

HUMAN ACTION AND MORAL REALISM

GABRIELE DE ANNA

Universität Bamberg (Germany) and Università di Udine (Italy)

gabriele.de-anna@uni-bamberg.de

ABSTRACT

Widely accepted accounts of human action strongly suggest that actions can only be identified from the first personal perspective, i.e. from the point of view of the reasons that motivate an agent. That view has important consequences for moral realism, since it seems to entail that values are subjective: constructivist views of value would then be the only viable accounts of moral experience that does justice to claims of objectivity. This essay suggests that moral realism can still be maintained, if it understood as the existence of a fitness between potentialities existing in reality and volitional powers of well-functioning human agents. On the basis of the acceptance of the first-personal account of reasons for action and of some basic normative intuitions, the essay argues that an agent has to make room for the possibility that the way in which he responds to facts may be inadequate. On the other hand, that is possible – it is argued – if and only if there are correct and incorrect ways of responding to facts, i.e. there are ways in which well-functioning human beings respond to facts

KEYWORDS

Action, good, reasons, constructivism, moral realism, first person, fitness.

1. STARTING FROM ACTION

Debates on realism and objectivity in ethics differ from debates on the same topics in epistemology, since they raise peculiar and difficult issues concerning subjectivity. Subjectivity is certainly implicated also in epistemic acts, both since an epistemic act is carried out by a subject and since such an act is performed on the basis of subjectively chosen reasons. In the domain of ethics, however, subjectivity gains an even greater role. In epistemology, the realist can claim that the true is what it is, regardless of what conditions any subject is in, whereas, in the field of ethics, even the realist must recognise that the good,

or whatever has normative value, must be *practical*, since it must be able to motivate a subject. The very idea of a good that is such regardless of its capacity – at least in principle – to motivate an agent seems incoherent. Aristotle himself claims that all communities originate for the sake of some good, since all men act in view of what *seems good to them*.¹

The importance of subjectivity in ethics gives us a reason to tackle the discussion of realism and objectivity from the point of view of the subject, namely by starting from an analysis of human action and its criteria. This way of proceeding does not mean that one renounces the possibility of offering a theoretical account of the good, i.e. an account which does not consider the point of view of any particular agent and which does not aim at guiding action. The point is that, even in order to grasp from a theoretical point of view what the good is, one must consider that it has the property of guiding action. In this paper, I intend to investigate how the subjective character of the good can be accounted for within a conception of it, which can be still considered realist and objectivist. Since the subjective character of the good depends on its being *practical*, the discussion will have to start from the consideration of human action.

2. HUMAN ACTION AND REASONS

In a previous essay, I suggested that a human action is essentially something that an agent does for a reason. My grounds to make that claim depended on Elizabeth Anscombe's contention that an action is such if there is a subject who owns it and takes responsibility for it. That happens when the agent can answer the question "Why did you do it?" and can do that by giving a reason he had to act (Anscombe, 1957, 9-12). Indeed, the agent would think that the "why did you do it?" question would not apply or he would answer it by giving a cause or a motivation rather than a reason, if he saw the thing he did as not belonging to him and if he felt no responsibility for it.

The view that I just suggested implies that a human action must be rational, and this could seem odd. However, I think that this implication is not implausible, if one considers that rationality here needs not be full. As Robert Audi (2001, 50) pointed out, rationality in action can come in different degrees. The highest degrees involve a commitment to finding out truth and doing the good. Lower degrees can involve a mere consonance with reason. The lowest degrees concern what Audi calls minimal rationality, the condition

¹ Aristotle, *Politics* 1252a, 1-3.

of someone who lacks most epistemic virtues but at least is not insane, and to that extent can own his doings and be held responsible for them. All that I have to say seems to me to hold for an action regardless of the degree of rationality that that action implies.

In my view, reasons for action need to have an objective and a subjective side. They need an objective side, since a reason involves normally a fact. (In making this claim I follow only partially Raz, according to whom reasons are facts (Raz 1975, 16-19), but I will say more about this below). The fact that the beggar is hungry is a reason to give him food. However, reasons need a subjective side, since they have to be practical, as we have seen. I am dissatisfied with attempts to deal with this problem that distinguish two parts or two components of reasons for action, typically a belief and a desire or some other pro-attitude (Davidson 1963). My dissatisfaction is due to the fact that these views imply that if one holds the belief that ϕ -ing is morally obligatory but one fails to desire to ϕ , then one does not have a reason to ϕ . In my view, this implication fails to account for our moral discourse, according to which one's belief that ϕ -ing is morally obligatory certainly gives one a reason to ϕ , even if this requires frustrating one's desires. Rather than taking a reason to be the conjunction of two elements, a subjective element (a desire or another pro-attitude) and an objective element (a content-bearing belief), I propose to account for the subjective and the objective sides of reasons at once. I take the two sides to be aspects of the same thing, while holding on to the intuition that a reason always implies both some content and an attitude towards that content (Audi 2010). How can one do that?

My proposal is to consider that the content of a reason for action is not simply a fact, as Raz proposed, although it normally involves also at least one fact. The content of a reason involves a fact but also an unrealised state of affairs, which depends somehow on that fact and which the agent believes himself to be able to actualise. The agent must see the fact as representing some goodness, or some positive value, but also as needing an improvement (i.e., the yet unrealised state of things that is connected to the fact), and he must see that improvement as something that he can bring about. John sees Mark starving, and the fact that Mark is starving can be a reason for John to give him a sandwich. But it is a reason to do so only if John sees Mark or his life as a good, his starvation as a deficiency to be settled and himself as having the power to do something about it. If John thinks Mark is vicious and deserves to starve, Mark's starvation is not a reason to feed him. If John respects Mark and his life but he thinks that Mark needs to lose weight for his good, Mark's starvation is no reason to feed him. If John sees the worth of

Mark and his starvation as a problem to be solved, but he is conscious that he cannot do anything about it, he has no reason to do anything. In short, that there are reasons for action implies that an agent sees reality as having some order in it – an order that is somehow incomplete and deficient and needs to be improved – and that the agent sees herself as having the powers required to improve the partially existing order.

The upshot of my proposal is that one can recognise reality as having an intrinsic order, which calls for completion and furnishes reasons for action, if and only if reality is not conceived as a mere bundle of facts, but as a bundle of facts, which have proper ways of perfection and which can solicit a subject to respond to its deficiency. I call “good” any fact and any yet unrealised state of affairs that contribute to constitute some reason. This notion of the good keeps the objective and the subjective aspects of reasons for action together.

Reasons that explain actions are motivational: that means that an agent sees them as normative, i.e. that she sees them as indicating what she should do, regardless of what desires she has.

3. NORMATIVE REASONS, OBJECTIVITY AND PRACTICAL REALISM

So far I have claimed that reasons explain action and that at least some of the reasons that explain actions are motivational for the agent, i.e. they are seen by her as normative. That only means, however, that those reasons seem good to the agent, not that *they are really so*. Motivational reasons might not be normative. One may not do what one has (normative) reasons to do, or do what one has (normative) reasons not to do. “Why did you hit that man?” “Well, he looked pretty rich, and I wanted to get his wallet: I need money”. A perfectly understandable reason is given as an answer to the why-question. The agent seems rational, and in fact he owns his action and is responsible for it: if there is a failure here, as it seems, we cannot blame it to rationality². At the same time, however, if she sees reasons for action where there are actually none, that failure must be (at least) also rational: if it cannot be a case of irrationality, it must be a case of limited rationality. She fails in the exercise of her powers, also her rational powers.

The practical reason of an agent could have two main kinds of limitations. There can be cognitive limitations and limitations in the adequacy of subjective responses. On the one hand, indeed, an agent could have cognitive

² Cf. Vogler 2002, Chs. 1 and 7; Audi 2001 and 2010, 275-285.

limitations, which do not allow him to access all the relevant facts subsisting in the world. I understand cognitive limitations in a wide sense, including also the mere lack of information, which can be completely non-culpable. Since our grasp of the world is limited, we humans are always subject to this kind of limitations, even if their degree can vary considerably. On the other hand, there are limitations in subjective responses in the sense that a subject can consider normative reasons which are not so or non-normative reasons which are instead so, even if he has all the relevant information. The point is that he responds in an inadequate manner to the order that he can recognise in reality, on the ground of the information, which is available to him. Let us imagine two men, Tony and Jeff, who are at a party and see a plate full of cash lying on a table near the sandwiches. "Maybe the money is here for the guests!", they comment, and they grab a hand full each. Let us slightly rewind and enter the possible world in which, before they extend their hands on the cash, Dominic says "Gosh, this is the cash that we have collected for the charity before: I wonder who left it here". Now, let us imagine that, on the basis of this information, Tony takes his hand back, whereas Jeff goes on and makes a handy large also in this possible world. This example seems to justify the distinction between cognitive limitations and limitations of the subjective response. In the world in which the agents lack information about the origin of the money they behave in the same way, whereas in the world where they have fuller information their actions diverge. Tony responds by changing his intention, whereas Jeff does not: he turns out to be insensitive to the reasons of private property. He does not lack information, but we still see his response as problematic, since we think that the way in which he thinks proper to complete the partially existing order (the order concerning the existence of money and conventions about its distribution, for example) is not really a completion, but even a kind of further dissolution. Hence, an agent can be motivated by an inadequate reason because of a lack of information or because, even if he has all the information needed to be motivated by normative reasons, he does not respond as he should to the facts that he considers in his deliberative processes.

Can we say anything about the inadequacy of subjective responses? The issue of the limitations of practical reason is complicated, since it cuts across different issues in moral psychology. As we noted, practical reason guides action, but it does that through complex interactions with other cognitive and volitional faculties. Practical reason can guide processes of reasoning and deliberation, which can concern either the consideration of singular reasons or the comparison among different competing reasons. Practical reason can

involve no reasoning, but only insight. Practical reason can operate on habits, which are acquired through education and social life (the subject matter of virtue theory), and on desires. Practical reason can interact with consciousness and with freedom. Because of its complex relations, practical reasons can fail in many ways: that means that there are many ways, in which an agent can act differently from what the best reasons to him available suggest.

We cannot discuss all those complex relations here, but a full treatment of them would be needed for a full understanding of the distinction between cognitive limitations and the response-limitations of practical reason. Two problems about the notion of limitations in subjective responses need to be addressed here, though. They are general problems, in the sense that they arise regardless of whether cognitive or volitional capacity is involved in the relation with practical reason. The first problem is whether there really are failures of response, which might not be imputed to a lack of information. The second problem is on the basis of what criterion, granted that there are limitations of response, which do not depend on lack of information, can we claim that a response is a failure, if no incompleteness or falsity of content is involved.

Let us consider the first problem. In order to press the objection that gives rise to it, let us further consider the example of Tony and Jeff, who found cash on a table at a party. We could ask what influenced the failure of Jeff's response. A possible scenario is that Jeff tends to undervalue charity, due to the influence of his father, who followed Spencer's views on the treatment of the poor. He despises charity to the point that he is blind to the fact that an act of torpedoing a charitable initiative can be a theft. In this case, a certain habit conditions the response, by shaping the way in which the new information is elaborated and interpreted, in the light of information that is already held. The limitation of the response, in this case, can be reduced to a cognitive limitation: an act of theft is not recognised for what it is. In this case practical reason itself fails, since an event, a theft, is not recognised for what it is. Let us imagine a second possibility: Jeff did not respond in the right way and took the money, since he deliberately performed the action describing it as the taking of the property of others. However, he does not believe that there is anything wrong in taking the property of others, since he grew up in a commune and has never reflected on the social importance of private property and its respect. In this case, the lack of a proper response also possibly depends on a cognitive failure: if Jeff were to think carefully about the importance of private property, maybe he would start responding properly to information akin to that of the example. Also in this case, practical reason itself fails: unlike before,

it does not fail to recognise to what kind a certain event belongs, but it fails to achieve a decent level of coherence among beliefs, due to the influence of habits developed during the process of growth of the agent. This kind of examples suggests the objection that we are considering: maybe all failures of response are reducible to lack of information, i.e. to cognitive deficiencies.

There are counterexamples, however, which show that such a reduction is not always possible. Let us imagine another scenario. Jeff is fully aware that taking money would be a theft; he knows that thefts are unacceptable, and he is aware that his reputation would be ruined if his friends new that he did something like that. However, he has a strong desire to take the money, since he really wants to buy a certain pair of trousers and he lacks the resources. He tries to resist, but when he realises that no one is looking at him and that he would not pay with his reputation for that action, he gives up and does it. We would agree that he failed, but, in this case, we cannot reduce his failure to a cognitive fiasco. The failure, in this case, has to do with Jeff's incapacity to control his desires through his reason.

Let us now turn to the second problem. The examples that we have just considered aim only at suggesting that there are aspects of our experience, which justify a distinction between limitations of our practical reason due to epistemic deficiencies and others depending on failures in the response. If all the failures of response could be reduced to insufficiencies in the information available to the agent, we could draw extreme conclusions about objectivity and realism in the practical domain. Reasons would be strongly objective, since they would be totally determined by their contents. Furthermore, in that case, practical reality would be made of an order that subjects can recognise and of predetermined forms of completion of that order. On the other hand, if the differences among responses depend on differences among the subjects, which cannot be reduced to differences in the information they have, then the hopes of practical objectivity and practical realism are less certain. In case of lack of information, we can make sense of the distinction between proper and improper responses in counterfactual terms: a proper response is one that the subject would have given, if he had full information. However, if we assume that an agent has full information, how can we make sense of his failing to respond? This is the second problem that we have to face.

The example of Jeff stealing money under the pressure of a strong desire seems to suggest an easy way out: Response-failures that are not dependent on information depend maybe on a clash between practical reason and other cognitive and volitional faculties. They are cases in which – as we are assuming – reason is fully functioning from a cognitive point of view, but it fails to rule

other faculties, which should be subject to it. However, this solution does not work, since it relies on features of the example that not all cases of failure under full information share. Let us imagine a further modification of our example. Jeff knows all he needs to know about the situation and recognises that taking the money would be stealing; furthermore, he has no strong desire to get the money. Still, he does it.

There are two possibilities here. The first possibility is that he recognises that stealing is wrong, and still does not act accordingly. He has and sees reasons not to take the money, but ignores them, although he is not under the pressure of contrasting desires or other interfering attitudes. In that case, he is irrational (Broome 2015). But there is a second case: he recognises that taking the money is stealing, but he thinks there is nothing wrong with it. The first case is not very troublesome: we can imagine the agent apologising or recognising that he did something wrong. It is a case of failure of practical reason, in an agent who is generally rational. The second case, however, constitutes a problem for realism and objectivity: it suggests that an agent with a fully functioning rationality could have all the relevant information and still fail to respond. This is a challenge for realism and objectivity, since one can ask: “why shall we say that Jeff in that case *failed* at all?”

The point of the challenge is that information – even when complete – can underdetermine what it is right to do. The thesis that there is a partially realised order in reality presupposes that what exists grounds potentialities, which an agent must actualise. However, of course, actualities bear potentialities, which can be actualised in many, incompatible ways. One can respond to a beggar asking for help by giving him money or by turning away, in the sense that both courses of action are viable in the situation. What makes some ways of responding right/correct/acceptable and others wrong/incorrect/unacceptable? If disagreement about what the correct response is continues also in epistemically ideal situations, e.g., when agents have and properly understand all the relevant information, on what grounds can one claim that there are objective criteria to judge which responses are correct and which are not? Or that there are values – broadly understood – in reality? How can one suggest that there are normatively relevant facts, which can be criteria for our practical judgements, as I did above?

It seems to me that the challenge of the second problem can be met. Let us recall that we are considering the standpoint of the agent, and that we need a solution that takes that standpoint in account. Let us consider how the problem arises. We have a strong intuition that certain actions like stealing or killing an innocent or not helping someone in need are wrong and that

someone who does them rationally must be responding to the wrong reasons (let us call it “normative intuition”). Then we can consider that no information about reality, no objective fact can settle the question for the agent and be a criterion to rule out unacceptable reasons. But sure we cannot give up our normative intuition and conclude that anything goes. If there must be constraints and they cannot come from objective facts or from reality, they must come from the subject himself. The way seems open to constructivist solutions and no hope seems left for realism. However, this result can be overthrown if we consider what constraints can come from the subject and if we consider the point of view of the agent which practical philosophy requires and that we endorse here.

From the point of view of the agent, the normative intuition that generates the problem takes this form: “fact f seems to me a reason to ϕ , but is it really so?” Our normative intuition is that one could be wrong in responding to situations in certain ways, and from the first-personal point of view that looks like a doubt concerning the correctness of one’s response. The point is that a subject can think that his response to a situation must not necessarily be correct and that the reasons that seem normative to him might not be so. That thought, however, is not a doubt about facts, but it is a doubt about the correct way of *responding to facts*. This means that the doubt amounts to asking oneself how an ideal agent would respond in the same situation. I will say more about the sense of ‘ideal agent’ in this sentence. For now suffice it to say that it is not an idealised subject, but a normal human being who functions well from both the cognitive and the volitional point of view. In practical situations, this can be a concrete person, as for example someone that the agent may trust and consider as flourishing.

All this means that, from the point view of the agent, the criteria that constrain his reason and save our normative intuition do indeed come from the subject, as the constructivist would have it. However, they do not come from the subject in the sense that they flow spontaneously from the agent, or from a transcendental structure of his practical reason. Rather they come from the comparison with the imagined responses of a flourishing human being who were to find himself in the same situation. This means that the agent sees his constraints in a matching between the normative structures of the world and a human subject, whose volition functions well: some ways of being of subjects and consequently some ways of responding by them are appropriate and others are not. On the other hand, not all human beings can be equally sensitive to normative reality. This gives an unexpected realist twist to the argument.

At this point, however, one has the problem of explaining the magic pre-set harmony between subjects, who respond correctly, and the order existing in reality. The position that has just been suggested is that the world has a normative structure and that that structure matches the volitional structure of “well-functioning” human agents. Where does this matching come from? The agent can follow this way of reasoning only if she can account for that harmony, by recognising a teleological structure of reality. That can be explained theologically, or through a form of non-reductive naturalism (Foot 2001 and Nagel 2011) or on an evolutionary basis (Casebeer 2003). This line of thinking is challenged by the objections of evolutionary moral sceptics: I believe that their objections are instructive, even if, in the end, not decisive, but I cannot argue for this point here.

4. FROM PRACTICAL REALISM TO MORAL REALISM

So far, I have been discussing the practical domain and reasons in general, not the moral domain. Sometimes I have used the word ‘moral’, but that was not for the sake of clarifying that term: it was only to offer examples of normativity, without any concern for what is special about morality. What I have been suggesting holds for the reasons we have to buy a nice painting, to play a nice piece of music, to go for a walk, or to choose a career. It holds for all that we do, in general. Moral actions and moral reasons are subsets of those that we have so far considered, but they are significant subsets, from the point of view of objectivity and realism.

The main difference between moral and non-moral normative practical reasons is that the moral imply an obligation that the non-moral ones lack. In a sense, non-moral normative reasons also indicate what one *ought* to do, but this is a weak sense of ‘ought:’ their indications can be ignored simply by ruling out an assumption on which they rest. The assumption might concern a certain result that one may want to attain, or being a certain kind of person that one wants to be. “If you want to play that passage on the violin well, you must study it in separate stings”. That statement might give Tony a normative reason to study the passage in separate strings. Tony could also decide not to study the piece that contains that passage at all, however. “If you want to be the kind of person who is always invited to parties, you *ought* to wear brand cloths and always smile”. Good, but if I *do not want* to be that kind of person I *can* avoid doing those things.

In the case of moral normative reasons, by contrast, obligation is stronger. “If you want to help that person who is starving you can give her a piece of

your sandwich.” If Tony said that he has no interest in helping that person – provided that he is the only one who can do it, and all other necessary amendments – we would think that he is not as justified as he was when he decided not to practice a certain passage on the violin in separate strings. “If you do not want to be a liar, you must avoid saying even small lies”. The answer “I do not care about not being a liar” does not seem as acceptable as the answer “I do not care about not going to parties.”

We can observe that some moral obligations are stronger than others. Some are so strong that they become even legal obligations: their violation cannot be tolerated and legitimate authorities can punish who violates them: “do not kill”, “do not steal”, etc. Here I will not deal with this distinction among different degrees of obligation of moral reasons and discuss only the distinction between moral and non-moral normative practical reasons.

Moral normative reasons are problematic from the point of view of realism and objectivity since they seem to request a stronger normative constraint in reality, than is required by non-moral normative reasons. Let us consider our previous example: “if you want to play that passage you must study it in separate strings”. Let us imagine that this is a genuine reason: it is true that the given order (i.e. the technical level of the player, the characteristics of the piece and so on) calls for a completion of a certain kind (studying in separate strings). That order, however, can be completed also in other ways, for example by studying an easier, but equally pleasant piece. The existing order calls for a response from the agent, but there is more than one way to complete that order and the subject can respond in many different ways.

When we deal with moral reasons, it seems that, among all possible ways of completing an order, some might be necessary and others might be unacceptable. That does not mean that a certain variety is not possible, but the stronger force of moral obligation seems to depend on the fact that in these examples there is a limitation in the range of viable possibilities for completing an order that in the case of non-moral normative reasons does not hold.

On what does this limitation depend? In the case of non-moral normative reasons, there is a range of open possibilities for the completion of an order – which may include also the option of leaving things as they are. The choice about which form the completion of order should take is left to the subjective response of the agent (regardless of the fact that the subjective response can depend from differences in information or not). In the case of moral reasons, by contrast, it seems that a certain response (or at least one of the responses in a certain set) must be given by whoever happens to be in a certain situation. The response (or a set of possible responses) is obligatory in the sense that it is

requested from anyone, regardless of the subjective characteristics of the agent. The point is that in non-moral normative reasons the obligation is conditional upon the assumption of the intention to be a certain kind of agent or of choosing a certain kind of end. In the case of moral normative reasons, the obligation holds regardless of the kind of person that an agent wants to be and the kind of ends that the agent wants to attain.

We have seen above, while discussing non-moral normative reasons, that an agent can always ask himself whether the reasons that seem normative to him are really such and that that question can take the form of a question about what a person whom the agent trusts and respects would do in a similar concrete circumstance. It is natural to expect that the chosen person may be a concrete or an ideal person, who in a certain way paradigmatically exemplifies the kind of person that the agent wants to be or who fully realises certain ends that he has chosen. In the case of moral reasons, the criterion can no longer be a (concrete or ideal) person of a kind to which the agent belongs: the question then is how any man would respond in the given circumstances, i.e. a man who can be fully counted among humans. I do not think here of a point of view of humanity. It is rather the point of view of a concrete man who has all and only the qualities that make him a well flourishing human, i.e. someone who chooses not on the ground of belonging to a certain specific category, but only under the generic assumption that he is human.

For the point of view of an agent who chooses, then, asking whether those which seem to one moral normative reasons are really such, means asking what is human in the concrete given circumstances and reasoning to this question means asking – in an intensional sense – what being human means.

This conclusion allows me to note that moral reasons limit the range of possibilities open to agents more than non-moral normative reasons, since they are more objective: the limitation is operated by a criterion that requires the recognition of what is human and the intuition of how humans can become in different conditions in which their lives can be realised. Let us remember that from the beginning we are reasoning from the point of view of the agent, i.e. of a subject who chooses what to do and how to do it: when an agent asks himself whether the reasons that seem normative to him are really such, he will never reach an absolute point of view. Even if he decided to follow some moral authority, this choice would still depend on reasons that he sees as such and thus the point of view that he could reach would not be absolute. It is a corollary that the moral objectivity and realism that we have reached is not absolute. However, the points I argued for still show, it seems to me, that an agent cannot but commit himself to what seems objective and real to him, i.e.

to the order that he recognises and to what being human means. I believe that his commitment is indeed absolute.

REFERENCES

- Anscombe, Elisabeth. 1957. *Intention*. London: Blackwell.
- Audi, Robert. 2001. *The Architecture of Reason. The structure and Substance of Rationality*. New York: Oxford University Press.
- Audi, Robert. 2010. "Reasons for action." in *The Routledge Companion to Ethics*, edited by John Skoropski, 275-285. London and New York: Routledge.
- Broome, John. 2013. *Rationality through Reasonableness*. Oxford: Wiley-Blackwell.
- Casebeer, William D. 2003. *Ethical Natural Facts. Evolution, Connectionism and Moral Cognition*. Cambridge (MA): MIT Press.
- Dancy, Jonathan. 2000. *Practical Reality*. Oxford and New York: Oxford University Press.
- Foot, Philippa. 2001. *Natural Goodness*. Oxford and New York: Oxford University Press.
- Nagel, Thomas. 2011. *Mind and Cosmos. Why the Materialist neo-Darwinian Conception of Nature is Almost Certainly False*. Oxford and New York: Oxford University Press.
- Raz, Joseph. 1975. *Practical Reason and Norms*. London: Hutchinson.
- Vogler, Candace. 2002. *Reasonably Vicious*. Cambridge (MA): Harvard University Press.