Ewing's Problem

Christian Piller
University of York
Department of Philosophy
York, YO10 5DD, UK

ABSTRACT
Two plausible claims seem to be inconsistent with each other. One is the idea that if one reasonably believes that one ought to *fi*, then, on pains of acting irrationally, one indeed ought to *fi*. The other is the view that we are fallible in respect to our beliefs about what we ought to do. Ewing's Problem is how to react to this apparent inconsistency. I reject two easy ways out. One is Ewing's own solution to his problem, which is to introduce two different notions of ought. The other is the view that Ewing's Problem rests on a simple confusion regarding the scope of the ought-operator. Then I discuss two hard ways out, which I label objectivism and subjectivism and for which GE Moore and Bishop Butler are introduced as historical witnesses. These are hard ways out because both of these views have strong counterintuitive consequences. After explaining why Ewing's Problem is so difficult, I show that there is conceptual room in-between Moore and Butler, but I remain sceptical whether Ewing's Problem is solvable within a realist framework of normative facts.
KEY WORDS: Rationality, normativity, fallibility, practical reasons, subjectivism, objectivism.

AC Ewing thought that a widely accepted moral principle is at odds with a fundamental assumption of moral discourse. According to this principle, each of us is his or her own authority in moral matters. We always ought to act in accordance with our beliefs about how we ought to act – let our conscience decide. Ewing, however, also accepted that in moral discourse we use a notion of moral truth that can be missed. People can be wrong about what they ought to do. His problem was how to solve the tension between our moral fallibility and the authority of our moral conscience.

I disagree with Ewing. I do not think that everyone is his own moral authority. Nor do I think that such a view is part of common-sense morality. Nevertheless Ewing has pointed to a deep problem. According to a common-sense theory of practical rationality, what is rational for an agent to do depends on his or her conception of the situation. In matters of rationality, we make an assumption, which runs parallel to Ewing's ethics of conscience: one could not act rationally if one acted contrary to one's belief about what one (rationally) ought to do. This principle, however, seems incompatible with a notion of normative truth, truth about what one (rationally) ought to do, that people might miss. How should we resolve this apparent inconsistency? This is what I call Ewing's Problem.

After setting up Ewing's Problem, I discuss two easy ways out. One is Ewing's own solution to his problem, which is to introduce two different notions of ought. The other is the view that Ewing's Problem rests on a simple confusion regarding the

scope of the ought-operator. Neither of these attempts, I will argue, provides us with a satisfactory solution to Ewing's Problem. Then I consider whether we could give up one of the principles in conflict and I discuss what consequences this would have. Finally, I introduce another solution, but being dissatisfied with it, I conclude that Ewing's Problem remains unsolved.

1. The Moral Case

In his book *Ethics* AC Ewing (1953, pp. 144f.) writes, "It is a recognized principle of ethics that it is always our duty to do what after proper consideration we think we ought to do, but suppose we are mistaken, then we by this principle ought to do something which is wrong and which therefore we ought not to do. Is not this a contradiction?"

Two ideas, both not implausible, seem to be in conflict with each other. The first is a principle of self-reliance. It says that in all moral matters one's conscience ought to decide. The second principle is an instance of the view that we are fallible, and it says that moral matters are no exception: one's conscience can lead one astray. Ewing's Problem is how to react to this apparent conflict between self-reliance and fallibility.

Ewing (1947, pp. 122f.) illustrates the conflict with the following example: "We may believe," he says, "that the soldiers who fight against us in a war are acting wrongly in fighting, yet every reasonable person will admit that, as long as they really think they ought to fight, they ought 'to obey their consciences' and fight." Observing what Americans and its allies call 'The War on Terrorism', a war whose nature excludes mutual feelings of military honour, puts Ewing's example in doubt. Some people might think that President Bush's actions are morally justified; other people might have sympathy for Al Quaida's agenda. There is, however, no one I know who thinks that both Bush and Al Quaida ought, from a moral point of view, to act on their convictions. Such a view would be cynical and, as such, strongly at odds with common-sense morality.

The case of anyone whom we regard as a misguided fanatic shows that our admiration for actions out of conscience will be limited by our sense of moral right and wrong. Common-sense morality does not support the principle suggested by Ewing, according to which it is always a person's duty to do what she thinks she ought to do. Not every crime is turned into a moral duty, nor does it become excusable, simply because it was done out of conviction. A belief that some action is right does not make it right, and, thus, conscience has to play a more modest role.

One accepted role for our conscience is to be our guide in situations where duties conflict without one clearly dominating the other. The true patriot will go to war and leave his mother behind, whereas the loving son will choose caring for his mother over serving his country. We would understand both choices and we would not condemn either one of them. In a situation like this, one will ask oneself, what one really identifies with. If one does come up with an answer, then, according to common-sense morality, one ought to follow what one sees as one's true path. In so doing, one establishes or preserves one's moral identity. The relevance of such considerations is, however, limited. The case of the misguided fanatic shows that

these considerations will loose their appeal once we lift the restriction which keeps the available options within the domain of what is morally acceptable.

I have introduced the figure of a misguided fanatic to draw a limit to Ewing's ethics of conscience. In the next section, I look at the idea of self-reliance – one ought to do what one thinks one ought to do – in the more general context of a theory of practical rationality. I will argue that, in this more general context, a principle of self-reliance looks more plausible.

2. The Case of Practical Rationality

A theory of practicality rationality aims at a general account of what one ought to do. It faces two tasks. First, it needs to tell us which considerations are relevant in determining what one ought to do. This first part gives us a theory of reasons. Secondly, a theory of rationality needs to explain how reasons determine what one ought to do. Reasons will stand in various relations to each other. They can undermine, support, outweigh, or exclude each other. This second part is a theory of how to aggregate reasons. We need this second part in order to determine what one, overall, ought to do.[1]

In the previous section, I have argued that Ewing's Problem – the incompatibility of self-reliance and fallibility – does not arise in the moral domain. We can be wrong about moral matters; and if we are, then it is not true that we ought to act in accordance with our moral convictions. Does the same idea – namely that no appeal to a principle of self-reliance could overturn the verdict of normative facts – apply equally in the domain of practical rationality? Take a case in which someone is wrong about what, for him, is rational to do. Suppose a man, fascinated by bungee jumping, has a go all by himself. He ties a rope around his ankle, fixes it to bridge's railing and jumps. He has not thought of the required elasticity, so what he does is certainly unwise. He thinks it is rational to jump, but he is plainly wrong. We know that he ought not to jump. Like in the case of morality, our commitment to fallibility seems to make any appeal to self-reliance futile.

In this case the man does something insensible because he has not thought matters through. The agent's belief that his suicidal set-up matches that of bungee jumping is unreasonable. In order to assess any principle of rationality, however, we should provide it with reasonable starting points. (Any form of good reasoning may lead one to crazy conclusions, if the starting points are crazy.) Taking this point on board, the Principle of Self-Reliance tells one to act on one's beliefs about what one ought to do, if such beliefs are reasonable.

[1] Theories of rationality differ in respect to the normative status they assign to moral considerations. At their strongest, moral reasons will outweigh all others. At their weakest, the normative force of moral considerations will depend on whether the agent to whom they apply is in the right psychological states. For example, according to some theories, the harm someone will suffer as a result of an agent's actions will only count against so acting if the agent cares about the affected person's well-being. According to these theories, moral considerations need not register amongst an agent's reasons. Their moral relevance, however, is sustained, as they still determine what is right and what is wrong. These issues, however, need not be decided here, as I will focus on examples which are morally neutral.

Acting on one's reasonable beliefs about what one ought to do might still fall short of success. Even if the equipment has been checked and one reasonably presumes it to be safe, it might still fail. Rationality does not guarantee success. Furthermore, acting successfully seems compatible with failing to do what one rationally ought to do. Had one not used the faulty equipment, despite believing that one ought to, one would have had a lucky escape. Such avoidance of harm, however, is not to be credited to one's rationality. Quite the contrary, it was irrational not to do what, as it turns out, would have been harmful. The Principle of Self-Reliance allows for this gap between rational and successful action. It arises because reasonable beliefs need not be true. This separation between rationality and success is very much part of our common-sense conception of rationality, both in the practical and in the theoretical domain.

A parallel move in the moral domain – namely to restrict the relevant principle to reasonable starting points – would certainly improve an ethics of conscience. The misguided fanatic is usually conceived as someone who has, judged from our own moral perspective, unreasonable or outrageous moral views. Thus, a serious version of Ewing's Problem might arise in the moral domain as well. Nevertheless, common-sense morality differs from a common-sense theory of rationality in at least the following respect: they assign different roles to the notion of success. Whether your well-intended action pleases or annoys those affected by it is something beyond your control, and even your most reasonable beliefs about its effects might be wrong. The effect, however, creates a moral difference in what you have done. It does not, however, create a difference in your action's rationality.

Following Nagel (1976), we can capture this point by saying that common-sense morality has a place for moral luck. Moral evaluations are, to some extent, success-related, the 'ought' of practical rationality, by contrast, seems to be independent of success. Luck affects the relation between what we try to do and what we end up doing. A rational agent weighs up the reasons for and against doing something and acts in accordance with her considered judgment about where the overall force of reasons lies. Acting rationally, I have claimed, does not guarantee acting successfully. We still need the cooperation of things outside our control to get things done the way we want. But whether we set ourselves in the right way is up to us. Bad luck strikes when our best efforts remain fruitless, and, depending on the circumstances, we might need more or less good luck to carry us from the rational to the successful. The very point of a theory of rationality seems to be to set us up in a way that minimizes the luck we need. If so, whether an action is rational or not is simply not part of the domain where good or bad luck could strike. Paraphrasing Kant we can say that the rational will is not rational because of what it effects or accomplishes; it is rational only through its willing, i.e. its rationality is not, in the relevant sense, success-related.[2]

This is why Ewing's Problem looks more serious as a problem for theories of practical rationality. The idea that an agent, in virtue of being rational, has to be guided by his or her own conception of the situation looks more plausible than its analogous moral principle. Suppose an agent has thought carefully about what to do and after proper considerations of all available facts comes to the conclusion that she

[2] I add 'in the relevant sense' to allow for cases in which the rationality of pursuing an end is itself something the agent aims at.

ought to *fi*. Ought she not, as a matter of rationally required consistency, to *fi* now? That is what Ewing's first principle, the Principle of Self-Reliance understood as a principle of practical reason, would demand. An agent, if rational, it seems, must rely on her own reasoning. What else could a rational agent rely on or be guided by?[3] According to the second idea, however, anyone is fallible in one's judgments about what to do. And if what one thinks is wrong, which always seems possible, the Principle of Self-Reliance would have it that one ought to *fi* whereas, by assumption, it is not the case that one ought to *fi*.[4]

3. Ewing's Solution

"Isn't this a contradiction?" Ewing (1953, p. 145) asked and he answered, "... it would be if we were not using two different senses of "ought"."

What are these two senses of 'ought'? The principles of self-reliance and fallibility define two different notions of ought. Let me start with the latter, which Ewing (1947,

[3] When I would say that an agent's choice satisfies the Principle of Self-Reliance, Pettit & Smith (1993) speak of 'orthonomous (or right) choice in the narrow sense', which is to act in accordance with and because of one's judgment about what one ought to do to. For them, the Principle of Self-Reliance is central to the idea of acting rationally. "If an agent judges that a certain option is to be done, if she sincerely sees that option as best, then any failure to take that judgment fully to heart is a failure of reason" (Pettit & Smith 1993, p.72). According to Pettit & Smith, self-governed or autonomous choice is a broader notion than rightly governed or orthonomous choice. "We see the non-heteronomous agent, the agent who is practically rational in the narrow sense, as someone in whom desire is appropriately governed [namely by one's judgment about what one ought to do], not just someone in whom government of desire is exercised by her' (*ibid*., p. 76). They illustrate their notion of orthonomy by contrasting them with two accounts of autonomous agency, which both fall short of orthonomy: Sartre's (1958) idea that any operative desire needs to be affirmed by radical choice and Frankfurt's (1971) idea which finds such affirmation on the level of second-order desires. They conclude, "The good government of desire is a regime under which desire is faithful to the rule of deliberation; being endogenously inspired and maintained is not enough, even if it is necessary" (*ibid.*, p.76). In this paper, I focus on the Principle of Self-Reliance and the problems arising from it. I leave it open how this principle relates to different accounts of autonomy. In section 8, I will present Bishop Butler as arguing for an understanding of autonomy in terms of the Principle of Self-Reliance. However, nothing would change argumentatively if I opted for the broader notion of autonomy suggested by Pettit & Smith.

[4] In a previous issue of this journal, van Willigenburg (2005) discusses a problem which, terminologically, looks similar. He asks: How is autonomy – making one's own choice – compatible with orthonomy – making the right choice? If by 'autonomy' he would mean acting in accordance with the Principle of Self-Reliance, he would deal with the same problem. His understanding of autonomy, however, differs from the Principle of Self-Reliance. It is given by the minimal account that if I choose autonomously I have to see such a choice as mine. He writes, 'How can we get a sense of making our own choices, if orthonomy implies that the criteria for doing so cannot simply be of our own making?' (van Willigenburg 2005, p. 76). The similarity to Ewing's Problem is, I think, only superficial. In contrast to Ewing's Problem, van Willigenburg's question has, in my view, an easy answer. How can we make our own choices if the standards for correct choice are not up to us and we want to make correct choices? Compare the following question. How can I choose something blue if the standards for blueness are not up to me? It is easy. As a rational person, I choose what I take it to be blue. Choosing what one takes to be blue is, if things go well, choosing something blue. Van Willigenburg's question receives an analogous answer. How is a correct choice my choice if its being correct is not up to me? I make what I take to be the correct choice and, if things work out, it will be a correct choice. I can make the right standards my own, simply by endorsing them. The existence of standards of right choice does not threaten my ability to make my own choices. If, however, one standard one has to follow is to act in accordance with one's own judgment, then Ewing's Problem does arise – how is this standard compatible with our fallibility in the normative domain?

p.118) calls the 'absolute' or the 'objective' ought: "'The action we ought to do' may mean that action which is really preferable, taking everything into account. This would be the action which an omniscient and perfectly wise being would advise us to perform; but it is impossible for us to take everything into account." Fallibility means that we might be wrong about what we ought to do in this 'objective' sense of ought. The second sense of 'ought' is subjective as its application depends on our beliefs about what we, in the objective sense, ought to do. Or so I will understand it. The belief that one objectively ought to do something grounds the subjective obligation to do it.[5]

If we accept these two senses of ought, Ewing's Problem disappears. There is nothing contradictory about its being the case that one objectively ought to do something that one ought not to do, subjectively speaking. The conflict between the principles of fallibility and self-reliance has been avoided because in each we employ a different notion of what we ought to do.[6]

What should we make of Ewing's solution? If we accept it we solve a problem, but at what price? Ewing's solution belongs to a group of views that divide the normative domain. Some divisions seem unproblematic. For example, we can distinguish those of an agent's reasons that relate to her self-interest from those she has in virtue of moral obligations that apply to her. This division is unproblematic as long as there are ways in which reasons from different sources can be aggregated to overall ought judgements. Should such a division be carried over to the level of aggregation, however, it violates a presumption that we usually make when we engage in practical reasoning. The presumption is that, cases of indifference aside, there is one thing that we ought to do. When faced with two options, should I stay or should I go now, I have to pick one of them. The point of engaging in practical reasoning is to pick the right one. But if our conception of what one ought to do were itself divided, so that there is always the possibility that in one sense of 'ought' I ought to stay and in another I ought to go now, the very point of figuring out what one ought to do would be undermined. The question that guides practical deliberation, what ought I to do, would itself be ambiguous. The fact that we always have to act in one way, thereby

[5] Sidgwick makes a similar point: "…it would, I conceive to be universally held that no act can be absolutely right, whatever its external aspect and relations, which is believed by the agent to be wrong. Such an act we may call 'subjectively' wrong', even though 'objectively' right" (Sidgwick1907, p.207).

[6] Ewing's position is actually more complicated than my outline suggests. There are two aspects to the subjective ought. One is its relation to an agent's perspective, which I took to be its defining feature. The other is its relation to blameworthiness. It is a substantial thesis on my account that blameworthiness is determined by the subjective ought in the following way: A person who does what she subjectively ought to do is not blameworthy if she thereby fails to fulfil her objective obligations. Ewing draws these relations differently. For him the subjective ought is defined by its connection with blameworthiness: ""Ought", both in philosophy and in ordinary discussion, is also used in a sense in which not to do what one ought, or to do what one ought not to do, is always morally blameworthy." About the idea that one ought subjectively to *fi* iff one believes that one ought to *fi* objectively, Ewing says: "To say that I ought to do A in this [second and subjective] sense is indeed not the same as saying that I believe I ought to do A, for the proposition that I ought to do what I believe I ought to do is synthetic, but it is, I think, synthetic a priori" (Ewing 1947, p.120). Let me also mention that in *The Definition of Good*, but not in his later work *Ethics*, Ewing talks about three senses of ought which are, first, what we really ought to do, secondly, what we believe we ought to do and, thirdly, what would be reasonable for us to believe that we ought to do. I have already restricted the Principle of Self-Reliance requires to cases in which one holds reasonable ought-beliefs.

closing off other options that would have been available, gives rise to the ideal of a unified account of normativity. This ideal expresses our conviction that what we do makes sense and can be justified on the basis of our reasons. Dividing the normative domain at the level of aggregating reasons runs counter to this conviction, as the unique sense of justifiability that we are after would be lost.

Am I not going too far when I talk about the commonly presupposed ideal of a unified account of normativity? Most philosophers would happily distinguish between theoretical rationality, rationality of belief, on the one hand and practical rationality, rationality of action, on the other. And once we have divided the normative domain, why not divide it further, if necessary? The practical pressure exerted by the division between an objective and a subjective ought is, after all, minimal.[7]

However, a closer look at the distinction between theoretical and practical rationality will, contrary to first appearances, support my scepticism about dividing the normative domain in the way suggested by Ewing. Think of an example in which the available evidence supports believing one thing, but believing another has great practical benefits. What ought one to believe? To say that in an epistemic sense of 'ought' one ought to believe the one, but thinking prudentially one ought to believe the other is ducking the question. It is like the uninformative advice that in your professional role you ought to do the one thing, but as a friend you ought to do something else. Usually this only states the problem, it does not solve it. You have to go one way or the other; that is the practical pressure you face when confronted with such a decision.

A common reaction to any apparent conflict between theoretical and practical rationality promises to answer this practical pressure. It suggests a separation of the domains to which these two distinct normative notions apply. Though it seemed that both sorts of reasons apply to what we ought to believe, we now distinguish between reasons for believing something, reasons captured by our notion of theoretical rationality, and reasons for bringing it about that one has a certain belief. The latter are practical reasons, reasons for doing as opposed to reasons for believing something. The plausibility of distinguishing between two notions of rationality depends, I want to suggest, on a separation of the domains to which they apply, beliefs in the one case and actions in the other. Only if beliefs and actions are different things, will this division be plausible. Thus, the division between theoretical and practical rationality is compatible with the ideal of a unified account of normativity. It simply renders the ideal domain specific. Its point is to express our unease with a situation in which reasons would point in different directions but are, as a matter of principle, not negotiable. Genuine moral dilemmas, if there are any, would be of that sort. But we do not expect a theory of rationality to engrain dilemmas into the very concept of rationality by allowing two notions of ought that apply to the

[7] Sidgwick, who, as it is well known, was worried about the split between duty and self-interest, was not nearly as worried about the split between objective and subjective ought because, as he put it, "…it can have only a limited and subordinate practical application. For no one, in considering what he ought himself to do in any particular case, can distinguish what he believes to be right from what really is so: the necessity for a practical choice between 'subjective' and 'objective' rightness can only present itself in respect of the conduct of another person whom it is in our power to influence" (Sidgwick 1907, p. 207).

same things. The price of accepting Ewing's solution would be to undermine the ideal of a unified account of normativity that is a presumption of all practical reasoning.

There is also meta-philosophical support for the position I am advocating. If we solved Ewing's Problem by disambiguating, what would prevent us from us from using such a strategy across a whole range of philosophical problems? Epistemological debates could be 'solved' by distinguishing between different senses of justification, and the debate between compatibilists and incompatibilists would be a debate between people who simply use different notions of freedom, none of which is any better than the other. Using disambiguation unrestrictedly would, it seems, dissolve philosophy itself. Each philosophical position about any matter could then be correct as it only spells out its own peculiar notions.

These considerations are not decisive. It might turn out that we cannot do any better than to accept two notions of ought. The disadvantage I pointed out, however, motivates the search for a different solution to the problem.

4. The Wide-Scope Solution.

The Principle of Self-Reliance says that if one believes that one ought to *fi*, one ought to *fi*.

(Self-Reliance) If Bel(O(fi)), then O(fi)

Understanding 'ought' as a sentential operator and reading 'you ought to *fi*' as 'you ought to ensure that you *fi*', we can distinguish the narrow-scope reading of the principle of self-reliance of above from a principle of self-reliance with a wide scope ought operator: You ought to ensure that, if you believe you ought to *fi*, you *fi*.

(Wide Scope SR) O(If Bel(O(fi)), then fi)

I will follow John Broome (2000) who calls such wide-scope ought sentences 'normative requirements' and contrasts them with narrow scope ought sentences the antecedents of which he calls 'reasons'.

When I talked about Ewing's problem as a problem within morality, I said it seems less severe than Ewing thought. In common sense morality, preserving one's integrity by acting in accordance with one's conscience is one consideration that determines what one ought to do, but it is not the only one: it can be *outweighed* by other things that happen. If we think of Ewing's problem as a problem of practical reasoning in general, this solution seems less attractive. If I have thought carefully about something and have come to the conclusion 'I ought to *fi*', what could outweigh this conclusion of my practical reasoning? The belief that I ought to *fi* is not a consideration on a par with others, like that *fi*-ing promises to achieve something I want. In coming to the judgment that I ought to *fi*, I take a stance on where the force of all relevant reasons lies. The belief that I ought to *fi* rests on all considered reasons and, thus, could not simply be outweighed by other considerations. Its normative force is thus like the force of a conclusive reason to *fi*.

Wide-scope ought sentences, i.e., normative requirements, offer a different way of weakening the principle of self-reliance. The conclusion of my practical deliberation cannot be outweighed, nevertheless we cannot infer from the existence of such a normative requirement and the fact that I believe that I ought to *fi* that I indeed ought to *fi*.

(Invalid) O [If Bel(O(fi)), then fi]
Bel (O(fi))
Thus, O(fi)

In analogy to modal sentences, a wide-scope ought sentence does not allow the detachment of its consequent, if its antecedent is given. Therefore, the principle of self-reliance, understood as a normative requirement, does not entail infallibility in normative matters.[8]

5. How the Problem Comes Back

Consider the following example[9]: If you believe that there might be traffic on the main road ought you not to stop or at least to slow down your car before entering the main road? If we understand this obligation as a normative requirement, it asks you to ensure that, if you believe there might be traffic, you stop. The fact that you believe that there might well be traffic on the main road, does not allow us to conclude that you ought to stop. The normative force of the ought-sentence has been weakened, and all we have is a disjunctive obligation: either stop believing that there might be traffic or stop the car. Driving instructors, however, simply tell you to stop, and, intuitively, rightly so.

Can we square this intuition with modal and deontic logic without simply insisting on a narrow-scope reading? What is noticeable in Prichard's example is that it would be crazy not to believe that there might be any traffic on the main road. After all, it is the main road. In believing that there might be oncoming traffic, you believe as you ought to believe. It seems that the analogy with modal logic now works against the wide-scope solution. If we ought to have certain beliefs, then the conclusion that we ought to stop might be detachable after all and Ewing's problem comes back. The crucial question is whether the following inference is rightly called 'Valid'.

(Valid) O [If Bel (O(fi)), then fi]
O [Bel (O(fi))]
Thus, O(fi)

Chisholm (1963) has challenged the above inference schema. Suppose you ought to visit your neighbours, and suppose, furthermore, that you ought to, if you visit them, call in advance to advise them of your coming. If above inference schema is valid, it would follow that you ought to call them. But what, Chisholm asks, if, your obligation notwithstanding, you will not visit them? Would it not aggravate the

[8] This reading of the Principle of Self-Reliance has been proposed as a response to Ewing's Problem in Dancy (1977) and in Gensler (1985).
[9] I take this example from HA Prichard (1949).

situation if you nevertheless called and announced your visit? I agree – this would be naughty.

The salient difference between Chisholm's and Prichard's examples is that in Prichard's example you fulfil the obligation you are under, i.e., you believe that there might be traffic on the main road, whereas you do not fulfil your obligation to visit the neighbours. If you did, we think, you certainly ought to call them beforehand. I do not mean to suggest that adding the premise Bel(O(fi)) would render an invalid inference valid. No, the inference is valid as it is, but we have to understand what it means:

> All deontically perfect worlds are such that if Bel(O(fi)) holds, then fi does as well. Bel(O(fi)) is the case in all deontically perfect worlds. Therefore, fi holds in all deontically perfect worlds.

The lesson to draw from Chisholm's counterexamples is that deontic logic tells us something about the structure of deontically perfect worlds and thus remains silent about contrary-to-duty worlds. The fact that you do believe that there might be traffic on the main road assures us that in all relevant respects you are in a deontically perfect world. Thus, 'Valid', properly understood, seems safe. We have to look for a solution of Ewing's Problem that goes beyond differences of scope.

6. Moore's Solution: Objectivism.

The principle of self-reliance, according to which our actions should be guided by our beliefs about what we ought to do, and the idea that our normative beliefs are fallible seem to be in conflict with each other. If they are, we have to abandon one of the two principles and develop a theory of what we ought to do solely on the basis of the other. GE Moore's theory of rationality sticks to fallibility and rejects the Principle of Self-Reliance.

In his book *Ethics*, Moore takes up the following objection to his version of utilitarianism. What we ought to do is, according to this objection, not what actually will have the best consequences but what we can reasonably expect to have the best consequences. We have already met one example that supports this objection, namely Prichard's case in which, intuitively, we ought to stop the car before we enter the main road. It is a precaution; being home a little bit earlier simply is not worth the risk. One ought to stop regardless of whether there is traffic or not. Moore would disagree. He sticks to his view: we ought to do what is actually best, and if, having stopped, we realize that, against all odds, there was no traffic, then we did not do what we ought to have done.

Karl Kraus once said 'Whenever you are in doubt about what to do, simply do the right thing'. There is something funny about giving advice like this. Is Moore simply not susceptible to Kraus's humour? 'Simply do the right thing!' sounds like advice, but it cannot play the role of advice because if the agent knew which action was right, no advice would be needed. Useful advice would need to help the agent to recognize which of his options, let us call them A1 or A2, is right.

Frank Jackson has objected to Moore's theory on the grounds that he fails to see the funny side of Kraus's remark. In other words, his theory of rationality does not give us any useful advice, and thereby it fails to do what a theory of rationality, presumably, sets out to do. John Broome thinks that Moore can be forced into giving substantial advice. But, as a result, Moore's objective account of ought is undermined. Jackson (1991, p.467) says, "We need, if you like, a story from the inside of an agent to be part of any theory which is properly a theory of ethics, and having the best consequences is a story from the outside." What is Moore's inside story?

Here is John Broome's attempt to extract such a story from Moore. Imagine a case in which you, facing a choice between A1 and A2, have overwhelming evidence that A1 is by far the better option. For Broome it is clear that the agent ought to stop, whereas Moore would say: "We can't really be sure what is best, it depends." Broome (1991, p. 128) replies, "Impatient with Moore's shilly-shallying, you ask him 'So what to you suggest I do, then? A1 or A2?' Pressed like this, Moore will certainly tell the agent to do A1. This is not simply an ungrounded whim on his part. He believes that practical reason, given the probabilities (on which Moore and the agent agree), requires the agent to do A1; it would be irrational on the agent's part not to. Another way of expressing this belief of Moore's is that you ought to do A1."

Broome is right. Moore would say that you ought to do A1. And Broome is right again that this is no whim on Moore's part: Moore does believe that the agent ought to do A1; A1 after all does look much better. But, according to Moore's theory, neither the agent's nor his own belief that he ought to do A1 is relevant for the question of whether he ought to do A1 or not. Moore (1912, p. 82) says, "… we may be justified in saying many things, which we do not know to be true, and which in fact are not so. And so in this case I do not see why we should not hold, that though we would be justified in saying that he ought to choose one course, yet it may not be really true that he ought."

To understand Moore's response we have to distinguish between Moore's theory of what we ought to do and the employment of the theory in a particular case. Simplifying somewhat, Moore's theory is that we ought to do what is best. Before the facts are in, it is not settled which of A1 and A2 is best, thus it is not settled what you ought to do. The objection was that if, even in clear-cut cases, that is all the theory says, then it fails as a theory of rationality because it does not tell us what to do in the sense of which of A1 and A2 we ought to choose. But the objection fails. Moore's theory can be employed. Advice is given according to one's beliefs about which action, A1 or A2, is best. And Moore, like everyone else, reasonably believes that A1 is best, or as we should rather say, that A1 will turn out to be best. But even the best advice, Moore insists, is fallible. The fact that in employing the theory in the circumstances of our example we get the result that you ought to do A1, does not entail that you ought to do A1, because in employing it, we have to rely on fallible beliefs. The crucial point is that employing the theory does not generate its own 'ought'. Employing it means thinking about what we ought to do. The result of even faultlessly employing the theory, within the informational constraints of the example,

is not a normative fact, a fact that one ought to do this or that, the result is always a belief, a fallible belief that one ought to act in a certain way.[10]

The first point we gain from this discussion is the following: Moore's theory is not impractical; Moore's theory can be employed. The second point is that Moore's theory avoids Ewing's problem. The principle of self-reliance says that an agent ought to act in accordance with his or her reasonable normative beliefs.[11] Moore would reject the principle of self-reliance. Even reasonable beliefs can be false. Is there anything he puts in its place? Not much. Moore can say that when you act in accordance with your belief about what you ought to do, you act in accordance with the best advice you can give yourself. But whether you indeed ought to act on the best advice you can give yourself is an open question.

Moore's view will not be dislodged by the objection raised by Jackson and Broome, but our discussion has highlighted an important point. If what we ought to do is always to bring about the best consequences, then, by acting as we ought to act, we will always act successfully. Understanding facts about what we ought to do as facts about where the balance of reasons lies and, thus, as facts about rationality, we loose the gap between what is rational and what is successful. This is, in fact, Moore's view. "The only possible reason that can justify any action," Moore (1993, p. 153) writes, "is that by it the greatest possible amount of what is good absolutely should be realized". And in *Ethics* (Moore 1912, p. 73) he says that the notions of expediency and of duty – which just are the notions of the successful and of what we ought to do – will always apply to the same action.

We could impose a different conceptual structure on Moore's theory in which, by definition, to act rationally is to act on the basis of one's belief about what will be best.[12] This would open up a gap between rationality and success. But then rationality would be an aspect of the theory's employment and we have seen that employing the theory doesn't generate its own 'ought'. Therefore a Moorean account of what we ought to do leaves us with the following alternative: either rationality is success or rationality is not normative, i.e., it need not be true that what we ought to do and what is rational to do coincide. Moore himself accepts the first alternative. In order to solve Ewing's Problem we abandon the principle of self-reliance, and, consequently, we lose the difference between what is rational and what is successful.[13]

---

[10] The situation becomes more complicated once we have accepted the lesson of decision theory, namely that rationality does not go by belief about what is best but by the highest expectation of goodness. In the present context this insight would change what it means to employ the theory correctly. Nevertheless the general issue re-emerges as the question whether we ought to do what has the highest expectation of goodness or what we reasonably believe to have the highest expectation of goodness. For a discussion of this problem see Piller (2000b).

[11] I say more about the restriction to reasonable starting points in Piller (2001). There I argue that only what does not itself violate standards of rationality can justify something else or make something else rational. This is a general fact that holds for rationality in general, be it practical or theoretical.

[12] Parfit (1997, p. 99) takes such a line. "As rational beings, we can ask: What do we have most reason to want, and do? What is it most rational for us to want, and do? These questions differ in only one way. While reasons are provided by the facts, the rationality of our desires and acts depends instead on what we believe, or – given the evidence, ought rationally to believe."

[13] Kolodny (2005) has recently defended Moore's view. He opts for the second alternative open to Moore and, consequently, denies that rationality is normative. I discuss Kolodny in my paper "The Normativity of Rationality" available at http://www-users.york.ac.uk/~cjp7/ .

7. Supporting Self Reliance: Subjectivism

I have called Moore's view 'objectivism' because it tells us that an action ought to be done if and only if it actually has the best consequences. In contrast, the Principle of Self-Reliance determines what one ought to do by solely relying on the agent's perspective, and 'subjectivism' is thus an appropriate label.

I have introduced the Principle of Self-reliance on the basis of the following intuition: Suppose someone has thought carefully about what to do. From reasonable starting points he reasons to the reasonable conclusion: 'I ought to fi'. Wouldn't it be highly irrational for this agent to do anything but to fi? I am certainly not alone in thinking that it would indeed be irrational for anyone to act against one's own reasonable conviction about what one ought to do. Tim Scanlon (1998, p. 25) writes, "Irrationality in the clearest sense occurs when a person's attitudes fail to conform to his or her own judgment: when, for example, … a person fails to form and act on an intention to do something even though he or she judges there to be overwhelmingly good reason to do it". Scanlon is here trying to say something everyone can agree with.[14]

The Principle of Self-Reliance has strong intuitive appeal and, furthermore, it is widely agreed upon. Thus, we might just accept it as a basic principle. An argument for the principle, however, would promise to deepen our understanding of it as it might explain why we find the principle so intuitively appealing.

If there is practical reasoning, then there has to be an endpoint to at least some of its instances, i.e., some processes of practical reasoning concerned with particular decision problem will yield practical conclusions. One legitimate way of expressing the conclusion of a process of practical reasoning is to say 'I ought to fi'. Thereby the

[14] Some philosophers, however, seem to disagree. In her paper 'On Acting Rationally against One's Best Judgment', Nomy Arpaly (2000) develops an idea Harry Frankfurt (1988, p.189) has expressed as follows: "Someone who made up his mind to sacrifice the world in order to spare his finger would thereby give a convincing indication of severe mental disorder. But the indication would be considerably more grave if he not only made this judgment but also showed that he was capable of actually carrying it out." Arpaly's main thesis is that if an agent's belief about what she ought to do is unreasonable, the agent will be even more unreasonable if she actually acts on her unreasonable conviction. She argues that it might be the reasons one overlooked, when forming the unreasonable ought-judgment, that actually prevent one from carrying out one's unreasonable plans. Note however that this view does not attack the Principle of Self-Reliance as I have introduced it here, because, as indicated above, justifying force only attaches to something that does not itself violate requirements of rationality or, in other words, beliefs about what one ought to do have to be reasonable for the Principle of Self-Reliance to apply. Taking this into account, we find that even Arpaly agrees with Scanlon's position. She says "… that every agent who acts against her best judgment is, as an agent, less than perfectly rational, as the schism between best judgment and desire indicates a failure of coherence in her mind" (Arpaly 2000, p. 491). Jonathan Dancy, who has recently revived Moore's objectivism in his *Practical Reality* (Dancy 2000), might be thought to be another critic of the view. His main thesis is that only facts, as opposed to elements of an agent's perspective on the world, are reasons. Dancy, however, endorses the wide-scope reading of the Principle of Self-Reliance. If what I say in section 5 is correct, no more is needed for Ewing's problem to arise. For a comparison between Dancy's and Moore's forms of objectivism see Piller (2003).

agent expresses her judgment or belief that she ought to fi.[15] This solves the decision problem for her. Now she knows what to do.

Remember what Moore, who denied the Principle of Self-Reliance, said about normative judgments. Even if we are justified in believing that we ought to *fi*, we might still be wrong. Thus, whether one ought to act on ones' judgment about what one ought to do is an open question. If one's judgment is true, one ought to act on it; if it is false one ought not to act on it. Practical reasoning, I said above, terminates in judgments about what one ought to do. An agent who accepted Moore's position would, consequently, have to regard it as an open question whether he really ought to do what he thinks he ought to do. But if that were still an open question for her, then the agent's belief that she ought to *fi* could not rightly be thought to be the state of the agent that concludes his practical deliberation. If the matter has to be regarded as being still open, even after this belief has been reached, the whole process has not been concluded. If 'I ought to *fi*' is not the conclusion of practical reasoning, then nothing will be. And if practical reasoning can never have a conclusion, there will not be any practical reasoning. Accepting the Principle of Self-Reliance, therefore, turns out to be a condition without which there could not be practical reasoning.[16]

Could Moore claim that acting on reasonable beliefs about what one ought to do has a higher chance of success than acting otherwise, and is thereby vindicated? For reasons of consistency, Moore will always have to deny that any fact that falls short of success, including the likelihood of success, generates its own ought. The only ought he accepts is the ought bound to successful action.

Moore could say that sometimes we ought to engage in practical deliberation, namely then when it actually is best to do so. Deliberating, for example, might postpone the time of action in a way that turns out to be beneficial or it might simply be fun. What Moore cannot say, however, is that practical deliberation is an essential aspect of reflective rational agency. Practical reasoning, on Moore's account, is an empty game. Its emptiness is witnessed by the fact that 'a conclusion of practical reasoning', namely an agent's belief about he or she ought to do, is bare of any normative significance. This, I think, amounts to a denial of practical reasoning. Thus, the acceptance of the Principle of Self-Reliance subjectivism turns out to be a pre-condition without which there could not be practical reasoning as we understand it.

---

[15] As long as the state one is in when one has reached the endpoint of practical deliberation can be conjoined with the judgment that one ought to fi, it does not really matter what the nature of the state actually is, be it a decision, the formation of an intention, a willing, a preference, a desire to act, or a judgment about what one ought to do. David Lewis has famously argued that the split between desires and beliefs, captured by the different ways in which expectations and beliefs evolve, makes the thesis of a necessary conjoinment of conative and doxastic states impossible. For my reasons for disagreeing with Lewis on this matter and for further references to this debate see Piller (2000a).

[16] We do not seem to face an analogous problem in the case of theoretical reasoning. The conclusion of a piece of theoretical reasoning is usually not that one ought to believe that p but simply that p. A theory of theoretical rationality tells us what we ought to believe and in concluding that p we might believe as we ought to. If, however, a 'preliminary' conclusion of our reasoning is that believing that p is best supported by our evidence or that we ought to believe that p, then we face the same problem. Not only practical reasoning but also theoretical reasoning would need to be made 'practical'. Scanlon (1998, p. 25) supports an analogous principle for the theoretical case: we ought to believe what we think we ought to believe.

8. Butler's Subjectivism

Moore's theory of practical rationality was easily summarized: 'Do what is best'. Butler's account can be captured by the slogan 'Follow human nature!'. Moore focuses on facts brought about by the agent, whereas Butler's theory looks 'inside' the agent to determine what he or she ought to do.

According to Butler, animals act on instinct and humans are also animals. Our instincts, passions and desires form the basic level of reasons for actions. When hungry, we have reason to eat, and when we are in rage, we have reason to destroy. But instincts and passions do not exhaust human nature. Humans can reason, which on the first level only means that we can think about what overall would satisfy our passions best. Butler calls the principle that captures such reasoning 'prudence' or 'self-love'. What if self-love advises you not to act on, for example, your rage? For Butler the resolution of such a conflict is not a matter of the relative strength of the particular passion on the one side and prudence on the other: "… there has to be some other difference or distinction to be made between these two principles, passion and cool self-love, … and this difference, not being a difference in strength or degree, I call a difference in *nature* and in *kind*" (Butler, 1991, p. 352). Butler suggests that it is in the nature of the principle of self-love that it has normative authority over the passions. Self-love, after all, respects the normative force of all passions; what it does is to calculate their best overall satisfaction. It is not a passion alongside rage or hunger; it is principle of aggregating the normative forces that arise from passions and desires. Self-love resides on a higher level of practical reflection and this fact gives it its authority.

There are different ways of aggregating the reasons arising from our passionate nature. The principle of self-love determines one of them, the principle of benevolence another. Again, resolving a conflict between different principles of aggregation is not a matter of power but of authority. Our conscience, which Butler also calls the principle of reflection (and I have called Principle of Self-Reliance), is the ultimate normative authority. "That principle, by which we survey, and either approve or disapprove our own heart, temper and actions, is not only to be considered as what is in its turn to have some influence; which may be said of every passion, of the lowest appetite: but likewise as being superior; as from its very nature manifestly claiming superiority over all others: insomuch as you cannot form a notion of this faculty, conscience, without taking in judgement, direction, superintendency… Had it strength, as it has right; had it power as it has manifest authority; it would absolutely govern the world" (Butler 1991, p. 354). Conscience is not our ear for God's voice. Or even if it were, that is not the source of its authority. It has authority because it is different from all other sources of practical reason in nature and in kind. It stands on the last level of practical reflection. Once all the reasons are in, a person has to decide, thereby either approving or disapproving of 'our own heart and temper' and thereby also approving or disapproving of different aggregation rules for these reasons. The fact that we can act reasonably or, more cautiously, the acceptance of the ideal of a unified account of normativity, assures us of, or commits us to, this highest level of practical reflection. Conscience or reflection rules supreme because nothing could oppose it. Conscience is whatever decides the conflicts that occur on a thereby lower level of reflection. "Every bias, instinct, propension within, is a real part of our

nature, but not the whole: add to these the superior faculty, whose office it is to adjust, manage, and preside over them and take in this its natural superiority, and you complete the idea of human nature" (Butler 1991, p. 356).

Butler argues that the structure of practical deliberation and its hierarchical order show us that the Principle of Self-Reliance is a pre-condition of autonomous agency. Acting autonomously, we ourselves determine what we ought to do. The nature of full-fledged human action is such that 'every man is naturally a law to himself'. The law is the Principle of Self-Reliance: acting in its paradigmatic sense, comes from the agent, i.e., from the agent's conviction of what he or she ought to do.[17]

9. Exploring the Principle of Self-Reliance

According to subjectivism, normative facts may vary with varying perspectives. The same action can be demanded or forbidden depending on how the circumstances are conceived by different agents (or the same agent at different times). Suppose someone has, unbeknownst to you, put some poison in your orange juice. Seeking refreshment, you reasonable think you ought to drink. But I know about the poison, thus I would firmly say that you ought not to drink. Am I not right?

It seems as if subjectivism is committed to the view that we are both right. Subjectivism won't find any fault with my judgment, and as long as your beliefs about the situation are reasonable, your conclusion, you ought to drink, will be reasonable as well. According to the Principle of Self-Reliance you then, indeed, ought to drink.

There are two problems with this view. One is its consistency: How can it be that your drinking is both such that it ought, and that it ought not, to be done? The subjectivist faces the task of explaining the semantics of ought-sentences, in particular the role perspectives play. If one says ought-sentences are only semantically complete if they mention a perspective, I cannot *disagree* with you about what to do. It would be like fighting about whether it is nice here or not when we refer to different locations. If one identifies agents with their perspectives, I cannot disagree *with you*, as I am talking about what a different person, namely you, if you knew about the poison, ought to do. Subjectivism inherits the relativist's problem of explaining what looks like genuine disagreement within a subjectivist framework. The second problem is that in the example above, we think that my perspective is superior to yours, as I just know more than you do. The subjectivist is challenged either to undermine this intuition or to reconstruct this notion of superiority in terms compatible with subjectivism.

These are important tasks, which I register without pursuing them further. The example of the poisoned orange juice could also be used to launch a direct challenge

[17] I have chosen Butler Sermons as historical support for the principle of Self-Reliance partly because Butler's argument bears some similarity to my argument above. In the context of Ewing's Problem, i.e. in the context of the tension between self-reliance and fallibility, we find strong support for self-reliance in Kant and his tradition. Kant, Fichte, and Fries all denied the possibility of an 'erring conscience', though their arguments as well as their views of what it would mean for conscience to err differ. For a summary of their views and for detailed references see Wood (1990, pp. 174-192).

to the Principle of Self-Reliance. This challenge seeks to locate an inconsistency within a perspective. Suppose I reasoned as follows: I know that you ought not to drink the juice. If I accepted the Principle of Self-Reliance, I would have to conclude that you ought to drink. Thus, the Principle of Self-Reliance is false.

If above criticism were valid, the following would be so as well. You reason from *p* and *if p then q* to *q*. Let us suppose that I know that *q* is false. If I said that I had to reject *modus ponens*, because it would commit me to believe something that I know to be false, my mistake would be to think that *modus ponens* is relevant to my reasoning. All I know is that you believe that *p* and that you believe that *if p then q*, but I cannot use *modus ponens* to derive anything from these premises. Similarly in the case in question: The Principle of Self-Reliance puts you under an obligation to *fi*, as you believe that you ought to *fi*, but it does not commit me, as an observer, to anything. Like *modus ponens* and other rules of reasoning, the Principle of Self-Reliance only applies to someone who fulfils the relevant conditions, which in the first case is to believe the two *modus-ponens* premises, and in the second case is to believe that one ought to *fi*.

You express your reasoning along the lines of *modus ponens* as follows: 'p, if p then q, therefore, q.' You can use *modus ponens* because you have the relevant beliefs in its premises. Your believing, however, is not part of the premises that occur in your reasoning. Similarly, you express your reasoning governed by the Principle of Self-Reliance as follows: 'I ought to *fi*. Thus, I *fi*.' Again, the belief that you ought to *fi* makes the Principle of Self-Reliance available to you, but this belief is not the premise from which you reason. (The reasoning expressed by 'I believe that I ought to *fi*, thus I ought to *fi*', is not an instance of the Principle of Self-Reliance. Such reasoning leads from an awareness of believing that p to the belief that p.) The Principle of Self-Reliance tells one to *fi*, if one believes that one ought to *fi*. It does not license any inference from my belief that you believe that you ought to *fi*. It is obvious that one does not have to reject *modus ponens*, in order to criticize a person's belief that this person has arrived at by using *modus ponens*. Similarly, if I tell you not to drink, I do not criticize your use of the Principle of Self-Reliance. I disagree with you about the premise from which your reasoning starts.

10. Could Subjectivism be True?

Is there any hope for a theory that sticks with the Principle of Self-Reliance and denies fallibility? One could not object simply by saying: believing something to be true does not make it so. Remember that only reasonable beliefs can justify. Thus, all the principle of Self-Reliance tells us is that if I reasonably believe that I ought to fi, then I ought to fi. The same move answers another objection, one that appeals to the priority of normative facts over normative beliefs. If the belief that one ought to fi would itself constitute the fact that one ought to fi, there would be nothing this belief could be about. But remember that not any belief that one ought to fi can make it the case that one ought to fi, only a reasonable belief can. Thereby the usual subject matter for this belief to be about is, arguably, put back in place. This subject matter is the conditions that have to be met for it to be reasonable to believe that one ought to fi. Principles of practical reason, like Kant's Imperatives, would still constitute this subject matter.

But even if some objections do not work as straightforwardly as might have been expected, a simple look at the principle alone will for many be enough of an objection. In the case of normative beliefs their reasonableness and their truth cannot come apart. That is, after all, exactly what the principle says. And do we not all know that they can come apart?

11. Why the Problem is Hard

We have two sets of notions. On the one hand we talk about rationality, reasons, and what ought to be done, on the other hand we have the notions of truth, fact and success, and we want a contrast between these domains of the normative, on the one hand, and the factual on the other. Rational beliefs need not be true, but in calling them rational we suggest that they have been formed in a way that is conducive to their being true. Similarly, rational actions need not be successful; rationality is our guide to success.

This distance between the factual and the normative is essential for there being a contrast between them and, thus, for being able to differentiate between the two domains. It is essential for the normative to step back from the factual. But, whatever our account of rationality, duty or what ought to be done, is, it will specify some facts that make it the case that the normative notion applies. Now we have two sets of facts that are both candidates for having genuine normative force. If we stick to the ideal of a unified account of normative force, consistency problems will arise, as illustrated by Ewing's Problem.

The principle of fallibility locates genuine normative force in facts removed from what I called the normative domain, whereas self-reliance is an example of investing the principles of practical reason with genuine normativity. But, as Ewing's Problem shows, we cannot have it both ways. The attempts to solve the problem that I have discussed so far, try to negate the normativity of one of the domains. According to Moore's position there is no success-independent notion of rationality. All normative force resides in such facts. The normative domain at a distance from these facts has been lost. Thus, in a sense, there is no practical reasoning and no practical rationality. Butler's position, which starts from the Principle of Self-Reliance, holds on to the normative domain at a distance and, consequently, denies that the normative can be outrun by the facts Moore appeals to. The problem is that if all normative force remains at a distance from Moore's facts, there seems to be too little to guide it.

12. A Way Out?

Both of the options I have offered so far have, in philosophical terms, drastic consequences. If we follow Moore, we have to give up on theories of reasoning and rationality. If we follow Butler, we have to accept that every one of us, as long as one remains reasonable, is infallible when it comes to what one ought to do. There seems to be, however, a third way the consequences of which, by contrast, seem less revisionist. This attempt to solve the problem exploits the difference between the following two claims:

(a) It is irrational to act against one's own normative judgment.
(b) It is rational to act in accordance with one's own normative judgment.

Claim (a) expresses what we could call the basic intuition behind the principle of self-reliance, the intuition we find, for example, in Scanlon. Claim (b) spells out the principle of self-reliance in positive terms, i.e., in terms of what is rational to do. We can move from claim (a) to claim (b) if we accept a further principle, namely that each decision problem has a rational solution, i.e., in each situation there is something I ought to do. What if we abandoned this principle?

The Principle of Self-Reliance in its original formulation provided a sufficient condition for what one ought to do.

(Self Reliance) If Bel(O(fi)), then O(fi).

We could instead claim that the normative belief is not a sufficient but only a necessary condition for the corresponding normative fact. Let us call this the Weak Principle of Self-Reliance.

(Weak Self-Reliance) If O(fi), then Bel(O(fi)).

The Weak Principle still explains what I have called above the Basic Intuition.

(Basic Intuition) not[O(fi)&Bel(O(psi))].

Rewriting the basic intuition, we get:

If O(fi), then not[Bel(O(psi)].

If we assume that the agent is consistent as well as reflective, i.e., we assume that the agent has relevant normative beliefs about what he ought to do, then the re-written basic intuition entails the Weak Principle of Self-Reliance: If O(fi), then Bel(O(fi)). Note furthermore that the basic intuition in the form 'If Bel(O(fi)), then not[O(psi)]', plus the assumption we have abandoned, namely that in each situation there is something one ought to do, which, simplifying somewhat, reads 'If not[O(psi)], then O(fi)', entails the Principle of Self-Reliance in its original formulation: If Bel(O(fi), then O(fi). That is why I speak of *Weak* Self-Reliance: abandoning the view that there always is something we ought to do, we get a principle that still explains the Basic Intuition.

The following picture emerges. (In what follows I leave the non-reflective agent aside, as Ewing's problem is a problem of reflection.) Normative beliefs, like my belief that I ought to *fi*, can be true or false. If it is false, then, obviously, there is no normative fact as that I ought to *fi*, but neither is there any other normative fact. If my belief that I ought to *fi* is true, then indeed I ought to *fi*. The Weak Principle of Self-Reliance tells us that believing that I ought to *fi* is a necessary condition for the normative fact that I ought to *fi* to obtain. I, as reflective agent, could not be under an obligation to *fi* without being aware of it.

Let us just assume for a moment that this solution of Ewing's Problem is correct. What could be learnt from it? First, there are decision problems that have no rational solution, and that is not because they are too hard but because the agent is in error about what he or she ought to do. And, second, although our normative beliefs can be false, there cannot be any hidden normative facts. As long as I have any beliefs about what I ought to do, it will never be the case that I ought to do something without me knowing that this is so. It sounds almost too good to be true.

Weakening Self-Reliance offers a somewhat shallow solution to what looked like a deep problem. So let me try to express the problem one more time. One the one hand we see ourselves as autonomous and self-governed agents. This view does not exclude accepting normative authorities. Something can, however, only have normative authority for me, if I accept its authority, which, in the end, would make what I ought to do depend on myself. The idea of being essentially self-governed gives us 'formal' authority over what we ought to do. This is what the Principle of Self-Reliance tries to capture. On the other hand, we also think that that there are things we ought to do and other things we must not do, independently of an agent's perspective on these matters. The third way, i.e. weakening the Principle of Self-Reliance, does not fully capture our commitment to normative facts. It does endorse our fallibility but understands it as something like an excusing condition: being wrong about what I ought to do renders any ought inapplicable to me. This way of understanding the problem points to a fourth reaction to Ewing's Problem. The first, and this was Ewing's own solution, was to give up on the unity of practical normativity. The second was to give up on practical normativity altogether, which, I have argued, is the consequence of Moorean objectivism. The third was to limit the applicability of ought by weakening self-reliance. The fourth solution would consist in making good on the Kantian programme. Then, self-determination is not purely formal. By its very nature, it would also provide us with the substance of what we ought to do.[18]

REFERENCES

Arpaly N. (2000), "On Acting Rationally against One's Best Judgment", *Ethics* 110, pp. 488-513

Broome J. (1991), *Weighing Goods*, Oxford: Blackwell

Broome J. (2000), "Normative Requirements", in Dancy 2000, pp. 78-99

Bishop Butler (1991), "Fifteen Sermons", in DD Raphael ed., *British Moralists 1650-1800*, Vol. I, Indianapolis: Hackett, pp. 325-377

[18] I want to thank audiences at the University of Valencia, Oxford University, University of Graz, University of Manchester, University of Glasgow, and the University of Helsinki for interesting discussions. For their helpful comments I want to thank John Broome, Krister Bykvist, Jonathan Dancy, Steve Everson, Matt Matravers, Samir Okasha, Philip Percival, Jim Pryor, John Skorupski, David Sosa, Peter Strasser, and Ralph Wedgwood.

Chisholm R.M. (1963), “Contrary-to-Duty Imperatives and Deontic Logic”, *Analysis* 24, pp. 33-36

Dancy J. (1977), “The Logical Conscience”, *Analysis* 37, pp. 81-84

Dancy J. (2000), *Practical Reality*, Oxford: Oxford University Press

Dancy J. ed. (2000), *Normativity*, Cambridge: Blackwell.

Ewing A.C. (1947), *The Definition of Good*, New York: Macmillan

Ewing A.C. (1953), *Ethics*, London: English Universities Press

Frankfurt H. (1971), “Freedom of the Will and the Concept of a Person”, *Journal of Philosophy* 68, pp. 5-20.

Frankfurt H. (1988), “Rationality and the Unthinkable”, in his *The Importance of What We Care About*, Cambridge: Cambridge University Press, pp. 177-190.

Gensler H. (1985), “Ethical Consistency Principles”, *Philosophical Quarterly* 35, pp. 156-170.

Jackson F. (1991), “Decision-theoretic Consequentialism and the Nearest and Dearest Objection”, *Ethics* 101, pp. 461-482

Kolodny N. (2005), “Why Be Rational?”, *Mind* 114, 509-563.

Moore GE (1993), *Principia Ethica* , ed. by Thomas Baldwin, Cambridge: Cambridge University Press

Moore GE (1912), *Ethics*, London: Williams & Norgate

Parfit D. (1997), “Reasons and Motivation”, Proceedings of the Aristotelian Society, Supplementary Vol. 71, pp. 99-130.

Pettit P. & Smith M. (1993), “Practical Unreason”, *Mind* 102, pp. 53-79.

Piller C. (2000a), “Doing What is Best”, *Philosophical Quarterly* 50, pp. 208-226.

Piller C. (2000b), “Belief and Expectation: A Problem of Practical Reason”, in B Brogaard ed., *Rationality and Irrationality*, Kirchberg: Austrian Ludwig Wittgenstein Society, pp. 79-83.

Piller C. (2001), “Normative Practical Reasoning”, *Proceedings of the Aristotelian Society*, Supplementary Vol. 75, pp. 195-216.

Piller C. (2003), “Two Accounts of Objective Reasons”’, *Philosophy and Phenomenological Research* 67, pp. 444-451.

Prichard H.A. (1949), “Duty and Ignorance of Fact”, in his *Moral Obligation*, Oxford: Oxford University Press, pp. 18-39

Sartre JP. (1958), *Being and Nothingness*, London: Methuen

Sidgwick H. (1907), *The Methods of Ethics*, 7th ed., London: Macmillan

Scanlon TM (1998), *What We Owe to Each Other*, Cambridge/Mass: Harvard University Press

Van Willigenburg T. (2005), “Reasons, Concerns, and Necessity”, *European Journal of Analytic Philosophy* 1, 75-87

Wood A. (1990), *Hegel’s Ethical Thought*, Cambridge: Cambridge University Press