And this disagreement will carry on irrespective of whether or not we can make an authoritative assessment of such risks in a given case. Even if we can't really make such assessments in a given case, the point is that we reasonably disagree about which %risk applies to provide the justificatory point at

which we may act paternalistically to preserve the person's (larger) interests by (in part) restricting their freedoms for a time.

But now we have simply returned to the world of plural and conflicting values whereby there may not be a uniquely right moral answer and conflicting options may each be morally and rationally acceptable ones. Both the concern to preserve another's more significant interests, such as their life, and the concern to not restrict their freedoms are significant values. Both these are important values and reasonable people who, as Gert requires, should respect each other's incompatible judgements, should then (again) find themselves conflicted. They should each see the force of the other's well-founded but incompatible view.

So while there is disagreement about how we rank the benefits and harms, the disagreement is about the conflicting values at stake which under-write the benefits and harms and we can both agree about the acceptability of these values. Seeing the conflicts in this way - in terms of conflicting values and uncertainty - also explains how we would expect those making the decision to feel, i.e., conflicted and uncertain.

Now consider Gert's first source of fundamental moral disagreement, i.e., the scope of morality. Take his key case: abortion. Here, he says, the fundamental disagreement is about whether morality protects the foetus. Of course, some conflicting views about the status of the foetus do present fundamental moral disagreements. If, for instance, one thinks morality applies only to persons or beings in so far as they possess person-like qualities, then one may think morality affords the foetus no protection at all. Of course, notoriously, on this view two year old babies would also be in trouble. On the other hand, one may reject this conception of the scope of morality applying only to persons by claiming a 'sanctity of life' approach. On this view the foetus is protected by morality since human life itself is sacred. Of course, notoriously, on this view the fertilised human egg would also seem sacred from the moment of conception. So while these two views present an example of fundamental conflict, neither are very plausible. Our moral responsibilities toward others are not exhausted by considerations of personhood and even if one imagined human life itself was protected by morality it is something else (even less plausible) to imagine it to be *absolutely* protected.

A more plausible account of some of the conflicting values that may be at stake would not limit the application of moral reasons to personhood and not claim the force of moral reasons protecting the foetus to be absolute. The fact is, many women do feel conflicted about their decision to abort and this sort of approach would make sense of it. Thus, a pregnant woman may feel the force of various reasons related to an interest in and concern for her foetus, but nevertheless, be faced with all sorts of countervailing reasons against becoming a mother – e.g., she may not be able to support a child, be too young to do so or not really want to do so. Such conflicting reasons may win the day without denying the legitimacy of reasons in of favour carrying the foetus. Presumably, this is why women quite commonly regret, say, that circumstances were not different and that they could not have proceeded with their pregnancy, while nevertheless believing that all things considered, given the circumstances, terminating the pregnancy was the right thing to do.

Thus, while I think there are conflicts in values presented by the abortion case and people can reasonably disagree about which values they favour, they can (and should) agree about the legitimacy of the conflicting values at stake. So while presenting conflicts and something of a dilemma, we are not presented with the fundamental moral disagreement concerning the scope of morality claimed by Gert.

In sum, I have made the following claims against Gert's related accounts of moral arrogance and fundamental moral disagreements. First, that it is not the case that on pain of moral arrogance, one must give up the view that there is only one correct answer, nor even the view that one has this answer. Thinking one's view uniquely right is neither necessary nor sufficient for moral arrogance. One just has to give up being arrogant by refusing to consider the reasons that would suggest one is mistaken. Second, against Gert's account of fundamental moral disagreements, I have argued that the sorts of disagreements Gert claims are fundamental are sourced either by plural and conflicting values about which we have, or should have, some fundamental agreement; uncertainties with respect to the circumstances within which we pursue these values, and/or; implausible views, which, while in fundamental conflict, since they are implausible, should be rejected.

Notes

¹ I leave out the empirical test –I'm not sure this really counts as a source of fundamental moral disagreement. For Gert's four sources of fundamental disagreement, see p. 140.

Reply to Dean Cocking: “Moral Arrogance and Moral Disagreement”

Bernard Gert
Dartmouth College

I have written so many books and articles on related topics in moral philosophy that I am often not clear where I have discussed a particular topic. I have passed along this confusion to many of the people with whom I have had extensive discussions on these topics. Dean Cocking is one of the persons with whom I have had extensive and productive conversations concerning many topics in moral philosophy, including conversations on moral arrogance. Thus it is not surprising that in part of his contribution to the symposium on *Common Morality* he discusses the topic of moral arrogance, which surprisingly, I do not mention in that book.

I should have discussed this topic in this book, and I do discuss many of the issues that are related to the topic of moral arrogance, but unfortunately I do not explicitly discuss this topic in *Common Morality*. It is the focus of a paper, “Moral Arrogance and Moral Theories,” that I read at a CAPPE seminar, and which was much discussed by Dean and myself.¹ I mention this explicitly at the beginning of this paper, for someone reading Dean’s paper and then looking at *Common Morality* may be puzzled why he is discussing a topic that I do not discuss in that book.

“Moral arrogance” is, to my knowledge, a phrase that I invented. However, although it does not have a standard meaning, it is not plausible to interpret it in the way that Dean does, “it is morally arrogant to hold that there is a uniquely right answer in cases where there is not, i.e., in circumstances where there are conflicting but equally morally acceptable or reasonable answers.” Neither I nor anyone else would regard someone as arrogant simply for holding this abstract position. But it is an essential feature of moral arrogance for a person to hold that he has a uniquely right answer in a case where there is no uniquely right answer. Moral arrogance is arrogance caused by the feature just mentioned. “Arrogance consists in viewing yourself as not subject to the constraints that morality imposes on all rational persons.” (Gert 2005, 307)

One of the constraints that morality imposes on all rational persons is that they not break a moral rule “when they would not want everyone to know that they are allowed to break the rule in the same circumstances” (135). To do this is to exhibit the standard kind of arrogance. The difference between moral arrogance and standard arrogance is that the morally arrogant person is more than willing for everyone to know that they are allowed to break the rule in the same circumstances; he may even hold that everyone is required to break the rule in these circumstances. That is, the morally arrogant person holds that he is justified in breaking a moral rule with regard to some person because that person acted or was planning to act in a way that the arrogant person regards as morally prohibited. Moral arrogance involves viewing yourself as not subject to the constraint that a person cannot use any beliefs other than those beliefs that all rational persons share in order to determine what kinds of acts morality prohibits, requires, discourages, encourages or allows. Standard arrogance involves the kind of violation of a moral rule such that no rational person would be willing for every moral agent to know that this kind of violation of a moral rule is allowed, even if they are using only those beliefs that all rational persons share. Moral arrogance

depends on using beliefs other than those that all rational persons share in determining what morality prohibits and requires.

A morally arrogant person most closely resembles a person who is arrogant on religious grounds. A religiously arrogant person holds that his unjustified violation of a moral rule is justified because he holds that doing what his God or his religion tells him to do overrides what morality requires or prohibits. The confusion about the relationship between religion and morality makes it difficult to distinguish between a morally arrogant person and a religiously arrogant one. A person may hold that the morally right action is the one that his God or his religion says is the morally right action and be unconcerned that not all moral agents know what his God or his religion says. He does not acknowledge that the kinds of actions that morality prohibits and requires must be known to all moral agents. He regards himself as knowing something about morality that not all moral agents know. He regards this knowledge as justifying his breaking a moral rule in circumstances that no moral agent limiting himself to beliefs shared by all moral agents, (except, of course, for beliefs about the facts of the particular case) would publicly allow.

Moral arrogance involves holding that your moral decision or judgment is the only correct decision or judgment on a controversial topic when you are aware that equally informed moral agents make different decisions or judgments. If you do not know about any such conflicting moral decisions or judgments you are simply mistaken, not morally arrogant. Moral arrogance occurs only on controversial topics, e.g., abortion, when it is well known that moral agents disagree and that no argument has been put forward that all of them agree resolves the issue.² Even though a morally arrogant person knows that other moral agents do not accept his answer, he regards himself as having the unique correct answer, and takes this to justify his breaking a moral rule with regard to another person. The morally arrogant agent does not acknowledge the constraint that he use only beliefs shared by all moral agents in determining what kinds of acts morality prohibits and requires and this makes it appropriate to regard moral arrogance as a kind of arrogance.

Dean and I agree that “being humble and not being arrogant are compatible with holding that there is only one right answer.” However, holding that there is a unique right answer to all moral questions provides support for holding that *you* have the uniquely right answer to a controversial moral question. Of course, a person who holds the general abstract view may be saved from this feature of moral arrogance by epistemic humility. This purely epistemological way to avoid moral arrogance may work, but it seems that only a commitment to total moral objectivity combined with a profound moral scepticism would lead anyone to hold that there is a unique right answer to every moral question and at the same time hold that there is no way to know what it is. Those persuaded by a philosopher’s claim to have the uniquely correct answer to a practical moral question are unlikely have such epistemic humility. My advice to philosophers is that they explicitly point out that many controversial moral problems have no unique correct answer. Were they to do this they would be making a valuable contribution to combat moral arrogance.

The concept of moral arrogance presupposes that there are some irresolvable moral disagreements, that is, that there are some topics about which equally informed impartial rational persons can make conflicting moral decisions or judgments. Moral arrogance also presupposes that there are some topics about which all equally informed impartial rational persons would make the same moral decision or judgment. Thus, the account of morality that underlies the concept of moral arrogance denies both that there is a unique correct answer to every moral question and also that there is no unique correct answer to any moral question.

In the course of his discussion of moral arrogance, Dean says, “It is a rather striking feature of Gert’s account of how there can be incompatible but equally morally acceptable answers that he does not talk in terms of plural and conflicting values, nor talk in terms of moral dilemmas.” First, it may be significant to point out that I never use the phrase “equally morally acceptable answers.” I do not use it because it is misleading; it suggests that there is a method for ranking morally acceptable answers, some of which will be better than others, but others of which will be equally morally acceptable. There is no such method; answers are either morally acceptable or they are not, there is no objective ranking within the class of morally acceptable answers. It is also misleading to talk in terms of plural and conflicting values. All rational persons agree on the basic values, that is, on what is a basic harm (evil), death, pain, disability, loss of freedom and loss of pleasure and what is a basic benefit (good), consciousness, ability, freedom, and pleasure; they disagree only on the rankings of these goods and evils, that is, on what is better and worse. This is significant, for it makes clear that all disagreements occur within a larger context of agreement.

Dean correctly claims that I think that “one person may simply judge or act in a certain way and another do so in an incompatible way, but neither need think they face any kind of dilemma.” If a person is aware that there are different rational rankings of the various evils, he can simply accept that he has a different ranking than someone else. He need not consider that either he or the other person is facing any dilemma, for both may realize that it is morally acceptable for each of them to act according to their own rationally acceptable ranking. However, Dean claims that “If each recognizes the equivalent moral acceptability of the other’s incompatible path, then they ought not be indifferent to the force of the other’s path. They ought to feel the force of being in something of a dilemma.” I do not see the force of Dean’s claim. Why should someone who knows that different people have different rational rankings of the various goods and evils be concerned when he discovers an instance of this?

Dean’s and Sartre’s example of competing obligations to family and to country is for me, not a basic disagreement; basic disagreements are about the rankings of the good and evils, not about competing obligations. But even taking this example, it turns out that everyone agrees that it is good to take care of your family and that it is good to defend your country from unjust aggression; what they disagree about is the relative weight of these competing claims. Dean’s view that accepting that there are conflicting moral claims must lead to moral dilemmas is plausible only if an agent holds that there is unique correct answer to every moral problem. Then, awareness that someone else accepts a different solution is troublesome, for at least one of them must be mistaken. This is another serious problem with the purely epistemological way to avoid moral arrogance. But a person who recognizes that there is more than one morally acceptable answer to some moral questions, can “continue on his merry path utterly unshaken by what his detractors are saying.” Such a person is not “merely patronizing the other by paying only ‘lip service’ to the idea that the other’s view is equally justified” because he is prepared for the other to also “continue on his merry path utterly unshaken by what” the agent is saying and doing. Such behaviour is not arrogant but tolerant, and is exactly the kind of behaviour that I am trying to encourage.

Dean may be misled by his continuing use of the phrase “equally justified,” which suggests that both parties hold that the goods and evils should be ranked in the same way. Whereas what is actually going on is that each has her own ranking of the goods and evils, while recognizing that it is rationally allowed to rank the goods and evils in a different way. I prefer to say that each of the parties knows that his behaviour is weakly justified and, if

neither is in any position of authority over the other, then each should be prepared to let the other act as he thinks best. I suspect that Dean continues to think about moral arrogance as an epistemological problem, for only if one thinks about moral arrogance in that way, will his objections seem persuasive.

1. Fundamental Moral Disagreements

Dean lists some of what I regard as key sources of fundamental conflict:

1. Dispute over the scope of morality –who is protected and who is not.
2. Dispute over the rankings we give to benefits and harms.
3. Differences in Interpretations of the Rules.

In discussing the “dispute over whether removing the respirator (with autonomous consent) counts as killing,” Dean says, “Now he must mean ‘killing’ here in the morally loaded sense, i.e., the killing as a harm or evil.” It is not using “killing” in the “morally loaded sense” to say that it always results in death; it is a tautology that killing someone always involves causing his death. However, the “dispute over whether removing the respirator (with autonomous consent) counts as killing” is morally significant, for killing is a violation of a moral rule and so needs justification in order not to be immoral. Abiding by a competent patient’s rational refusal of treatment is not killing, and so does not need a justification in order to avoid being immoral.

It is clear that Dean is not using “a harm or evil” in the same sense that I do. Killing causes a harm, viz., death, but it is not itself an evil or a harm, that is, “something that all rational persons would avoid unless they had an adequate reason not to.” Dean does not approve of this objective use of “evil” or “harm.” He prefers to use these terms in such a way that no one can be in a situation, e.g., being faced with either continuing pain or death, where he chooses what he ranks as the lesser of two evils. For Dean if he chooses death he is not choosing an evil at all. I do not see any advantage in using “harm” or “evil” in this relativistic way rather than in the objective sense in which I am using it. Using my sense, one can simply list the five basic evils, and allow people to differ in their rankings of these evils. On my view if a person’s ranking of the evils is rational, it is usually unjustified to prevent him from acting on that ranking.

Dean also seems to object to my claim that rational people can rank the evils in different ways, preferring to consider the disagreement about whether to commit a person who has a specified risk of killing himself as a disagreement based on uncertainty about the facts. Of course, there is such uncertainty; and disagreement about the facts, including predictions about what will happen, is the major source of moral disagreement, but this kind of disagreement does not cause any philosophical problems. It is irresolvable non- factual disagreement that most philosophers seem to deny. But after seeming to object to my view, Dean says the disagreement is simply a case of “plural and conflicting values.” But this is simply Dean’s misleading way of talking about a disagreement in the rankings of the various evils or harms. People do not have plural and conflicting values, as if neither acknowledges the values of the other, rather both regard death and loss of freedom as evils, and disagree whether a significant but low risk of death ranks higher or lower than a certain but temporary loss of freedom. Dean and I have no real disagreement here, only a difference in preferred terminology, and I see no advantages, and some disadvantages in Dean’s terminological preferences.

Dean's disagreement with me about the scope of morality also seems to be primarily a preference for a different terminology. I do not say that a person must hold either that a foetus is impartially protected or not protected at all, rather, I allow for rational persons to take either of these positions or to take some intermediate position, e.g., that morality protects foetuses to some degree but not as much as it protects moral agents. This involves a disagreement about the scope of morality, even though it also involves a ranking of the various goods and evil involved. Thus with regards to foetuses or non-human animals, different people can weigh the limitation on the freedom of moral agents as higher or lower than the death of non-moral agents. If this same limitation of the freedom of a moral agent were being ranked against the death of another moral agent, then there would be no question; all rational persons would rank the death of the moral agent as higher, and would not allow the moral agent to be killed.

Dean has raised some important questions, and has forced me to become clearer about what I mean by "moral arrogance." However, I think that one of the most important lessons to be learned from Dean's paper is the importance of terminology. There can seem to be important substantive differences that are really only differences in terminology. However, it is not insignificant which terminology one uses. Some terminology makes it almost impossible to express a view that everyone holds, e.g., that we are sometimes in situations where it is rational to choose the lesser of two evils. Other terminology, e.g., "plural and conflicting values" leads one to overlook that all rational persons agree on what counts as the basic goods (benefits) and evils (harms). When these differences in terminology are properly understood, I think that there is far less difference in substantive views between Dean and myself than initially seems to be the case.

Notes

¹ An extensively revised version of this paper will be published in *Noûs* Supplementary Volume 15 on Applied Ethics. Many of the revisions were due to my discussions with Dean.

² See my paper "Moral Disagreement and Abortion," *Australian Journal of Professional and Applied Ethics*, Volume 6, Number 1, June 2004. pp. 1-19.

References

Gert, B., *Morality: Its Nature and Justification*, Revised Ed. Oxford University Press: London.

Timeless Wisdom or Moral Arrogance?

Thomas Pogge

Australian National University

Gert's view about morality has been evolving for several decades now and, since it has been quite influential, it has been intensively discussed and criticized. Being a novice to these debates, I don't feel confident that I master the terminology Gert is employing and can fully appreciate the objections and responses that have been exchanged heretofore. So I come to the game with some humility. But, in the spirit of a good discussion, I will keep this humility to myself and try to raise some forceful objections to the view Gert has here so clearly and economically stated.

Gert sets himself apart from other moral theorists by attributing to them the view that "morality always provides a uniquely correct answer to every moral question" (145-6, *et passim*). In a recent paper, Gert has extensively criticized this view as an instance of "moral arrogance" (Gert 2005).

I doubt that the theorists Gert singles out for critique, Immanuel Kant and John Rawls, actually hold the view Gert attributes to them. In Kant's defence, I would invoke *inter alia* the third formula of the categorical imperative, which introduces the notion of a realm of ends, which is a unified, harmonious whole of all ends — "both of rational beings as ends in themselves and also of the particular ends that each may set for itself" (Kant 1912, 433). The basic idea of this formulation of the categorical imperative is that each agent is to think of herself as a legislator who is to determine, once and for all, which maxims it is to be permissible and impermissible for persons to act on. Qua private agents, we may act only on the maxims that, qua legislator, we would make available to all.

One important contribution this formula makes to clarifying the content of the categorical imperative is that it requires coherence: An agent must consider all the maxims she deems to be permissible *together*, that is, the universal laws corresponding to these maxims must fit together into ["zu"] a possible realm of ends¹ (Kant 1912, 436). She can rationally will any one of her maxims only if she can will its universal availability, and she can rationally will any one of her maxims to be universally available only if she can will all of them to be universally available jointly.

Applied in the area of imperfect duties, to the issue of mutual aid for example, the categorical imperative thus prohibits a double standard: The maxim governing one's requests for help from others must not be more expansive than the maxim governing one's own helping of others. If one refuses to give help **H** (of a certain sort, in certain circumstances in which one is not naturally, perhaps from sympathy, inclined to help), then one may not request such help for oneself. For in the world in which everyone felt free to refuse such requests, they would be refused and the end of receiving such help could thus not be attained and the maxim of requesting it could then not be willed. This thought experiment shows that a maxim may be rejected because its universal availability would defeat the material end — not of *this* maxim, but — of some other maxim to the availability of which the agent is committed. The thought experiment does not show which of the incompatible maxims the agent must give up. So, while Kant's morality requires each agent to endorse a single standard, governing both her helping and requests for help, it does not fix what this standard must be. And Kant does not then have a uniquely correct answer to the question how much

help to others is called for. Rather, his categorical imperative leaves open a wide range of answers — from the belief that help is called for only in genuine emergencies to a utilitarian policy prescribing help whenever the benefit to the helped exceeds the cost to the helper.²

John Rawls is a similarly inappropriate target for Gert's criticism. Endorsing the method of reflective equilibrium, Rawls leaves open the possibility that different people, even if raised in the same culture, may conscientiously arrive at different equilibria. He hopes, of course, that most of his compatriots will upon due reflection be able to endorse his theory of justice. But he also expresses considerable scepticism on this point: "I doubt, however, that the principles of justice (as I have defined them) will be the preferred conception on anything resembling a complete list... Even if the argument I have offered is sound, it only shows that a finally adequate theory (if such exists) will look more like the contract view than any of the other doctrines we have discussed" (Rawls 1971, 509)³ And he voices serious doubts about whether such a consensus can be extended beyond his own theory, which covers only the basic institutional order of a national society. For example: "It is too much to suppose that there exists for all or even most moral problems a reasonable solution. Perhaps only a few can be satisfactorily answered. In any case, social wisdom consists in framing institutions so that intractable difficulties do not often arise and in accepting the need for clear and simple principles" (Rawls 1971, 77-78).

Leaving Gert's ill-targeted criticisms to one side, let us consider his own view, or rather the view that he claims to find widely shared among his contemporaries. Gert calls this view Common Morality, but without clarifying the spatiotemporal reach of the word "common." On behalf of this Common Morality, Gert rejects, of course, the notion that every moral question has a uniquely correct answer. He formulates this denial as follows: "It is true that common morality is systematic enough that a computer could be programmed so that, provided with the facts of the case, it always comes up with acceptable moral answers. However, another computer could be programmed differently and still always come up with acceptable answers. There is no computer program that can tell you which of the competing computer programs you should adopt" (148-9).

Let us analyse Gert's view in set-theoretic terms. Gert holds that possible acts can be morally evaluated under five mutually exclusive labels: *prohibited*, *required*, *discouraged*, *encouraged*, *indifferent* (8). In Gert's caricaturistic description of his predecessors, they all sought to assign each and every possible act to one and only one of these labels. What Gert wants to do instead is to place each such assignment under one of two labels: *admissible* and *inadmissible*. In this way, Gert allows for greater flexibility of moral assessments: Instead of demanding that each act must be rigidly assigned to just one label, he allows that two or even three such assignments might be admissible. (For example, helping a poor woman find an affordable abortion clinic might admissibly be assigned the labels *discouraged*, *encouraged*, or *indifferent*.) Still, his account is rigid on the meta-level by counting each such assignment as being either admissible or inadmissible. For each and every moral judgment (moral evaluation of an act as: prohibited, required, discouraged, encouraged, or indifferent) there is a uniquely correct answer to the question whether this judgment is admissible or not.

We can make this same point somewhat differently by describing Gert as expanding the list of mutually exclusive moral labels from five to 31. For while, according to Gert, some acts fit under a single label exclusively, others fit under two or more. This yields a grand total of 31 theoretical possibilities: $2^5 - 1$. There are five possibilities of acts fitting under exactly one label: prohibited (P), required (R), discouraged (D), encouraged (E), indifferent (I). There are another ten possibilities of acts fitting under exactly two labels: (PR), (PD), (PE), (PI),

(RD), (RE), (RI), (DE), (DI), (EI). There are ten possibilities of acts fitting under exactly three labels: (PRD), (PRE), (PRI), (PDE), (PDI), (PEI), (RDE), (RDI), (REI), (DEI). There are five possibilities of acts fitting under exactly four labels: (PRDE), (PRDI), (PREI), (PDEI), (RDEI). And there is one possibility of acts fitting under all five labels: (PRDEI).⁴ Criticizing his predecessors for assuming that each act fits uniquely under one of five moral labels, Gert himself assumes that each act fits uniquely under one of 31 moral labels.

It is worth noting that Gert may here display, if anything, *greater* rigidity than, say, Kant. As we have seen, Kant allows for some flexibility in the assessment of moral judgments. Take the moral judgment that one may refuse to give certain help **H** to others. This moral judgment is admissible when combined with the moral judgment that it is prohibited to request such help from others, and inadmissible otherwise. On Kant's account, the admissibility of moral judgments is then decided holistically, while Gert (as I read him) assumes that the admissibility of any moral judgments can be decided in isolation from the others.

In any case, Gert's view exemplifies similar moral arrogance as he diagnoses in his predecessors by assuming that each act can be assigned to one of his 31 labels in a uniquely correct way. Equivalently, Gert presumes to lay down in advance what moral judgments are admissible. This moral arrogance is displayed also when Gert tightly constrains the parameters of admissible moral disagreement to a mere four: "(1) the scope of morality, such as whether foetuses are impartially protected, or protected at all, by morality; (2) the rankings of the various goods and evils, such as whether a specified loss of freedom is worse than a specified risk of death; (3) the harmful and beneficial consequences of a violation being publicly allowed, when these are based on ideological views that are not open to empirical investigation; and (4) the interpretation of a moral rule, such as whether turning off the respirator of a ventilator-dependent patient who has refused further treatment counts as killing him. ... These disagreements are responsible for *all* of the morally unresolvable questions" (140; my emphasis). By monologically legislating in advance tight constraints on reasonable disagreement, Gert thus substitutes a somewhat modified moral arrogance (meta-arrogance) for the classical variety he diagnoses in his predecessors and seeks to overcome.

How consequential this modification is depends in part on what percentage (roughly speaking) of all moral judgments Gert would rule inadmissible. Here Gert's rhetoric strongly suggests that he would disqualify the vast majority of moral judgments as inadmissible. He believes that his Common Morality reflects "overwhelming agreement on most moral matters" (148), and that 99 percent of all acts fit under one of the five traditional labels in a uniquely correct way. The modification is thus, by Gert's own account, relatively minor.

Let us now look at the postulated core of Gert's Common Morality and, in particular, at the moral judgments which supposedly we all agree on (each of which uniquely assigns some act to one of the five primary moral labels). Here two immediately pressing questions are: Who are "we"? And what is the status of the claim Gert is advancing?

On one reading, "we" comprises all "moral agents" (20). This expression is defined and used in various conflicting ways throughout the book. According to one definition, agents are moral agents just in case they are "subject to moral judgment" (87). According to another definition, agents are moral agents just in case they "know [the moral rules] and can guide their conduct accordingly" (26). A third definition says that "moral agents are rational persons with sufficient knowledge and intelligence to know what the moral rules prohibit and require, and who have the ability to guide their behaviour by their knowledge" (91). Leaving

aside the substantial discrepancies between these definitions, there is an obvious problem of circularity in the last two proposals: If moral agents are defined as those who agree with Gert's favoured rules of Common Morality, then it is not remarkable that they all agree with these rules.

Gert is more charitably read, I believe, as making an *empirical* claim about the moral judgments ordinary people actually make or would make if queried. But how far is this set of ordinary people meant to extend in space and time? Sometimes it sounds as though Gert means to make a very expansive claim that covers all of humankind across much of human history. Thus he writes in the Preface: "Now that the people with the most diverse cultural and religious views must interact with each other, it is important for everyone to be aware of their common morality" (x). Now if Gert really means to make an empirical claim about the moral judgments made by a large and diverse population of past and present human beings, then it is astonishing that he makes no effort of any kind to support his claim. There are dozens of passages assuring the reader that *everyone* would make this, and *no one* would make that other judgment, that *almost no one* would say such a thing, and so on. Yet there is never the slightest attempt to support any of these assertions by looking at empirical work (of which there exists a great deal!) about the moralities of different cultures and subcultures.

To be sure, I have not looked carefully at such empirical work either, and so am not in a good position to decide whether Gert's claims about his Common Morality are true or false. But I suspect that they are false in many important respects. I suspect this on the basis of extensive travels to developing countries, such as China, where people seem to think quite differently about interpersonal morality — though they may soon be sucked into the Western cultural orbit. But I suspect this even on the basis of my familiarity with my own culture. And I will proceed to give some examples from here, as these should make it easier for us to judge whether Gert's Common Morality is really common.

Gert holds that there are ten and only "ten general moral rules [that] account for all of the kinds of actions that are morally prohibited and required" (20). These rules are (20):

- Do not kill.
- Do not cause pain.
- Do not disable.
- Do not deprive of freedom.
- Do not deprive of pleasure.
- Do not deceive.
- Keep your promises.
- Do not cheat.
- Obey the law.
- Do your duty.

Is this really a complete account of the morality human beings hold in common?

Gert's account completely leaves out a traditionally hugely important part of morality focused on sexual relations. Contrary to Gert's bald assertion that "moral rules always concern our behaviour toward others" (45), it used to be widely believed that masturbation is morally prohibited. The condemnation of other kinds of sexual deviance is even stronger, and ongoing. Few behaviours elicit as intense and widespread horror and disgust as bestiality, incest, and sodomy, for example, or the culinary use of human flesh. But Gert's Common Morality finds such behaviours acceptable so long as they do not cause pain, which they

certainly need not do. (Sex with one's mother or a donkey can surely be effected without causing pain, and to eat human beings one need not accelerate their death.)

Faced with this attempted counterexample, Gert might say — as he in fact writes about polygamy (49) — that bestiality, incest, sodomy, and cannibalism are morally prohibited (under his ninth rule) because and insofar as they are legally prohibited. But this response would be doubly problematic. First, very few people believe that sex with one's mother or with animals, or cannibalism, are morally permissible wherever they are not illegal. Second, this response leaves unclear why Gert's moral critique should target polygamists, cannibals, or sexual deviants rather than the law that renders their behaviour illegal. If there were a law prohibiting sucking one's thumb or using the word "balcony," then Gert would surely find fault with this law and call for its repeal. It should be legal to use the word "balcony," he would say, with universal approval. But any call for the legalization of bestiality, incest, sodomy, or cannibalism would be widely condemned on the ground that such deeply immoral conduct must not be tolerated by the law. I conclude then that there are forms of conduct that are widely judged to be morally prohibited but not proscribed by any of Gert's rules.

Another counterexample showing underinclusiveness involves the prohibition on spreading embarrassing truths about others behind their back. One can engage in this conduct without violating any of Gert's ten rules, and such conduct would be widely condemned.

A third and last counterexample is more complex and aims to show over-inclusiveness. Gert's rules, as he explicates and discusses them, state prohibitions that are insensitive to the degree of compliance these rules actually elicit in the agent's social environment. Thus, failing to keep a promise is prohibited regardless of how frequently promises are kept in the agent's vicinity and regardless of the extent to which the promisee has fulfilled his or her promises. Here again most actual people would see things differently. They would not consider it morally wrong to break a promise made to someone who notoriously breaks her own promises; and they would consider it less wrong to cheat others in a social environment in which cheating is widespread than in one in which cheating is very rare — especially in situations where compliance with moral rules would make one a "sucker."⁵

Gert's *Common Morality* is elegant and highly ambitious. I hope that the criticisms I have formulated here reflect my respect for Gert's work and offer some measure of reciprocation for what I have learned from it.

Notes

¹ Kant says just this, but unfortunately it gets mistranslated. The point is not that all my maxims must harmonise *with* a possible realm of ends – as if such a realm of ends were somehow independently given.

² Cf. Thomas Pogge: "The Categorical Imperative" reprinted in Paul Guyer, ed.: *Kant's Groundwork of the Metaphysics of Morals* (Totowa: Rowman and Littlefield 1998), 189-213.

³ Cf. also 46 and 267.

⁴ Of course, Gert is not committed to the view that each of these labels is actually instantiated. There may be no act that can admissibly be labelled *encouraged* and *prohibited* and nothing else.

⁵ This widely recognized "sucker exemption" is extensively discussed in section 2.2 of my essay "Historical Injustice: The Other Two Domains" in Lukas Meyer, ed. *Justice in Time. Responding to Historical Injustice* (Baden-Baden: NOMOS 2004), 117-134.

Reply to Thomas Pogge's: "Timeless Wisdom or Moral Arrogance?"

Bernard Gert
Dartmouth College

Thomas Pogge, like Dean Cocking, is reacting not only to *Common Morality*, but also to the paper "Moral Arrogance and Moral Theories" that I presented at a CAPPE seminar at the Australian National University. Indeed, Pogge's remarks during that seminar led me to limit my criticism of Kant, and I now claim only that Kant holds that there is a unique answer to all moral questions about what he calls duties of perfect obligation and I call morally prohibited or required actions.¹ With regard to imperfect duties or about what I call morally encouraged actions, Kant may hold that each individual decides for himself which maxims he can consistently and coherently hold. In *Common Morality* I admit that Kant allows people to disagree about when to do actions that are imperfect duties. However, Pogge does not dispute that for Kant everyone must agree about perfect duties or about what actions are morally prohibited or required.

However, I have some lingering doubts about whether Kant can allow purely rational beings to disagree about anything. As I understand Kant's account of purely rational beings, they do not differ from each other in any way. This is the way that Rawls understands Kant, when he says, "My suggestion is that we think of the original position as the point of view from which noumenal selves see the world." (Rawls 1971, 255) In the original position, Rawls claims, "each is convinced by the same arguments. Therefore, we can view the choice in the original position from the standpoint of one person selected at random." (Rawls 1971, 139) This strongly suggests that, according to Rawls, Kant views all noumenal selves as seeing the world in exactly the same way. If that is the case, then it is not clear that Kant can allow for disagreement even with regard to imperfect duties. However Pogge's account of Kant is so plausible that I am inclined to think that Rawls is wrong or that the maxims that are put forward are not those of purely rational beings. Thus, my criticism of Kant is now limited to his view that everyone agrees about the maxims that yield perfect duties.

Since Rawls was a teacher of mine in the late 1950's, and we exchanged manuscripts of *A Theory of Justice* and *The Moral Rules: A New Rational Foundation for Morality*, in the late 1960's, I am not equally inclined to accept Pogge's authority as an interpreter of Rawls, especially of the original version (1971) of *A Theory of Justice*.² However, I do not want to get into a battle of quotes with Pogge, even though the quotes I offer in the previous paragraph seem to make it clear that Rawls cannot, without inconsistency, deny that all equally informed rational persons will make the same moral judgments. Rawls was a modest person, and he probably was genuinely doubtful that his theory of justice was completely adequate. Since his theory of morality was far less developed, he was probably even more doubtful about its adequacy. He probably sincerely held that attempting to attain complete reflective equilibrium would reveal serious problems in his theories. However, that is all irrelevant with regard to the theory that he actually put forward. On that theory, all equally informed rational persons in the original position make the same moral judgments. Pogge says nothing that conflicts with this criticism of Rawls's theory.³

Pogge labels my description of most previous moral theories as caricatures, but that may be because he accepts my view that it is absurd to hold that there are unique answers to every moral question. Pogge actually addresses real world moral problems, so that it is not

surprising that he accepts my view. However, many, if not most, moral philosophers do not attempt to apply their theories to any real world problems. These philosophers do not think it is a caricature of a moral theory to describe it as offering unique correct answers to every moral problem. Indeed, some anthologies in ethics explicitly claim that a moral theory must provide a unique correct answer to every moral question in order to be adequate. Rawls at an early stage of his career clearly held that one of the strong points of consequentialism was that it offered a decision procedure for settling all moral questions, and *A Theory of Justice* was an attempt to provide a non-consequentialist decision procedure.

Pogge claims that I hold “that possible acts can be morally evaluated under five mutually exclusive labels: prohibited, required, discouraged, encouraged, indifferent.” This is not quite what I say. The first and most straightforward mistake is substituting the word “indifferent” for the word that I use, which is “allowed.” This mistake is interesting, for “indifferent” is the word commonly used by those consequentialists who hold that every moral question has a unique correct answer, namely, prohibited (wrong), required (right), or indifferent.⁴ If two acts are equal in whatever kinds of consequences the particular theorist claims are the kind of consequences that determine what counts as the best consequences, then it is morally indifferent which act is done. But, contrary to the consequentialists, although there may be agreement on the kinds of consequences that count, two equally informed impartial rational persons need not rank these consequences in the same way. To use indifferent as a basic category suggests that all equally informed impartial rational persons agree in their ranking of these consequences, which is why I do not use indifferent as one of my five basic categories.

The second and more important misleading claim is that these five labels are “mutually exclusive.” This strongly suggests that an act for which one of these labels is admissible can never be an act for which another one of these labels is also admissible. But Pogge realizes that on my account, people can sometimes differ about which label to apply to an action. So he puts his point in a slightly different way. He claims that I hold that there is a unique correct answer, namely admissible or inadmissible, for whether a label can be applied to a given act. I do claim that some decisions or judgments are such that no informed impartial rational person can apply that label to the act in question, and other decisions or judgments are such that all informed impartial rational persons can apply that label to the act in question. For example, it is inadmissible to hold that killing a person just for the fun of it can be labeled anything other than prohibited.

Pogge is aware that for many cases, I allow several different labels to be admissible. He constructs 31 possibilities or combinations of labels that theoretically could be regarded as admissible. He realizes that this allows for people to disagree quite significantly about which label to apply, but accuses me of greater rigidity than some whom I criticize, because I hold that all informed rational persons will agree on what labels are admissible and inadmissible. But I do not hold there is complete agreement about what equally informed rational persons would regard as admissible or inadmissible. Rather, as a methodological strategy, I count as admissible any moral decision or judgment that any significant number of otherwise rational persons could make under the blindfold of justice. I do not want to rule out as an admissible moral judgment any moral judgment that any significant number of otherwise rational persons could make under the blindfold of justice. I fail to see any objectionable rigidity or moral arrogance in this view.

Pogge claims, “Kant allows for some flexibility in the assessment of moral judgments. Take the moral judgment that one may refuse to give certain help **H** to others.

This moral judgment is admissible when combined with the moral judgment that it is prohibited to request such help from others and inadmissible otherwise.⁵ On Kant's account, the admissibility of moral judgments is then decided holistically." This suggests that Kant does not put any objective limits on what maxim a person might put forward; rather he is concerned only with consistency. Pogge makes this point quite clearly in the following paragraph:

"The maxim governing one's requests for help from others must not be more expansive than the maxim governing one's own helping of others. If one refuses to give help H (of a certain sort, in certain circumstances in which one is not naturally, perhaps from sympathy, inclined to help), then one may not request such help for oneself. For in the world in which everyone felt free to refuse such requests, they would be refused and the end of receiving such help could thus not be attained and the maxim of requesting it could then not be willed⁶ This thought experiment shows that a maxim may be rejected because its universal availability would defeat the material end — not of this maxim, but — of some other maxim to the availability of which the agent is committed. The thought experiment does not show which of the incompatible maxims the agent must give up. So, while Kant's morality requires each agent to endorse a single standard, governing both her helping and requests for help, it does not fix what this standard must be."

Although Pogge may have correctly formulated Kant's formal approach to determining what is an admissible maxim, in the remainder of this paragraph, he does seem to accept this completely formal approach. The final two sentences of the paragraph go as follows. "And Kant does not then have a uniquely correct answer to the question how much help to others is called for. Rather, his categorical imperative leaves open a wide range of answers — from the belief that help is called for only in genuine emergencies to a utilitarian policy prescribing help whenever the benefit to the helped exceeds the cost to the helper." If Pogge accepted Kant's purely formal approach, he would not limit the wide range of answers in the way that he does. Why can't a rational person hold that help is not called for even in a genuine emergency? Why can't a rational person hold that help is called for even when the cost to the helper exceeds the benefit to the helped? Pogge limits the range of admissible answers because he correctly holds that it would be irrational to hold that help is not called for even in genuine emergencies, and that it would also be irrational to hold that help is called for even when the cost to the helper exceeds the benefit to the helped. If he does hold these correct views, then he rejects Kant's purely formal view and accepts my view that certain moral judgments are inadmissible.

Pogge claims that I also display moral arrogance when I "tightly constrain[s] the parameters of admissible moral disagreement to a mere four." If I do, it is certainly less than the moral arrogance of those who not allow any admissible moral disagreement. But I do not see any moral arrogance in maintaining that the four parameters I list exhaust the parameters of unresolvable moral disagreement among equally informed impartial rational persons. Pogge does not even try to offer a fifth parameter, so it may be that he is simply objecting to my explicitly putting content into my moral theory instead of claiming, as philosophers are wont to do, that everything can be decided in some formal way.

Pogge claims that there exists a great deal of empirical work "about the moralities of different cultures and subcultures." It is interesting that Pogge does not seem to think that there is any problem in determining what counts as "the moralities of different cultures and subcultures." He seems to think that morality can be defined as something like "the code of conduct that a culture or subculture adopts to guide the behaviour of its members." In the larger versions of my moral theory, I discuss this kind of societal definition of morality, as

well as other commonly offered definitions.⁷ I do not explicitly discuss the definition of morality in *Common Morality*, so that it is understandable if Pogge takes me to be accepting some kind of sociological or formal definition of morality, similar to the kind of definition that he seems to accept. I do not accept such a definition. I take as a defining condition of morality that it applies to all normal human adults, and that all of them know what kinds of actions morality prohibits, requires, discourages, encourages, and allows. For me, as for Hobbes, Kant, Mill, and Rawls, morality is universal and is to be distinguished from those local customs that a given society or religion might mistakenly lump together with it. When Pogge says that I leave out “a traditionally hugely important part of morality focused on sexual relations”, it is these local customs that he is talking about.

A well-known example of the widespread confusion caused by the failure to distinguish non-moral guides from the moral guide is The Ten Commandments. The Ten Commandments were originally intended solely for those who accepted Judaism as their religion, not as a moral guide to be followed by all persons. However, for various reasons they came to be regarded as a moral code, even though some of the commandments, such as prohibiting the making of graven images and requiring one to keep the Sabbath holy, have nothing to do with morality at all. Even more disturbing, slavery is not condemned by The Ten Commandments, but rather is accepted, with a mild attempt to make it less harsh. Most people who accept The Ten Commandments as the moral law do not even know what these commandments say. In giving lectures on this topic, I often ask, “How many of you think that The Ten Commandments should always be followed?” About 95% of the audience raises their hands. Then I ask, “How many of you know what The Ten Commandments say?” and less than 5% of the audience raises their hands.

To take as relevant the fact that “it used to be widely believed that masturbation is morally prohibited,” indicates a failure to distinguish morality from religion. It also used to be widely believed that homosexual behaviour is morally prohibited, but now almost every college and university in America prohibits discrimination against homosexuals. Sexual matters lead many people to hold inconsistent views, not just about morality, but about other matters as well. For many years the *Diagnostic and Statistical Manual of Mental Disorders* (DSM) of the American Psychiatric Association stated that people who engaged in deviant sexual behaviour had mental disorders, regardless of whether they suffered any harm. However, in 1987, DSM-III-R adopted a definition of mental disorder such that no one had a mental disorder unless they had a condition that was the cause of their suffering an evil or of being at increased risk of suffering such an evil. This definition was included in all subsequent DSM's, but deviancy was still accepted as sufficient to label some sexual deviations as mental disorders. These DSM's were inconsistent.⁸

Pogge is probably correct that “Few behaviours elicit as strong and widespread horror and disgust as bestiality, incest, and sodomy,” at least among those who do not participate in such activities. However, even if this claim is true, I do not know what is proved by it, unless Pogge equate behaviours that elicit strong and widespread horror and disgust with immoral behaviour. But if, as I do, one takes morality to be essentially related to rationality, then bestiality, incest, and sodomy would have to be prohibited by all impartial rational persons for these behaviours to count as immoral. But whether all impartial rational persons would prohibit bestiality, incest, and sodomy depends on what they thought would be the consequences of these activities. If they believed that no one would be hurt, either physically or emotionally, and that no essential institutions such as the family would be damaged, then they need not prohibit such activities. But, of course, many believe that there would be harmful consequences, either directly or indirectly from such activities. Whether rational

persons, under the blindfold of justice, regard such behaviour as immoral depends on their believing these activities to have harmful consequences. Without such beliefs, only those who mistakenly base morality on their feelings and emotions or their religion would prohibit such activities.

I shall conclude my reply to Pogge by discussing the three counterexamples that he mentions. Regarding his first counterexample, Pogge asks why I regard the law prohibiting polygamy as a justified law when I would not regard a law prohibiting sucking your thumb or using the word, “balcony” as justified. If some other word were substituted for the word “balcony,” for example, a word like “fuck,” which many people in the society find offensive, then it might be justified to have a law prohibiting using that word. But the word “balcony” is not an offensive word, so that a law prohibiting its use is simply bizarre. Similarly, sucking your thumb does not bother many or any other people, but a law might justifiably prohibit some behaviour that does offend a great number of people, such as going out in public with no clothes on. This makes it clear why it might be justified to prohibit behaviours that “elicit as strong and widespread horror and disgust as bestiality, incest, and sodomy... or the culinary use of human flesh.”

However, although I think it might be justifiable to have laws prohibiting bestiality and sodomy if they were done in public, I am much less convinced that it is justifiable to have such laws if they prohibit such behaviours if they are done privately and if there is no reasonable expectation that such private behaviour would even increase the probability of harmful consequences for anyone. Many think that there is no reasonable expectation that private acts of bestiality or sodomy do increase the probability of harm for anyone, but they may hold that simply knowing that incest was allowed might affect the family structure in harmful ways.

It is not clear what Pogge means by “cannibalism.” Normally “cannibalism” involves killing a person in order to eat their flesh, but this cannot be what Pogge means by “cannibalism,” for he says, “to eat human beings one need not accelerate their death.” So he must simply mean by “cannibalism,” “the culinary use of human flesh” after a person has died in some way completely unrelated to cannibalism, for example, after he is killed in an automobile accident. Pogge is correct that even this kind of cannibalism is viewed with widespread horror and disgust. People who are in a situation where they are starving, and in which there is no other food available, are still extremely reluctant to eat the flesh of a person who died. But if they did not kill the person, it is not immoral for them to eat the flesh of that person although I can understand why they would be reluctant to do so.

Nonetheless, a law prohibiting “the culinary use of human flesh,” even after a person is killed in an automobile accident, can be justified by the fact that its being legally allowed would cause strong and widespread horror and disgust. In addition to this justification for a law prohibiting cannibalism, it is plausible to hold that publicly allowing the eating of human flesh would lessen the respect that everyone wants people to have with regard to human life. It is not that human flesh tastes so good that it provides another motive for killing people, but that eating human flesh is treating dead human beings like we treat other dead animals, and so lessens the psychological force with which the rule against killing human beings is now invested.

With regard to his second counterexample, I do not understand why Pogge thinks that “spreading embarrassing truths about others behind their back” does not violate any of the moral rules. Each of these rules not only prohibits causing harm, but also doing those actions

that significantly increase the probability of people suffering harm. “Spreading embarrassing truths about others behind their back” clearly significantly increases the probability of suffering harm for the persons about whom you are spreading these embarrassing truths. This explains why “such conduct would be widely condemned.”

Talk about acts that are morally prohibited should not be talk about unrealistic particular instances of such acts tailored so that we know that they have no bad consequences. Moral rules prohibit the kinds of acts that cause harm or increase the probability of someone suffering harm. It is immoral to do any act of this kind unless one has an adequate justification for doing it. Because morality applies to fallible people, it is not an adequate justification that you justifiably believe that a particular act of this kind will not have bad consequences. If that were an adequate justification then a student could justifiably cheat on some pass-fail exam in a required philosophy course if he was sure that he would not be caught.

Pogge’s third counterexample is the only example that is not an example of my having left out some kind of behaviour that should be morally prohibited. This example is supposed to show that my moral rules prohibit some behaviour that most people would not consider to be immoral. Pogge says that as I explicate and discuss the moral rules, they “state prohibitions that are insensitive to the degree of compliance these rules actually elicit in the agent’s social environment.” It is true that I separate my discussion of the moral rules from my discussion of how one determines whether a violation of these rules is justified. Pogge may be accustomed to philosophical discussions of moral rules in which all of the discussion is concentrated on the rules themselves. I think it preferable to separate discussion of the rules from discussion of that part of the moral system that is used to determine when it is justifiable to break the rule. It is in that discussion, which includes a list of the morally relevant features of the violation that “the degree of compliance these rules actually elicit in the agent’s social environment” would occur, insofar as the different social environments affected the consequences of breaking the rule.

Pogge claims that many “would not consider it morally wrong to break a promise made to someone who notoriously breaks her own promises; and they would consider it less wrong to cheat others in a social environment in which cheating is widespread than in one in which cheating is very rare — especially in situations where fulfilment of one’s promise would make one a ‘sucker’.” Pogge may be correct about people’s views about promise breaking. However, if someone held this view, then she would be willing for everyone to know that anyone who makes a promise “to someone who notoriously breaks her own promises” is allowed to break that promise, “especially in situations where compliance with moral rules would make one a ‘sucker’.” If this is true, then this counterexample is not a counterexample at all, for my theory does not require people to keep such promises. But the question arises, why would one make such a promise, if everyone knows that they are allowed to break it? Perhaps, it is advantageous to make it, but one may hold, like some act utilitarians, that the fact that one made the promise is not considered, by itself, a reason to keep it; the only reason for keeping it is that it may turn out to be advantageous to keep it.

I am not sure what Pogge means by saying, “and they would consider it less wrong to cheat others in a social environment in which cheating is widespread than in one in which cheating is very rare. I do not know what it is for an action to be “less wrong,” any more than I know what it means for one pregnant woman to be less pregnant than another. One woman can be farther along in her pregnancy, and one wrong action can be “less serious” than another, but that is a different matter. In an environment in which cheating is rare, cheating

may have more harmful consequences than an environment in which cheating is widespread, so it is taken more seriously. But in both cases, the cheating is morally wrong.

I am uneasy about this last counterexample because it seems to me so close to the view that it is admissible to violate a moral rule “because everyone is doing it”. This seems to me to be an inadmissible view. With regard to violations of a moral rule, the claim that “everyone is doing it” is always false. To use that claim as one’s sole reason for violating the rule oneself is not admissible because if one does not think the consequences would be better if everyone knew that they were allowed to break the rule in those circumstances, then obviously one is making a special exception for oneself. One is not being impartial. Pogge, in the beginning of his paper makes clear that Kant requires impartiality and rules out as unacceptable making special exceptions for oneself.

I am pleased that Pogge took the risk of criticizing my views even though he did not feel confident about doing so. I found his criticisms quite challenging and they forced me to be clearer than I had been about various features of my theory. They also made me realize that the theory needs even more clarification if it is to be persuasive to those whom I most want to persuade, namely, philosophers like Pogge.

Notes

¹ See the revised version of “Moral Arrogance and Moral Theories” in *Noûs* Supplementary Volume 15 on Applied Ethics.

² Pogge’s quotes from Rawls seem to come from the 1999 edition, which differs from the 1971 edition in some significant ways,

³ For those interested in further support for my criticism of *A Theory of Justice*, see Chapter 13, “Versions of Morality as Impartial Rationality”, in *Morality: A New Justification of the Moral Rules*, (Oxford University Press, 1988)

⁴ Such consequentialists do not have the categories of discouraged or encouraged. However, other consequentialists who distinguish between acts that have both (1) the best consequences and (2) would have the best consequences if non-performance of those acts were punished, and acts for which (1) is true, but (2) is false, can have the categories of discouraged and encouraged. But such consequentialists would still have the category of indifferent, if two acts had the same quality of consequences.

⁵ I do not agree that the moral judgment “it is permissible to refuse to give certain help **H**” is admissible only if “combined with the moral judgment that it is prohibited to request such help from others.” I see nothing incompatible with requesting aid and at the same time holding that it is permissible to refuse that request.

⁶ Pogge’s assumption that “in the world in which everyone felt free to refuse such requests, they would be refused” seems false and so his conclusion “and the end of receiving such help could thus not be attained and the maxim of requesting it could then not be willed.” This assumption may explain why Pogge makes the claim discussed in the previous note.

⁷ See the first chapter of several of the full-length accounts of my moral theory: *Morality: A New Justification of the Moral Rules*, (Oxford University Press, 1988); *Morality: Its Nature and Justification*, (Oxford University Press, 1998); *Morality: Its Nature and Justification*, Revised Edition, (Oxford University Press, 2005). I have also written the online entry of the “Definition of Morality” for the *Stanford Encyclopedia of Philosophy*.

⁸ For a full account of this problem, see “A Sex Caused Inconsistency in DSM-III-R: The Definition of Mental Disorder and the Definition of Paraphilias,” *The Journal of Medicine and Philosophy*, Vol. 17, No. 2, April 1992, pp. 155-171

Discussion

Government and Faith-based Organisations in a Pluralist Society

Bruce Langtry

University of Melbourne

1. Neutrality towards religious outlooks

In this paper I sketch three principles, with supporting rationales, that should guide the government of a pluralist society such as our own, in its dealing with faith-based organisations that it funds to provide educational or welfare services, such as foster-parent programs for children.

I begin indirectly: my first principle enjoins a specific sort of religious neutrality. The kind I have in mind is not neutrality between religions because I intend to characterise the neutrality via the concept of being friendly to liberal democracy, and most contemporary religions (such as Christianity and Islam) cannot sensibly be said either to be friendly or to be unfriendly to liberal democracy. For example, while Christian thinking in Australia today tends to be broadly supportive of the foregoing values, there have been other times and places where Christian thinking has been, by and large, opposed to them; the same is true of Catholic thinking, Presbyterian thinking, etc. So to formulate the target principle, I need to introduce a technical idea religious outlook, and to distinguish between two different kinds of religious outlook. The principle formulated in these terms will have a secondary but important application to how the government should treat faith-based organisations.

The concept of a religious outlook is vague, but it is not necessary to make it more precise for the purposes of this paper. What I have in mind are different combinations of theological, moral and political principles arising within a single religion, individuated in a medium-grained way. For example, the set of beliefs held in common by evangelical Anglicans who are liberal democrats counts as a religious outlook.

The distinction I wish to draw is between religious outlooks that are friendly to the fundamental political principles characteristic of liberal democracy, and those that are hostile to, or knowingly subversive of, those principles. What principles? They include equal basic civil and political rights, such as those bound up with freedom of thought, freedom of speech, and freedom, within broad limits, to choose one's own lifestyle; they include the power of the people to effectively hold governments accountable and remove them from office.

The government ought to tolerate religious outlooks which are hostile to or knowingly subversive of liberal democracy, as long as their adherents do not violate the law, including those laws which uphold other people's civil and political rights. Amongst their adherents' civil and political rights are rights to hold and proclaim illiberal and anti-democratic views. Furthermore, the government and its agencies should not discriminate against individuals with respect to employment, welfare services and so on, merely on the grounds that they hold and proclaim such views. There will be exceptions, perhaps, but these will need very careful thought. After all, even a person who favours racist laws and practices might affirm and

conform to a professional ethic according to which one's personal views on such matters should not, given the laws and conventions that actually prevail today, influence the way one makes decisions in one's role as a public servant.

Similarly the government should not discriminate against organisations that are corporately committed to 'unfriendly' religious outlooks, with respect to matters like taxation, building permits, or basic services such as garbage removal.

Nevertheless, I doubt that the government has a general duty to treat organisations corporately committed to 'hostile' religious outlooks on the same basis as it treats organisations holding friendly outlooks. Nor need the government make the same effort to avoid offending the religious beliefs and feelings of groups holding hostile outlooks as it makes to avoid offending the religious beliefs and feelings of groups holding friendly outlooks.

Suppose that there are three faith-based organisations O, P and Q tendering to supply welfare services to some group of immigrants; organisation O is the only one committed to a 'hostile' religious outlook, but it is the one that best understands the needs of the immigrants, can best communicate with them, and so on. What should the government do? The neutrality principle I have advanced does not determine the answer. The general permissibility of discriminating against O in virtue of the 'hostile' outlook must be balanced against the fine-grained nature of government's moral duties to care for members of the immigrant group. The government will also need to take into account its duty to abide by any published rules of procedural fairness that it finds in place.

Obviously a lot more needs to be said about the scope and limits of permissible government discrimination against organisations committed to 'hostile' religious outlooks, but developing a theory is a big project that I must defer until some later time.

Let us now leave aside the hostile views, and consider only religious outlooks that are friendly to liberal-democratic political principles. There are quite a lot of these. The neutrality principle I am proposing is this: the state should not do anything with the intention of favouring any friendly religious outlook over any other friendly religious outlook. For example, the government might intentionally give money to a Christian school in Brunswick, in preference to a Muslim school in Coburg, but not in order to favour a friendly Christian outlook over a friendly Muslim one. (One favours outlook A over B if one favours some individual or group over another wholly or partly on the grounds that the favoured individual or group subscribes to A and the other subscribes to B.)

Suppose that the government issues a state school curriculum one of whose aims is to encourage students to believe the theory of biological evolution. Suppose also that friendly religious outlook A affirms, or at least is consistent with, the theory whereas friendly religious outlook B rejects it. It does not follow that the government intends to favour A over B. The government may not have thought about A or B at all. If the government has noticed that in one respect its curriculum advantages A over B in the marketplace of ideas, it need not have intended this. And, a fortiori, it need not have intended that its curriculum gives an overall advantage to A over B.

The neutrality principle should be accepted because it fits in with basic values and principles of liberal democracy. The people who live in Australia today are deeply divided on religious, philosophical and moral issues. Despite this fact, we might hope that liberal

democracy in this country is viable in the long term: that a liberal-democratic constitution and social structure might attract lasting support from a broad spectrum of religious and moral standpoints. Not from every standpoint, perhaps, but from enough. If that is to be so, then people from a wide variety of religious outlooks must see Australian laws and regulations as a fair basis for social co-operation - a basis that is fair to all outlooks that are friendly to liberal democracy. It is plausible to suppose that they will be greatly encouraged to see things this way if the constitution and political practice embody the neutrality principle, and will be otherwise discouraged.

State compliance with the neutrality principle may still leave adherents of many friendly religious outlooks feeling that Australian laws and regulations, and the arguments that governments advance in their favour, are subversive of metaphysical beliefs and values that are central to their position. This is a limitation but not a defect in the neutrality principle. Adherents of friendly outlooks ought to be able to see that in any liberal society it will be impossible to ensure that no friendly position is weakened as a result of government policy. So adherents of this or that damaged friendly outlook may well classify the damage to their position as regrettable but not a sign that the system is unfair.

2. Respect for corporate identities, missions, and values

The second principle that should guide the government is in dealing with 'friendly' faith-based organisations providing educational and social welfare services is that it should be respectful of the organisations' identities, missions and values, including the religious aspects of these.

Let me begin with some substantial assumptions about moral duties owed by and to individuals and organisations. Individuals' moral duties include civic duties. Although these civic duties are not requirements of law, and although they can often be formulated only vaguely and with lots of escape clauses, their widespread fulfillment is vital to society. They include a duty to treat people of different ethnic and religious backgrounds with respect and courtesy. They include a moral duty to contribute to the good of the community, over and above what is legally required. Organisations, such as sporting clubs, schools, churches and businesses, have similar civic duties. For the survival and prosperity of most such bodies depend on relationships of mutual trust, fairness, co-operation and unselfishness both within the organisation and also with outsiders. Hence corporate civic duties are in part duties of gratitude: the organisation has been nourished by the goodwill of individuals and other community bodies, acting from motives other than their own commercial advantage, and so it should be prepared to benefit society at some cost to itself.

The weight and scope of an organisation's civic duties depend partly on whether its aims are narrowly focused or broadly focused. The aims of tennis clubs and corner stores are in most cases narrowly focused. In contrast, those of schools are typically broadly focused: their aims extend beyond the intellectual development of its students, and include their physical, social and moral development. Accordingly, we expect a school to have a much deeper sense of civic responsibility than a tennis club or a corner store.

Likewise the weight and scope of civic duties that individuals and groups have to act generously to help sustain and assist other associations - e.g., by advertising meetings, lending equipment, providing advice etc - depend partly on whether the aims of these other associations are narrowly focused or broadly focused, and so on the weight and scope of their civic duties. There is a kind of reciprocity here. Leaving aside emergencies and special ties,

we are not talking about duties to assist any particular community organisation in any specific way, but rather what philosophers call imperfect duties, duties involving a good deal of discretion about when, to whom and how one should make a public-spirited contribution.

Federal and state governments possess civic duties of the kind I have been talking about. For one thing, government programmes typically depend partly on the goodwill and voluntary co-operation of individual citizens and community groups, and so the government acquires duties of gratitude. Even more fundamentally, the government's core mission is to promote Australia's well-being, which plainly requires a profusion of voluntary associations pursuing private and public goods. The government does not have to adopt or endorse the aims and values of all these organisations, even those holding religious outlooks friendly to liberal democracy – for one thing, the aims and values are still sometimes incompatible with those of other bodies, or are directly contrary to what the government believes to be in the common interest. But it does have to create and sustain conditions favourable to the flourishing of a rich array of voluntary associations, and to be respectful of their corporate identities, missions and values. These duties are especially strong towards organisations with which the government itself is engaged in voluntary co-operative activity.

Such co-operative activity includes contractual arrangements with both commercial firms and not-for-profit organisations that provide services to it or on its behalf. The relevant aims of most not-for-profit organisations, however, are much more broadly focused than those of most businesses, and the civic duties owed by them and to them are correspondingly more extensive and more important than those owed by and to commercial operators.¹ The government should treat both schools and cleaning firms with civility, but both it and the entire community require much more of schools and so it should be much more respectful and considerate of schools.

In what ways should a government be respectful of the corporate identities, missions and values of organisations whose aims and activities it believes are contrary to the common interest – and, indeed, organisations that are unfriendly to liberal democracy?

Moral disagreement and political disagreement do not typically suffice to morally justify contempt towards the other individual. Nor does the fact that the individual has committed serious moral wrongs. (Think about a variety of cases involving fighting in a war.) The relevant attitude of respect typically involves, as a bare minimum, a recognition that the organisations and their members are not morally vicious and are not in disgrace. It also involves, very frequently, treating the objectionable aspects of the organisations' aims and activities as irrelevant to a large set of decisions that public and private bodies officials make with respect to them.

3. Discrimination between faith-based organisations

I have already pointed out that some faith-based organisations might be friendly to the fundamental political principles characteristic of liberal democracy and yet have aims, values and political views that are directly contrary to what the government believes to be in the common interest. (After all, the government of Australia usually thinks that the political party or coalition constituting the Opposition fits both these conditions!) Obviously a faith-based agency that counsels and trains unemployed people might be liberal-democratic in orientation but be strongly opposed to the government's labour market policies. The government might regard the views and public statements of this agency as making it an unsuitable partner for the implementation of employment-related programmes, and therefore exclude it from

consideration as a possible state-funded participant in these programmes. So it might prefer some other faith-based agency instead.

Such discrimination is consistent with the neutrality principle I have stated, to the effect that the state should not do anything with the intention of favouring any friendly religious outlook over any other one. For to discriminate between faith-based organisations on the basis of their aims, values and political beliefs need not involve favouring one religious outlook over another. After all, individuals and organisations that share the same religious outlook can differ about important political issues e.g., those bound up with the war in Iraq, or logging old growth forests in Tasmania. Such discrimination is also consistent with the respect I talked about in the previous section.

Nevertheless, acknowledging the consistency does not settle all the philosophical issues that might be worrying us at this point. Specifically, there is an issue that applies not only to faith-based social welfare agencies but to organisations of any kind that are interested in state-funding to provide goods and services to, or on behalf of, government: Considering only those organisations that are friendly to liberal-democratic political values, and whose service provision would meet appropriate standards of fairness, quality, and accountability, is it acceptable in principle for the government to prefer one candidate provider to another on the grounds that the latter disagrees with, and has spoken against, various government values, goals and policies related to the same programme, or perhaps more broadly to the same department or agency?

This is a hard question. People are likely to feel differently about different examples.

Here is an argument in favour of the answer Yes. A government deciding whether to fund an agency to provide an ongoing range of counseling services is surely entitled to take into account a variety of ethical criteria, including the extent to which the agency is humane and fair in its dealings with its employees, and whether it has made amends for past misdeeds. Therefore the government's judgment of the agency's ethical suitability to partner the state in providing social welfare services may properly take into account its moral stance on sensitive public issues.

This argument needs some fine tuning. Sometimes when another individual disagrees with us on an important moral issue, we think their holding their opinion reflects adversely on their character. If so, then presumably the government can be justified in forming and acting on an adverse moral judgment about the ethos of a corporate entity, evidenced by its views on public issues. But it would not always be justified in doing so. Sometimes the government definitely should not regard an agency's opposite views as an indication that it is ethically unfit, or even less fit than some alternative body, to partner the state in providing social welfare services.

The criteria for deciding when some of your moral views indicate that you have a defective character are vague and contestable. Nevertheless the foregoing demarcation creates an opportunity to advance my third principle: it is not acceptable in principle for the government to prefer one candidate provider to another merely on the grounds that the latter disagrees with, and has spoken against, various government values, goals and policies.

Here is why I hold this view. I pointed out earlier that the government's fundamental task is to promote Australia's well-being, which requires lively and well-informed public discussion of both specific policy issues and also broader questions of value, justice and

meaning. Such discussion is obviously enhanced by the participation of organisations with a variety of well-considered value-schemes, with expertise (and indeed wisdom) in relevant policy areas, and whose views in these areas are not significantly influenced by their own self-interest. Expert critical scrutiny of the conduct and assertions of governments and other powerful players gives rise to a form of accountability that would otherwise be lacking. It functions as a check on both bad policy-making and the abuse of power, and thereby contributes to justified public trust in the system. Expert lobbying on behalf of unduly neglected or unfairly treated individuals is a corrective that the system needs to not only accommodate but encourage. Faith-based organisations are often well placed to carry out these roles. It is understandable that governments do not welcome exposure of their errors, and opposition to their proposals, but when they step back from their own corporate self-interest and consider the large-scale well-being of society, they should welcome and foster the participation of faith-based social welfare agencies, exercising the foregoing capacities, in the policy arena. It is obvious that discrimination against such organisations on the grounds that they disagree with government values, goals and policies related to their area of expertise would constitute a big disincentive to expressing disagreement, and in that respect would work against important aspects of the common good.

Notes

¹ A firm may take pride in the fact that its activities provide a genuine benefit to the community, and that it is fair and considerate to stakeholders including employees, but these are not typically independent aims that strongly motivate the firm's central operations. Narrowly focused aims tend to be accompanied by narrowly focused values centred on the quality of product or service and craftsmanship with which it is generated and delivered.

Critical Notice

Neil Levy's *What Makes Us Moral: Crossing Boundaries of Biology*

(Oxford: Oneworld, 2004. ISBN 1851683410. Price \$AU39.95.)

Richard Joyce

Australian National University

“Beer Gut Gene Discovered” announced the *Sydney Morning Herald* in 2003 (January 9)—yet another media declaration that scientists have uncovered the “gene for” such-and-such. Claims such as these are, in the popular consciousness, often conflated with proposals from sociobiologists and evolutionary psychologists regarding the innateness of certain human traits: infanticide, rape, or intelligence correlated with gender or race. When these traits are nasty or politically disconcerting (as are the three listed) then those pressing the claims are usually quick to point out that to identify any such tendency as the manifestation of an evolutionary adaptation is in no sense to exonerate the behaviour or to justify any political arrangement designed to accommodate it. Often, however, though we may not be quite able to articulate where this defence fails, we are left feeling uneasy.

Neil Levy's small but ambitious book, *What Makes Us Moral*, wades into this quagmire with the heartfelt intention of bringing some clarity to these consequential yet all too often confused issues. The result is an interesting and honest book: certainly not the “future classic” advertised on the bookflap, but a thought-provoking contribution nonetheless. The topic of the evolution of human morality lies at the intersection of many disciplines: not just evolutionary psychology and biology, but cross-cultural anthropology, economics, neuroscience, developmental and social psychology, genetics, primatology, and, of course, moral philosophy. The book is accessible to researchers from any of these fields — indeed, it can be read by a popular audience—for it is written in simple, jargon-free English. This resolve to write a book accessible to a general audience is to be praised, though it must be admitted that it results in some scene-setting which the more experienced reader will find all too familiar. We are, for example, given a potted chronicle of the impact of Darwinism, including the infamous debate between Bishop Wilberforce and Thomas Huxley, the Scopes “monkey trial,” and the history of the eugenics movement that ended in the Nazi death camps; we are treated to high school-level explanations of the process of natural selection and Mendelian genetics; kin selection and reciprocal altruism are covered in the customary manner (and group selection is rejected for the usual reasons), along with the inevitable discussion of the prisoner's dilemma and Axelrod's work on the tit-for-tat strategy. For a novice reader for whom these topics sound mysterious and interesting, *What Makes Us Moral* will be a readable and stimulating introduction. For those who've been through it all several times before, moments of impatience should be anticipated. Nevertheless, tolerance will be rewarded, for the author has some genuinely interesting arguments, and a sensible agenda to promote.

Levy's concerns have a scientific and a philosophical component. The scientific component addresses the question of whether the speculative hypotheses concerning the adaptive function of various human psychological and behavioural traits (put forward these

days principally by evolutionary psychologists) are empirically well-founded. Though Levy is no *tabula rasa* theorist—accepting that humans have “built-in biases and heuristics of which we are largely unaware” (199)—his primary concern is to counter these evolutionary arguments by emphasising instead the malleability and culture-oriented aspect of human psychology. The philosophical component addresses the question of what practical implications might follow if certain psychological or behavioural traits can be given an adaptive explanation. In particular, might the discovery that human morality is an adaptation upset any of our moral judgements or moral institutions? Levy identifies three possibilities. First, such a discovery might have an unexpected or unpalatable positive normative output; we might realise that we have obligations that we didn’t know we have (and that we wish we didn’t have). Second, the discovery that moral judgement is an adaptation—that we think this way only because doing so helped our ancestors make more babies—might have an undermining effect, revealing morality to be an illusion. Third, an evolved human psychology might place a constraint on what social and political arrangements are possible for us; certain ways of organising our relations with each other which seem *prima facie* fair and just might turn out to be inaccessible utopias, due to the recalcitrant biased nature of the human mind.

To the first possibility Levy devotes the opening chapter of his book (which comprises five chapters in total), focusing on the Social Darwinist movement as espoused by Herbert Spencer, and picked up on by Andrew Carnegie, J.D. Rockefeller, and Adolf Hitler. Those of us who know that such views are misguided hogwash—not just offensive but scientifically misinformed—will be comforted to see the author agree at some length. Chapter 2 discusses the second possibility, though most of it is in fact concerned with outlining the descriptive hypothesis that human morality could be the product of natural selection; the metaethical issue of whether such an evolutionary history shows morality to be an illusion receives a fairly brisk treatment at the chapter’s close (an argument to which I’ll return below). The third possibility is more of a ubiquitous theme in the book; entwined in his general critique of evolutionary psychology (which comprises Chapters 3 and 4) is one of the author’s take-home messages: that a confirmed evolutionary psychology would indeed represent a constraint on human choice, and thus any claim that such research has no practical, moral or political implications is naive. In fact, Levy argues, evolutionary psychologists, if they are correct, “are identifying significant and perhaps (for all practical purposes) immovable obstacles that stand in the way of some of our most cherished hopes for peace, equality, harmony, and happiness” (131). (This too is an argument I will return to below.) Thankfully, Levy distances himself from the intellectually repugnant view that because evolutionary psychology may have uncomfortable social implications we should reject its findings in advance of examining them with an open mind (138-9), though it must be confessed that at times I thought I detected the whiff of such a rhetorical strategy in the air. He might have been better advised to have kept these two issues—whether evolutionary psychology is empirically supported, and whether evolutionary psychology has unsettling practical implications—more clearly separated.

The scientific component of Levy’s case is, in my opinion, the more successful, though his critique of evolutionary psychology is perhaps not as far-reaching as he intends. Chapter 3 outlines a number of claims that have been made by evolutionary psychologists: that men are naturally more inclined towards promiscuity than women (Trivers’ work on parental investment), that some motivations in favour of infanticide and child abuse are adaptations, or at least the by-product of adaptations (Daly and Wilson’s research on step-parents and abuse), that rape is a male adaptive conditional strategy (Thornhill and Palmer’s work), that the human female brain is more wired for empathy than the male brain (Baron-Cohen’s research). Chapter 4 takes on these claims in turn, and does a fairly effective job of

criticising each. But of course evolutionary psychology turns on much more than just these particular issues. All the specific hypotheses mentioned could turn out to be groundless while the research program as a whole remains robust and viable. So numerous are the claims made on behalf of “evolutionary psychology” that many of them are sure to be problematic—some are likely to be downright silly—and perhaps this field attracts more than its fair share of silliness. Frequently Levy makes assertions of the form “Evolutionary psychologists claim that...,” when in fact it is at best *some* evolutionary psychologists who make the claim in question, and at worst just *one* researcher who has managed to get his possibly wacky evolutionary speculations into print. (Compare “Philosophers claim that...”)

Levy is aware of the limitations of a piecemeal approach, and thus in the final section of Chapter 4 attempts to launch a more general critique of evolutionary psychology. He is keen to resurrect the so-called “Standard Social Science Model” (SSSM) of the human mind—a label that evolutionary psychologists cannot utter but with contempt. According to the straw man SSSM (the one so easily scoffed at), the human mind is a blank slate that is engraved solely by environmental experience. According to Levy’s sensible version of the SSSM, by contrast, many human preferences and desires are indeed shaped by evolution, but culture determines much of the detail of how these desires play out in concrete form. “In general, and for many of the most significant aspects of human life, nature only sets boundaries: social norms and history settle what the way of life will be within them” (164). Unfortunately, Levy has not really allowed himself space to develop this hypothesis in any detail, raising the question of whether he was wise to devote his earlier energies to responding piecemeal to the more sensational offerings from evolutionary psychologists (rape as an adaptation, etc.). Absent from his discussion—perhaps also as a result of being pressed for space—is any acknowledgement that there is in fact an enormous literature in opposition to evolutionary psychology (a section on “further reading” at the end of the book cites just three works by critics of evolutionary psychology), so Levy doesn’t really succeed in situating his view in a live academic context. His most interesting contribution to this debate is his attempt to explain the apparent universals observed in many areas of human life without recourse to nativist theorising. Many such universals, he argues (drawing on Brian Skyrms’ work), are the result of coordination problems being solved by strategies that enjoy a only small natural advantage over competitor solutions, but whose tie-breaking advantage may then become massively amplified, thus stabilising the strategy across many environments. A relatively modest initial asymmetry, such as the physical difference between men and women, can lay the foundation for an entire edifice of asymmetrical genderised social norms, thus explaining the ubiquity of human patriarchy.

Given that Levy’s laudable intention is to champion a sensible middle ground between two unproductively dichotomised views, it is a little surprising that he contributes to the polarisation by describing his own position as a version of “the SSSM.” (After all, one could not reasonably claim his view to be orthodoxy among the social sciences.) But regardless of what we label his position, it is a viable alternative to the more extreme “massively modular” vision of the human mind advocated by many evolutionary psychologists, and it is a hypothesis that deserves to be developed. However, what is lacking in Levy’s presentation is any discussion of what evidence might be sought, what research undertaken, in order to decide between the two hypotheses. Evolutionary hypotheses are notoriously accused of being just-so stories: rich with creativity and broad plausibility, light on evidence. Levy has outlined an alternative—a worthwhile undertaking, to be sure—but has offered no hard evidence that should tempt us to think it actually true. One would like, at least, some idea of where we should look to locate such evidence (neuroscience? developmental psychology?).

Let me now turn to the more philosophical component of *What Makes Us Moral*. Above I outlined three possibilities identified by the author regarding ways in which evolutionary findings might have practical impact. In what remains I will offer some critical remarks concerning his discussion of the second and third possibilities.

In order to discuss the evolution of something we need to have a pretty good grasp of what that thing is. All too often discussions purportedly about “the evolution of morality” give no indication of whether “morality” denotes *prosocial behaviour*, or *social sentiments*, or *moral judgements*—resulting in a heap of confusion. Levy is an exception to this rule—making it clear that he is referring to the evolution of *moral judgement*, and going to some effort to spell out what is distinctive about this phenomenon. Morality, he thinks, has both Kantian and Humean elements. It is a system of prescriptions that are held to be unconditionally binding upon all rational agents (thus satisfying Kant), and these prescriptions are supposed to be intrinsically motivating (thus Hume). In addition, morality has a contentful constraint: It concerns the welfare of others, roughly speaking. I’m sure Levy is well aware that for each of these items there are numerous moral philosophers who will roar in protest, and his desire to sidestep this labyrinth of argumentation is understandable—but it has to be admitted that the result is somewhat dogmatic. If one is inclined simply to doubt that morality has (say) a Kantian element—as many respectable philosophers are—then one won’t find any arguments here to persuade the reader otherwise.

In any case, Levy goes on to identify the usual puzzle: How could something so authoritatively *other-promoting* evolve by natural selection? Surely natural selection is a process that will always favour the selfish over those motivated by the welfare of others? We now know how the answer to this puzzle begins (though perhaps we don’t yet know the whole answer): The processes of kin selection and reciprocal altruism (so-called) can result in organisms that behave in ways helpful to their fellows. Note, though, that neither process results in *evolutionary altruism*: the tendency for an organism to reduce its own reproductive fitness while raising another’s. (Evolutionary altruism requires group selection, which Levy dismisses on the grounds that the circumstances conducive to its occurrence are unlikely to transpire in nature.) In acting helpfully due to either kin selection or reciprocity, an organism is advancing its own fitness relative to a counterpart who does not act helpfully, and thus such behaviour is evolutionarily selfish.

For some reason this troubles Levy; he worries that altruism is only “altruism” (78). But it is difficult to see what the worry is. Let us begin by identifying three things that might pass under the name of “altruism.” The first, evolutionary altruism, I’ve just mentioned: It is a term of art concerning the relative reproductive fitness of organisms, and it is quite possible that there is no such thing. Second, we might just mean (somewhat vaguely), *helpful behaviour*—such that bees and ants are altruistic in this sense.¹ Third—and in line with vernacular English—we might mean actions that are performed with a certain other-promoting motivation (such that the terms “selfish and “altruistic” in this sense cannot be applied to bees and ants, but only to creatures with a degree of cognitive sophistication). If it turned out that all human social behaviour is really selfish in the latter psychological sense, this would indeed be a blow to our aspirations for morality. But the discovery that it’s all selfish *in the evolutionary sense* has no impact at all. The important thing to realise is that there is no reliable connection between the evolutionary and psychological types of selfishness. From the fact that a pattern of helpful behaviour is to be explained by the forces of reciprocity, for example—and from the fact that the behaviour is therefore evolutionarily selfish—*nothing* follows about the motivations of the creatures designed to engage in the

exchange. Reciprocal partners may enter into such exchanges for selfish motives, for altruistic motives, from a sense of moral duty, or their exchanges may be mere conditioned or hardwired reflexes properly described neither as “selfish” nor “altruistic” in the psychological sense. There are no grounds for insisting that such creatures must be “*really* selfish” in any psychological sense, that they are “really” motivated (at some subconscious level, perhaps) by their genetic interests, and that they are just self-deceived in this matter. Evolutionarily selfish creatures can be as sincerely altruistic, loving, and morally motivated as you please.

Given this, it was irksome to read Levy interpret my own “melancholy” view as holding that morality is a myth because really it’s all disguised selfishness (78-9). Certainly I have argued that morality is a myth, but my grounds for this conclusion had nothing to do with problems over altruism and selfishness. Rather, the argument is that we can give an evolutionary account of the origin of our moral judgements—an account that, moreover, may be empirically supported—which at no point presupposes their truth. Thus moral judgements are shown not to be false, but to be epistemologically unjustified. To this Levy responds that Michael Ruse and I are “measuring morality against an inappropriate standard” (79) in that we expect morality to be “out there,” as part of the furniture of the universe. Levy likens this dialectic to that concerning the existence of colour: Colours exist, he says, but not independently of our perceptual equipment. I remain unmoved by the analogy, since in fact I’m inclined to doubt that colours exist. Levy would respond that “we can all agree upon them [the colours], and we have much the same experience of them, ... [and] the fact that we can use colours for such important tasks as controlling traffic demonstrates that we have no qualms about their existence” (80). But this, it seems to me, is lamentably weak as an existential test. Scientists of the eighteenth century had no problem identifying escaping phlogiston: Anyone could point to an open flame and declare “There’s the phlogiston escaping!” If someone pointed to a bowl of water and said “This contains phlogiston” she would be making a mistake. Phlogiston theory could be successfully employed for all sorts of important tasks, such as cooking food or keeping warm. All this, yet there was no phlogiston! If it is possible for an individual to make a mistake about what the world is like, then we must allow that it is possible for a community of individuals to make such a mistake. Using collective agreement as an existential test, as Levy seems to, would belie this platitude.

Levy goes on to admit that his arguments here would, if successful, vindicate only some of morality: “If morality is real in so far as, and because, the emotions that underlie it are real and generally shared, then it is only its Humean side that is vindicated” (82). The Kantian element also stands in need of defence. (It is worth noting that the “inappropriate standard” to which Levy accuses me of holding morality is really nothing more ontologically extravagant than seeing the satisfaction of the Kantian component as vital. The “objectivity” with which I think moral judgements are imbued has less to do with their “out-there-ness” and more to do with the unconditional authority of the prescriptions.) Unfortunately, at this crucial point in proceedings Levy’s metaethical argument begins to fragment and unravel. He claims that even if we were unable to accommodate the Kantian aspect of morality, “evolution will have undermined not morality *per se*, but at least our commonsense concept of it” (82). But this sudden softening on the Kantian desideratum seems unaccountable. When earlier he defined morality for us, the Kantian element was front and centre; no hint was given that it is a negotiable and dispensable aspect of moral judgement, that it might be extirpated while leaving us with something still deserving of the name “morality.” (After all, without the Kantian element it is unclear that there is anything distinctive about the moral realm.) Furthermore, when Levy proceeds to attempt to accommodate the Kantian intuition (83ff.), what he in fact focuses on is *the sphere* of moral concern; no attempt is made to defend the distinctive *authority* of a Kantian moral prescription. Though the two properties

may be related, their connection is at best an obscure one. Moreover, he reverts to a genealogical mood—setting himself the task of explaining why natural selection may have furnished us with these Kantian intuitions (concerning which he has some genuinely interesting speculations)—seemingly having forgotten that the Kantian element of morality stands in need not merely of explanation, but *vindication*. We are, in the end, left waiting for a vindication that never comes. Despite the chapter closing with a section boldly entitled “Evolved morality is real morality,” we are in fact supplied with no reason to think that it is “real”—if by this is meant that moral pronouncements might actually be *true*. Indeed, in supplying an evolutionary genealogy of the Kantian element of morality Levy contributes to exactly the argument he set out to counter. If we have a complete, empirically confirmed explanation for why humans are inclined to categorise certain actions as unconditionally forbidden (say), but this explanation at no point presupposes that any such categorisations are true, then we should be left doubting whether they are in fact true at all.

Let me finally turn to the third way that Levy identifies of evolutionary explanations possibly having practical ramifications: that such findings indicate constraints on human choice, thus implying that certain social and political arrangements are better than others. Evolutionary psychologists are usually quick to deny that they are endorsing “genetic determinism”—the thesis that certain phenotypic traits of an organism are the product solely of its genotype. (Of course, nobody sensible has ever been a genetic determinist, for not even the number of legs that a person has is determined solely by the genotype.) Levy’s recurrent line here is that although his opponents may not be genetic determinists, they are still claiming that the genotype contributes the lion’s share to the phenotype—“that we understand human behaviour better by focusing on genes, mental modules, and evolved desires, than by looking at cultures and social norms” (129)—and this he interprets as implying that phenotype can be altered “only through great effort, and at great cost” (130). But it is doubtful that evolutionary psychologists are committed even to this watered down version of determinism.

First, consider the claim that evolved traits are *difficult* to alter. The fact that Levy persists in pushing this point despite the fact that evolutionary psychologists have denied it on numerous occasions leads one to suspect that they must be talking past each other. The diagnosis of the misunderstanding, I think, is that there are competing and non-equivalent notions of *innateness* in play. Even though Levy doesn’t discuss matters using the word “innate,” I believe light is shed on the dialectic by considering the useful disambiguation of the concept offered by Paul Griffiths.² Sometimes innate traits are considered to be those that are essential to being a member of a kind, sometimes they are those that can be given an adaptive explanation, sometimes they are those that exhibit developmental fixity in the face of environmental variation. In assuming evolved traits to be hard-to-alter, Levy is employing something like this last notion of innateness, having apparently drawn inspiration from Philip Kitcher’s *Vaulting Ambition* (1985), where it is claimed that sociobiologists predict that variation in many human traits will be relatively flat across the range of environments that are practically accessible. But the sociobiologists and evolutionary psychologists, I hazard to suggest, typically have in mind the second notion: that a trait is innate if it is an adaptation. And according to this view, innate traits may possibly be altered *easily* through environmental variation, especially if that variation is a type that did not occur in the historical environment in which the natural selection of that trait took place (e.g., Pleistocene Africa). As Griffiths reminds us: “There is no intrinsic tendency for evolved traits to be buffered against variation in environmental inputs to development. ... The constructive role of environmental factors in the development of evolved traits should come as no surprise.

Selection cannot favour a trait that compensates for the loss of a developmental input that is, as a matter of fact, reliably available.”³

Levy seems dimly aware of the existence of this competing notion of innateness, but it is casually consigned to a couple of footnotes. First, in footnote 95, he cites Richard Dawkins' claim that genetic influences may be easily altered by environmental factors. (In fact, in paraphrasing Dawkins, Levy has him conceding that genetic dispositions may “even, quite conceivably” be altered by environmental factors; but actually there is no such dampening qualification in Dawkins.) Second, in footnote 125, he cites Janet Radcliffe Richards' assertion (in *Human Nature After Darwin* [2000]) that adaptations may be no more difficult to alter or eradicate than traits that are the result of socialization. Levy grudgingly concedes that this is “true, in the abstract.” But instead of acknowledging and investigating the competing notion of innateness that underlies Dawkins' and Richards' remarks, he insists that evolutionary psychologists' own arguments show a commitment to adapted traits being robust, hard-to-alter traits. Their central claim, he observes, “is that these traits have proved impossible to eradicate across all environments so far, including environments in which the elimination of the trait has been an explicit aim ... The evidence of universality is explicitly advanced as evidence that the preferences in question are hard to alter” (225). But Levy's reasoning is fallacious on this point. Even if recalcitrance in the face of environmental variation can be used as evidence of a trait's being an evolutionary adaptation, it would not follow that all evolved traits exhibit such recalcitrance. Yet it is this latter general claim that lies at the heart of Levy's argument.

We should also give consideration to the second aspect of Levy's claim: that evolved traits will be *costly* to alter. His argument here seems based on generalizing from a particular example of how an (allegedly) evolved disposition can be altered. Apparently men the world over show a marked preference for a certain waist-to-hip ratio (WTR) in the female figure. But there is an exception: Men of the Yomybato tribe of Peru prefer the highest possible WTR. One hypothesis on behalf of evolutionary psychology is that the Yomybato have traditionally lived in an impoverished environment, where large fat reserves were a good indicator of fertility. Levy concludes that we could thus alter men's WTR preferences in general “only in one direction, and only by taking steps which would be disastrous and immoral: that is, by causing widespread famine” (p.130). Why he concludes that the alteration could go only in one direction is unclear. Why he thinks that the only way to accomplish it is via *famine* is also unclear. And, most importantly, even if it is true that the only way to alter men's WTR preferences would involve “great cost,” no grounds are supplied for thinking that what holds for this single case can be generalised to the conclusion that such alterations must *always* be costly. Even if it would take great effort to alter a particular trait, there is no obvious reason for doubting that the means might be congenial and worthwhile. (It might be claimed that *effort* must necessarily involve some cost. But then just about anything counts as costly: making a cup of tea, getting out of bed in the morning, etc. If this is the kind of “costliness” that Levy has in mind, it is hardly a kind we should be troubled by.)

Evolutionary psychologists are not committed to the view that evolved traits may be altered “only through great effort, and at great cost.” Some evolved psychological traits may be easy to shift through environmental influence, some may be moved only with great and unpleasant effort, some may be altered through great but desirable effort, and some may refuse to budge come hell or high water. All must be decided empirically on a case-by-case basis. Might a successful evolutionary psychology have political implications? Sure it could. Should we expect that it will? The jury awaits the empirical data.

Notes

¹ This is the sense in which what Trivers dubbed “reciprocal altruism” is altruism.

² Paul Griffiths, “What is Innateness?” *Monist* 85 (2002), pp. 70-85.

³ *Ibid.*, pp. 74-75.

Notes on Contributors

Andrew Alexandra is Senior Research Fellow in the Centre for Applied Philosophy and Public Ethics at the University of Melbourne
Email: a.alexandra@unimelb.edu.au

Dean Cocking is Research Fellow in the Centre for Applied Philosophy and Public Ethics at Charles Sturt University

Bernard Gert is Stone Professor of Intellectual and Moral Philosophy at Dartmouth College
Email: Bernard.Gert@Dartmouth.edu

Michael Hauskeller is Research Fellow at the ' ESRC Centre for Genomics in Society ' at the University of Exeter
Email: M.Hauskeller@exeter.ac.uk

Jeroen van der Hoven is Professorial Fellow in the Centre for Applied Philosophy and Public Ethics at the Australian National University
Email: vandenhoven@fwb.eur.nl

Justin Jeffrey is PhD candidate at University of North Carolina at Chapel Hill
Email: jjeffrey@email.unc.edu

Richard Joyce is Research Fellow in the Research School of Social Sciences at the Australian National University
Email: rjjoyce@coombs.anu.edu.au

Seumas Miller is Professorial Fellow in the Centre for Applied Philosophy and Public Ethics at the Australian National University
Email: smiller@csu.edu.au

Thomas Pogge is Professorial Fellow, Centre for Applied Philosophy and Public Ethics, Australian National University
Email: Thomas.Pogge@anu.edu.au
