

damage, flourishing, and two sides of morality

Adam Morton – adam.morton@ubc.ca

almost-final draft of paper to appear in [Eshare](#)

Humans are at their best when they are making things: families, social systems, music, mathematics,... . This is human flourishing, to use the word in the somewhat un-idiomatic way that has come to be standard in translating Aristotle and developing views like his. We admire well-made things of all these kinds, and the people who make them well. And although "happiness" is not a good translation of Aristotle's *edudaimonia*, it is a plausible conjecture about human psychology that people are happiest — most content, most satisfied with their lives, least troubled — when they are accomplishing, making, things of all these kinds, from families to mathematics. And they are miserable when they cannot. One kind of misery comes when one's efforts are not successful. Families fail, music is detested, "theorems" have counterexamples. Another kind of misery comes when one is blocked from being able to achieve any of the things that human life is shaped around. The focus of this paper is on ways that people's actions can make other people incapable of achieving properly human lives. This is what I call damage. I think its importance has only recently come to be appreciated; the delay in acknowledging it as a central moral concept has been particularly long in philosophy. And in human cultures worldwide an appreciation of how vulnerable we are to psychological damage is very recent.

damage

Consider first rape and the sexual abuse of children. These are dramatically awful examples of what I have in mind, and I will consider less extreme phenomena soon. My aim is to draw attention to the real reason that they are so wrong: they damage people. These reasons waited for our time: my lifetime, I think, and I was typical in coming slowly to the realization. Not long ago while decent people would have thought of them as wrong, often very wrong, they would have traced the wrongness to violation of autonomy and infliction of short term pain. The perpetrator is doing something to someone against their will, and it hurts. In other cultures the wrongness is also traced to factors that now seem to us perverse. In Roman culture the rape of a daughter or a servant is taken as an offence against the *paterfamilias*, and in Greek culture the rape of a woman in the temple of a goddess will usually lead to the goddess' anger at the woman for defiling her space, rather than at the rapist. One is reminded of reports of contemporary cultures in which rape victims are charged with adultery. But we have come to see a basic thing that is missing from these reactions. The victims are often *damaged* in a deep and long term way, which is sometimes seen as akin to post-traumatic stress. They can be prone to depression, irrational feelings of guilt, a sense of being bad and unworthy, and in some cases suicidal tendencies. Their capacity to flourish is drastically reduced (Resick 1993). I shall speak of a wide range of injuries to peoples capacities to lead satisfying lives as damage. I shall avoid the word "harm", because for my purposes it is awkwardly between pain and the damage that concerns me.

The authorities of the Catholic Church are generally decent and sympathetic people, and they never dreamt of anything but condemnation of abusive priests. But they took the

grounds for the condemnation to be forbidden sex rather than terrible wounding. Though this is a conjecture, some support is given by the papal document *Sacramentum Poenitentiae*, which takes the crime to be a violation of the commandment against adultery.

There are many ways in which people can be damaged, and many of them were invisible until recently. Post-traumatic stress, first noticed as "shell shock" after the first world war, is an example. We now see it as occurring also in milder forms. Similarly we thought of torture as the infliction of great pain, which it usually is. But in so doing we ignored the great injury to a person's conception of herself and ability to function, of which there is now abundant evidence. Torture, like post-traumatic stress, can also take milder forms, and can be subtle and psychological rather than overtly physical. (Bernstein 2015, Kashdana, Todd, and others 2006.) There are many other undramatic kinds of damage. We have learned that corporal punishment of children does not make them become well-adjusted and considerate adults. A vitally important topic is that of subtle implicit prejudice. There is now a lot of evidence that having one's attention drawn to one's membership in a group thought to be less capable reduces one's performance on tasks requiring attention and skill. (Schmader, Johns, and Forbes 2008.) One functions less well.

Recent work by both economists and psychologists suggests that stress coming from long-term deprivation, frustration, and envy of others disrupts the reward system in the brain and can lead to despair. The economists Case and Deaton (2017) use the phrase "deaths of despair" to refer to the epidemic of drug and alcohol abuse and suicide among working-class white Americans in the past decade. And the psychologist, Martinez (2010)

gives empirical evidence that feeling oppressed and unappreciated interferes with the reward system in the brain, involving the neurotransmitter dopamine, in a way that makes a person prone to addiction. There is also a fair amount of data that bringing attention to the fact that one is from a low status group that is not expected to succeed at difficult tasks disrupts short-term memory in a way that does indeed interfere with difficult tasks (Steele and Aronson 1995, Schmader 2008). So prejudice is in this way self confirming.

utilitarianism and Kantian ethics

Damaged people, according to this evidence, have less pleasant lives and accomplish less of what they want. Thus according to most versions of utilitarianism it will be wrong to inflict the damage. But this misses a point. The real damage is to their capacities, and not to their experiences. A person who overcomes harm that is done to them in order to achieve a satisfactory life has still been harmed, and the infliction of this harm is still wrong. It is certainly possible that the wrongness can be squeezed into a utilitarian framework. Almost any moral consideration can, with enough effort. But the aim would seem to be mistaken, since the victim is the person herself and not her experience. In fact, much of the harm that is done will not be reflected in the person's experience, since it will consist in projects not attempted, potential satisfactions never gained, and accomplishments that were never possible for the person in question..

We can also make connections with Kantian, deontological, ethics (Williams 1973b, 1985). Take the heart of the ethic in the form that one should take every person as an end rather than as a means. (And there is an implicit principle that this is nonnegotiable,

and does not get weighed against other good and bad things.) Then it is in the same general territory as the golden rule in Christianity (Matthew 7:12: "love thy neighbour as thyself"), or in Islam several passages in the Hadith, such as Kitab al-Kafi, vol. 2, p. 146, where the prophet gives as a rule "as you would have people do to you, do to them; and what you dislike to be done to you, do not do to them." So given that you do not want to be damaged, you do not want to inflict damage.

But again there are differences. They are very similar to the differences from utilitarianism. What you want is largely available to you in terms of what you consciously want, and the focus now includes forms of damage that you are not consciously aware of. Otherwise we would not need the empirical evidence that such events are as harmful as we now realize. And in fact a person may think that she wants something that is in fact harmful to her, such as the continuing availability of her drug, or her deference to the men in her life. There are other similarities and differences too. The concern is for the person herself and her capacities in both Kantian ethics and damage-avoidance. On the other hand, damage obviously comes in larger and smaller forms, ranging from the catastrophic to the trivial. So any damage-avoiding ethics will have to include ways of balancing lesser and greater benefits and evils.

Benefits as well as evils: you can increase as well as diminish a person's capacity to flourish. A parent talking seriously and articulately to a child does the child a service that she may never appreciate. And taking an abused person seriously and compassionately, showing that you appreciate the harm that has been done to them, goes a small way to mitigating that harm. (Being nice to people is no substitute for not hurting them in the

first place, though. And recovery from serious damage is not accomplished by niceness alone.)

the two neighbours

Why were these things ignored? I suspect three related factors. One is the special nature of our capacities for imagining each other. They have developed to mediate standard forms of social life and cooperative activity. So they focus on grasping other people's intentions and actions: the desires actions aim to satisfy and the beliefs that shape them. The other is emphasis on features of people's minds that they are conscious of. But people often do not know even that they are damaged.

The third factor connects the first two. Our intuitive ways of getting on with one another and our official codes of ethics typically serve to foster cooperation between individuals and groups in shared projects. I will help you harvest your crops if you will watch over my children. And these cooperative routines and conventions can themselves be implicated in damage. People who are not very ambitious for themselves and do not have much faith in their own capacities tend to be content to play a small role in the plans of others and the enterprises of larger groups. So people kept in a state in which they are somewhat limited and conformist are likely to be faithful participants in shared plans. They do not have the confidence in their own abilities — in fact many of these abilities may well longer exist — or the originality needed to make plans of their own.

This situation can be illustrated with an example I have used elsewhere (Morton 2009). (The description of the situation also gives the example a rationale, so that it becomes more than an intuition that may well not be shared.) You have two neighbours. The person on your right is a model neighbour, returning borrowed tools on time, sheltering your children when they come back from school early, and so on. The person on your left is far from this, occasionally taking tools from your garage without consent, coming

home singing loud happy songs late at night, and putting out trash in an insecure way that animals get at. In your country a dominant majority maintain their position by suppression of a minority. You have defended the rights of the minority, but the situation has moved from political debate to physical action, and members of the minority and those who support them are being rounded up for no-one knows what mistreatment. You plan to disappear and work in secret opposition. But you need a safe haven for your children, in a hurry. You could appeal to one of your neighbours. Which one should you approach?

One appealing way of reasoning is treacherous. It is to think "The neighbour on the right has always been friendly and cooperative, so the evidence is that she is a good person, so she will take personal risks, if need be, to protect my children." The personal characteristics that support judgements of moral character in routine everyday life concern a limited variety of situations. Social psychology tells us that human behaviour is less consistent from case to case than we tend to assume, even within a given social context, and inference from one context to a very different one is even more dubious. Moreover the fact that your neighbour on the right is cooperative in small everyday matters suggests that she may appreciate the good will of those around her, and be uncomfortable with the lack of it. Cooperative people are often conformists, and indeed a preference for conformity makes many everyday interactions proceed more smoothly. In the situation at hand too great a concern for one's image would be a liability.

One might indeed reason in the opposite direction. "My neighbour on the left is nonconformist and independent-minded. He makes up his own ideas about what to do,

not particularly trying to please those around him. So if either of the two neighbours is able to act contrary to the dominant public mood, it is more likely to be him.”

This contrasting reasoning might also be misleading. The neighbour on the left might be a less than model citizen not because he is concerned with more important things and thinks for himself, but simply because he is thoughtless and self-centred. But at any rate the contrast presents considerations to block the first way of reasoning. You need a more sensitive test.

imagination versus cooperation

I think the more sensitive test of whether these requirements can be met can be expressed in terms of imagination. You want to choose someone who can maintain the goodwill of conventional members of society while possessing enough insight to see how their conformity is based on prejudice. The second of these needs a particularly difficult kind of imagination, permitting the person to grasp the un-articulated motives and dispositions behind the smooth workings of society. If you can know that either neighbour has this capacity, then that is the one that you should choose. If you have no alternative, you should choose the one who you think is most likely to have this deep imagination (Murdoch 2001).

In the case of psychological damage as I was describing it earlier, any use of any normal imaginative skill within its usual range is likely to be inadequate. That is a basic reason why we continue to inflict many kinds of damage on one another without realizing quite what we are doing. Imagination has to work together with psychological evidence: once

we have reason to believe that particular acts can have particular consequences we have a chance of being able to imagine the effects on particular people.

A less demanding example of morally relevant imagination that tends in the same general direction is given by Jonathan Bennett's famous case of Huckleberry Finn (Bennett 1973). Huck has helped the slave Jim to escape, and realizes that this goes against the moral code in which he was raised. He is depriving Jim's owner of her property. Bennett quotes Mark Twain as attributing the following thoughts to Huck

(a) I couldn't get that out of my conscience, no how nor no way. ... I tried to make out to myself that I warn't to blame, because I didn't run Jim off from his rightful owner; but it warn't no use, conscience up and say, every time: 'But you knowed he was running for his freedom, and you could a paddled ashore and told somebody.'

(b) I knowed very well I had done wrong, ... Then I thought a minute, and says to myself, hold on—s'pose you'd a done right and give Jim up; would you feel better than what you do now? No, says I, I'd feel bad—I'd feel just the same way I do now.

Two forces are competing in Huck's mind. On the one side there is the code of respect for people in his society like those of his family (essentially, propertied white people) and cooperation with their projects, and in particular their use of their property. He identifies this with right as opposed to wrong action, and with conscience. On the other side there is his capacity to imagine the effects of his actions and his failures to act on the lives of particular other people, in this case Jim. We, unlike Huck, see that this has an equal claim to be labelled as moral. But they conflict. He cannot be a cooperative member of

his society and also minimize the bad consequences for Jim's life. At any rate he cannot fully yield to either of these without neglecting the other. (There are sophisticated compromises that will occur to sophisticated thinkers. But he is not one.)

Both the two neighbours example and the Huckleberry Finn example involved societies in crisis. (The society in the Huck Finn case does not know that it is in crisis.) Political crises will make moral crises for individuals. Tensions between the same competing forces can occur for individuals in non-crisis situations, for example when honouring a promise to one person about a matter of middle sized importance would mean inflicting serious pain — or serious damage — on another. (The possibility of damage is important, because it shows that these conflicts are not always easily conceptualized in terms of general rules versus particular consequences.) The rival considerations in such cases compete not only for influence on a particular person's actions, but also for the status of morality.

the dark side of morality

The psychology that fosters cooperation in real human beings, and the psychology that fosters the imagination of another person's condition, have their sinister aspects. People readily cooperate with others in their own social groups, and often the cooperation is in competition with those of other groups. In fact, this may be intrinsic to cooperation in finite beings, since the problem of discovering courses of action that are of most aggregate benefit to all people in the world is just too hard for our limited minds. So cooperation is usually motivated by identity, tribal loyalty, and precedent. It is then directed away from members of other groups, and lack of cooperation and outright enmity is often directed towards other groups. Members of subgroups that are not part of

the dominant cooperative scheme, such as children and women, can then easily find their claims ignored. In actual imperfect human beings the cooperative impulse can feed xenophobia, and indeed abuse.

A related set of problems accompanies empathy and imagination. This aspect too can be selective and biased, since imagining another person accurately is not a trivial task. In everyday life one tends to have a detailed and accurate grasp of what will help or injure only a small number of people. For all others one uses very rough heuristics about what people want and what does them harm, with the result that familiar people are treated mechanically and unfamiliar people are taken to be mysteries. I do not think we could manage our routine everyday life with others on the basis of one-to-one imagination, so that fixed cooperative routines will inevitably take much of the burden of our dealings with one another (Harris 2000, Morton 2013). So wherever we place the emphasis, there will be situations where acting in ways that often result in mutual benefit in fact produce disaster, injustice, damage.

The worrying side of the psychology that produces our best behaviour can be made dramatic by describing two extreme possible societies. In society A promises are kept, compromises are made, efforts are distributed sensibly, and in other ways people work together to achieve common aims. Very few think of their society as unjust. But because of family structure, the treatment of children, the relations between the sexes, ingrained habits of denigration and pessimism, plus many other hard to describe experiences that everyone undergoes, there is widespread depression, victimization of vulnerable people, and general unhappiness. Few people have wholehearted affection for others, and few

people attempt anything adventurous or challenging. It is a miserable place, though it thinks of itself as highly moral.

In society B all people are treated with great respect, children are raised carefully and affectionately, there is much soul-searching about the effects of practices on the well-being and capacities of individuals, and the realization of each person's potentiality to its fullest extent is sacred. But this takes time and energy, and as a result few can play a full conscientious role in shared projects, promises are kept in only a very cursory way, and little attention is paid to working out efficient ways of distributing different people's efforts. Here, too, people generally think of their society as highly moral, though the aspects they cite are very different to those that members of society A cite. They would describe their society in terms of consideration and kindness rather than in terms of justice and fairness. (For similar dystopias see Hinckfuss 1987.)

Society A, the puritans, and society B, the hippies, are both unsatisfactory. This is in spite of the fact that members of each think of their core commitments as defining the way people collectively ought to be. Is there a best compromise or mixture of them? Is there an ideal against which they can both be compared? "Best" and "ideal" here seem to beg the question, but I do not know even how to pose these questions more carefully.

atrocities

The contrast here can be put crudely as the cooperation of a good citizen versus the empathy of a sensitive person, especially to individuals she is sharing projects with. We need both, but they can conflict and they rest on different sides of human nature, at any

rate for instinctive reactions and heuristics. One place where the contrast is vivid is in issues about mass atrocity, public evil. At any rate it is if one takes a certain line on its origins.

The line comes from Hannah Arendt. (Arendt 1951, 1964, 1971). The central idea, as I read her (Morton 2004) is that many of the people essential to large-scale atrocity are disturbingly like the rest of us. They are not monsters, though a sprinkling of monsters may grease the wheels of atrocity. Rather, they are unimaginative civic-minded people, cooperating in a social project that, though they are blind to this, involves terrible events for many, especially people that they do not regularly interact with. The sprinkling of monsters would cause the project to fail if it were not for the mass of sensible, rationally cooperating, citizens who choose effective means to accomplish the project.

There are many ways of not getting the point. One is to think that all praiseworthy actions come from careful deliberation in the light of all morally relevant factors. Then of course a high regard for neighbourly cooperation will not lead you to complicity in genocide. But simple everyday helpful behaviour will not then count as moral either. Perhaps on this view we should actually condemn keeping promises and helping colleagues, if they do not stem from the most meticulously pure motives. (Kant actually suggests something like this, in chapter 10 of Kant 1785/1997.) Another way of missing the point is to insist that the ideal social contract would require people at the same time to minimize pain to one another, do one another little psychological damage, and to lend a hand in mutually agreed projects. We can have it all, if we formulate things carefully enough. Well, no one ever has formulated such a contract, and it remains to be seen whether it is even possible, given the variety of human situations. We can accept that

people often do act from an implicit grasp of an understanding with others around them of what is allowed is not, and that this regulates many of their actions. But there are always large gaps in the areas covered by such implicit contracts, including some crucial to human welfare. And inasmuch as they regulate the actions of real people in real life they concern behaviour with respect to a limited range of other people in a limited range of situations. That is one reason why societies seemingly committed to universal benevolence have regularly been involved in slavery, prejudice, and war.

the disunity of the moral

The argument has been in terms of the psychology of moral behaviour, taking both "psychology" and "moral" in fairly wide senses: the capacities that allow us to participate in the practices that make human life productive are not homogeneous. That leaves a more abstract question open, whether there is an idealized concept of right behaviour to which these disparate capacities allow us to approximate in our fallible clumsy ways. I do not think there is, as long as moral concepts are supposed to apply to the full complexity of human situations.

Return to the Huckleberry Finn example, and the story of the two neighbours. We read the Huck Finn case so that he would be following the morally better course in helping Jim to escape. And if the less conventional and less well behaved neighbour is the one who will shelter your children in times of crisis then he is the better person. But the stories can easily be made more nuanced, as in real life they would be. Until things get grim the unconventional neighbour may be someone you reasonably complain about and want to have little to do with. And for all you know until the crisis strikes you may

continue to have reason to think badly of him. (And aiding supporters of the minority may bring catastrophe on that minority by ruining economic cooperation between segments of the society.) And there may be a good chance that Jim will be caught and beaten to death as an example to others. Then Huck will think that he did a great disservice to Jim in helping him to escape. Few cases are really clear-cut. (Remember that Mark Twain was writing in the 1880's, long after the time of the fictional events, when the history had become simple and mythical. I should add, to preserve the reader's opinion of me, that as Twain tells the story there is no doubt about what on balance Huck should do.)

Is there always a morally best choice in such cases? Does damage-avoidance always trump cooperation, or vice versa? It seems very implausible that one factor will always dominate, even with a subtle weighting scheme. Accepting that these are dilemmas which will not yield to mechanical solutions (Morton 1996: I am afraid this last section is peppered with references to my own work), there is the further question whether they are moral dilemmas or general decision-making elements. Bernard Williams, as I read him, takes there to be incommensurable moral considerations at play in these and many similar cases. Whatever one does has a moral argument against it, so that a sensitive agent will pick what seems to be the least objectionable option and feel a specifically moral regret that the other option was neglected.

I am inclined to a starker reaction. The conflict is between principles that claim the label of morality, each with some justice because each often serves to allow us to live together for mutual benefit without doing one another too much harm. When there is an irresolvable conflict between them there is no determinate answer to which one of them

can best claim the label. They both normally coexist, and for reasons described above we normally are blind to the tensions between them. When the tensions surface and cannot be evaded, individual agents have hard problems about what to do. One consideration is which category of consideration a particular person can more effectively follow (Morton 1990). There are people who are good at cooperation and people who are good at principled rebellion, rather like the scientists who are at their best in Kuhnian normal science and the scientists who are at their best leading a scientific revolution. We need both.

Of course there are many people in the middle, neither ethically normal (conventional, conformist) nor ethically revolutionary (innovative, imaginative, radical), just as there are in the scientific case. But even given this range it is possible for two people who occupy the same position in it to use very different words to describe their dilemmas. One person may take herself to be reconciling a moral duty to live up to other people's reasonable expectations of her, against her desire to be helpful to people she cares for. Another person, in exactly the same situation, may take herself to be reconciling a moral duty to react to the plight of people in difficult circumstances, against her desire to maintain her image in the eyes of her associates. Suppose that they both resolve the tension with the same action. Both will feel regret, though one will describe it as regret at not being able to do the right thing, and another will describe it as regret at not being able to give help when it is needed.

A third person, also in exactly the same situation, may think of herself as torn between two moral considerations. And when she resolves the tension in exactly the way that the others did she also will feel regret. But she will think of it as regret that she could not do

the right thing because she had to do a different right thing. This is what Williams refers to as "agent regret" (Williams 1973a, Williams 1981). But it is directed at the choice between the same options as the other two people, and may well feel exactly the same. It is not easy to individuate emotions, but the way I would describe the situation or three people have the same first order emotion, and react the knowledge of what they feel with different second-order emotions, differing in the extent to which they take their emotional dispositions as unworthy. (For second-order emotions see Mendonça 2013.)

All three people use the phrase "moral duty" differently. Is one misusing it? When we strip away the contentious labels all three are reacting to the same situation in the same way. All three have to balance the same competing considerations, and arrive at the same resolution. A hard question is what cognitive resources there are for resolving conflicts between incomparable desiderata (Morton 1990 again) and what makes one mode of resolution better than another. I do not see any reason to suspect that resolving conflicts between incomparables is sensitive to which side of the conflict is labelled as moral. That seems to me the absolutely central question. If there is no such sensitivity then we should doubt that there is any tight unity to the considerations that we call moral.

Adam Morton

University of British Columbia

bibliography

- Arendt, Hannah (1951) *The Origins of Totalitarianism*. New York, Harcourt Brace & World.
- Arendt, Hannah (1964) *Eichmann in Jerusalem: A report on the banality of evil*. New York, NY: Viking.
- Arendt, Hannah (1971) Thinking and moral considerations: a lecture. *Social Research*, 38, 3: 417-446.
- Bennett, Jonathan (1974) The Conscience of Huckleberry Finn. *Philosophy* 4: 123–34.
- Bernstein, J. M. (2015) *Torture and Dignity: An Essay on Moral Injury*. University of Chicago Press
- Case, Anne and Angus Deaton (2017) Mortality and morbidity in the 21st century. *Brookings Institute Papers on Economic Activity* Conference Drafts, March 23–24, 2017
- Harris, Paul (2000) *The Work of the Imagination*. Oxford: Blackwell
- Hinckfuss, Ian (1987) *The moral society, its structure and effects*. Preprint series in environmental philosophy, Australian national University, no 16.
- Kant, Immanuel (1785/1997) *Groundwork of the Metaphysics of Morals*, trans. Mary Gregor Cambridge: Cambridge University Press.
- Kashdana, Todd, and others (2006) Fragile self-esteem and affective instability in post-traumatic stress disorder. *Behaviour Research and Therapy* 44: 1609–1619
- Martinez, Diana and others (2010) Dopamine Type 2/3 Receptor Availability in the Striatum and Social Status in Human Volunteers. *Biological Psychiatry* 67: 275–278
- Mendonça, Dina (2013) Emotions about Emotions. *Emotion Review* 5, 4: 390–396
- Morton, Adam (1990) *Disasters and Dilemmas*. Blackwell.

- Morton, Adam (1996) The disunity of the moral. In Jan Bransen and Marc Slors, eds. *The Problematic Reality of Values*, Amsterdam: Van Gorcum: 142–55
- Morton, Adam (2009) Good citizens and Moral Heroes. In Pedro Tabensky, ed., *The Positive Function of Evil*. Basingstoke: Palgrave: 127–38
- Morton, Adam (2013) *Emotion and Imagination*. Cambridge: Polity Press
- Morton, Adam (to appear) Cousins of Regret. In Anna Gotlib, ed. *The Moral Psychology of Regret*. Rowman and Littlefield.
- Murdoch, Iris (2001). *The Sovereignty of Good*. London: Routledge.
- Orth, U. and R. Robins (2013) Understanding the Link between Low Self-Esteem and Depression. *Current Directions in Psychological Science* 22, 64: 55-460.
- Resick, Patricia (1993) The Psychological Impact of Rape. *Journal of Interpersonal Violence* 8, 2: 223-255.
- Schmader, T., M. Johns, and C. Forbes (2008) An Integrated Process Model of Stereotype Threat Effects on Performance. *Psychological Review* 115, 2, 336–356
- Steele, C. and Joshua Aronson (1995) Stereotype Threat and the Intellectual Test Performance of African Americans. *Journal of Personality and Social Psychology*, 5: 797-811
- Williams, Bernard (1973a) 'Ethical consistency'. In *Problems of the Self: Philosophical Papers 1956–1972*. Cambridge: Cambridge University Press, 166–86
- Williams, Bernard (1973b) *Utilitarianism: For and Against*, with J.J.C. Smart. Cambridge: Cambridge University Press
- Williams, Bernard (1981) Moral luck. In his *Moral Luck*. Cambridge: Cambridge University Press, 20–39.
- Williams, Bernard (1985) *Ethics and the Limits of Philosophy*. London: Fontana