

How Outlandish Can Imaginary Cases Be?

JAKOB ELSTER

ABSTRACT *It is common in moral philosophy to test the validity of moral principles by proposing counter-examples in the form of cases where the application of the principle does not give the conclusion we intuitively find valid. These cases are often imaginary and sometimes rather 'outlandish', involving ray guns, non-existent creatures, etc. I discuss whether we can test moral principles with the help of outlandish cases, or if only realistic cases are admissible. I consider two types of argument against outlandish cases: 1) Since moral principles are meant for guiding action in this world, cases drawn from other worlds are irrelevant. 2) We lack the capacity to apply our intuitive moral competence to outlandish cases. I argue that while the first approach is importantly flawed, the second approach is plausible, not because our moral competence per se is limited to cases from this world, but because we lack the capacity to imagine outlandish cases, and we cannot apply our moral competence to a case we fail to imagine properly.*

1. Introduction

The following type of philosophical dialogue should be familiar from both conference and classrooms discussions and published writings:

- Philosopher A: 'Based on the previous analysis, I propose the following moral principle
P: Actions of type X are permissible if and only if conditions x, y and z are met.'
- Philosopher B: 'While your analysis seems sound, P must be rejected, because here is a counter-example in the form of a case where conditions x, y and z are met, but because conditions f, g and h also obtain, we have the intuition that actions of type X are impermissible. I grant that in real life, conditions f, g and h will never obtain, but the case I propose is nevertheless not inconceivable.'

At this point, depending on her philosophical temperament, philosopher A might either admit defeat, since she accepts the method of testing moral principles against intuitions about cases; or she might reply, even if she accepts this general method: 'My principle P can only be tested by means of intuitions about situations which can occur in this world and it is thus not refuted by intuitions about an outlandish case involving conditions such as f, g and h, which will in all probability never obtain.'

This article is an attempt to examine the validity of this reply. Indeed, anyone familiar with contemporary moral philosophy will have noticed that the cases used to elicit moral intuitions are often quite exotic, ranging from the use of non-existent technical devices such as ray guns or neuroscientific mind control, through features of human biology or society which are not and perhaps cannot be found in real life, to situations which, though they might occur in the world as we know it, are highly strange and unusual. This feature of analytic ethics has become so common that many philosophers hardly notice

it and if they do, find it unproblematic. Other philosophers, however, react against the use of such cases, and argue either that we should only use actual cases from real life or, more often, that while we can use imaginary cases in ethics, these cases must be such that they could (with greater or smaller probability) occur in our world as it is.¹ We can thus distinguish three paradigmatic positions as to the kind of case we may use as intuition-pumps:

- **Conceivabilism:** As long as a case is conceivable, it is legitimate to use it to elicit intuitions for testing moral principles.²
- **Realism:** Only cases which could plausibly occur given the world as it is should be used to elicit intuitions.³
- **Actualism:** Only cases which have actually occurred should be used to elicit intuitions.⁴

For the purpose of this paper, I will set actualism aside, and principally consider the debate between conceivabilism and realism. One reason for doing so is that the use of imaginary cases in ethics has some clear advantages, notably that it allows for what Frances Kamm calls ‘the technique of equalizing cases’, where we vary a given case in precise ways in order to isolate morally relevant factors;⁵ real life cases will be so rich in detail, and differ between each other in so many details that they cannot be used in this way.⁶ However, by examining the problems of conceivabilism, I hope to also show that we need not adopt actualism in order to avoid these problems.

I will consider two ways of defending the realist position.⁷ First, a defence built on the nature and role of moral principles: on this view, since the purpose of moral principles is to guide us in *this* world, they have a limited scope and are not meant to be valid in all possible worlds; hence they cannot be refuted by means of intuitions about far-fetched cases. I will argue that this approach, which is common in the literature, is flawed in important ways. The second and more promising line of defence builds on the nature of our competence for forming moral intuitions: on this view, the capacity we have for forming correct moral intuitions about cases does not extend to far-fetched cases. I consider these lines of argument in Sections 5 and 6 respectively. Before investigating these lines of argument, however, I need to clear the ground a bit, first (Section 2) by stating what is common ground in the debate between conceivabilists and realists, next (Section 3) by discussing how the term ‘outlandish’ should be understood, and finally (Section 4) by showing the different ways in which outlandish cases purportedly can be used to refute otherwise valid moral principles.

2. Common Ground

The dispute between conceivabilists and realists is in many ways a family quarrel. Both parties agree on the role which intuitions about imaginary cases can play in moral inquiry. These intuitions allow us to test the validity of moral principles in the following manner: if principle P says that actions of type X has deontic status D whenever conditions x, y and z obtain, and we can present a case where x, y and z obtain, but where we intuitively think that the action has a different deontic status, this intuition counts as evidence against P being valid. Inversely, P is supported by intuitions stating that the action has the deontic status which the principle states that it has.⁸ While some theorists use intuitions only to *reject* principles, others also appeal to intuitions as support when establishing valid moral principles and building moral theories consisting of such prin-

ciples. These theorists typically adopt some version of the method of reflective equilibrium introduced by John Rawls, although they do not necessarily follow Rawls' strictures about the kind of cases which can be used.⁹

Both positions in this debate thus reject a wholesale scepticism against the use of intuitions, which has become increasingly common in recent years, in part as a result of empirical research on intuitions and the processes by which we come to have them.¹⁰ While confronting this general scepticism about intuitions is an important project, it should not make us overlook the independent need for a sound *methodology* of intuitions, which would allow us to know which intuitions we can appeal to and how we can use them, if we are to use intuitions at all. Indeed, there is a practical need for such a methodology in a way that there is not for meeting the sceptical challenges against intuitions, since philosophers will in all probability continue to appeal to moral intuitions, either because they believe the sceptical challenge can be met, or because the use of intuitions appears inevitable. (Compare: scepticism about the external world does not obviate the practical need for learning how to use a map and compass to navigate.) I will therefore discuss *how* intuitions should be used, even though we lack (or so many would say) a cogent refutation of the sceptical claim that intuitions should not at all be used.¹¹

3. When is a Case 'Outlandish'?

Two paradigmatic cases will give the flavour of what I mean when I call a case 'outlandish'; they are both drawn from a recent article by Kasper Lippert-Rasmussen, where he argues against the principle of self-ownership by way of various cases (many of) which he quite explicitly admits are outlandish.¹² I will quote them at full length, as I will refer to them throughout the paper:

'Four-eyes'

[L]et us consider an eye redistribution scheme in which half the population is born with two pairs of eyes and the other half with no eyes. In sighted individuals, one pair of eyes is located normally and fulfills the usual function. The other pair is located inside the human body, say, in the shoulder. Although this latter pair would enable those who have them to see if they were surgically moved to the eye sockets, they play no role where they are. Indeed they cannot perform any visual or other bodily function without being moved. Suppose further that the body of a person born with two pairs of eyes will expel the spare pair when that person reaches twenty years of age. The pair can then easily be reabsorbed into the shoulder of its owner, or the owner can transfer his spare eyes to a blind person. Typically people will feel alienated from their spare pair of eyes in the same way that the main character in Sartre's *Nausea*, Antoine Roquentin, feels alienated from his hand when he perceives it as a crab lying on its back moving its legs and claws. As it happens the state implements a compulsory redistributive scheme in which whenever a person's body expels a nonoperative pair of eyes, these are appropriated and transplanted into a blind person, giving him normal sight. (p. 98)

‘Two hundred legs and arms’

Suppose, for instance, that people are born with huge bodies they can barely move, bodies with two hundred legs and arms. At any given moment, they can at best sense and control 1 percent of their bodies, although they can readily determine which percent that is. Since their bodies heal very easily, their ability to control their lives is promoted best if 99 percent of each body is removed in such a way that these abnormal individuals end up with what are, for us, normal human bodies. (p. 109)

What underlies the realist position is usually more a sense that cases like these are too ‘far out’ than an a priori definition of what counts as outlandish. Thus, instead of trying to find a definition of ‘outlandish’ which captures our sense of outlandishness, and then argue that all cases which are outlandish according to this definition should be excluded, I will proceed by examining the reasons we might have for wishing to avoid outlandish cases. In order to pursue this strategy, our general sense of which cases are outlandish suffices, and I will, for the sake of simplicity, talk about cases from ‘our world’ as distinguished from cases from ‘outlandish worlds’ to refer to this general sense. Once we have seen what reasons we have for excluding outlandish cases, we might discover that not *all* outlandish cases are necessarily inadmissible in moral reasoning; some outlandish cases might be immune to the arguments for exclusion. (And for this reason, a definition of outlandishness becomes superfluous.)

One general distinction needs to be made at the outset, however: A case is *superficially* outlandish if the unrealistic features involved are not necessary for the moral intuitions we have about the case, i.e. if they could be replaced by more realistic features playing an analogue role. For example, Robert Nozick’s example of self-defence against someone falling on top of you involves a disintegration gun;¹³ this science-fiction element is, however, not necessary for the case to work and could be replaced by a more realistic alternative at the cost of some imaginative work. What is essential to the case is the general structure of the self-defence killing, not its details. By contrast, a case is *essentially* outlandish if the outlandish features are necessary in order for us to have the moral intuitions we have. For example, in ‘Four-eyes’, the outlandish assumption is necessary in order for us to reject our standard intuition that it is wrong to forcibly take sighted people’s eyes in order to give them to the blind.¹⁴ In the following discussion, I will only be concerned with essentially outlandish cases. (If the arguments for excluding essentially outlandish cases also extend to superficially outlandish cases, this is unproblematic: we can always rephrase the latter so that they are no longer outlandish.¹⁵)

4. How can Principles be Invalidated by Outlandish Cases?

Before we consider whether we should allow intuitions based on outlandish cases to invalidate otherwise valid moral principles, we need to ask how it is possible that a principle can be valid for *this* world (in the sense that it explicates all of our relevant intuitions about cases taken from this world), and yet can be rejected on the basis of intuitions about cases from other worlds. I can see two ways of explaining this phenomenon. First, the feature X, singled out as morally salient by principle P, might in this world always co-exist with feature Y. However, it might actually be Y which is morally relevant, and we can only see this from outlandish examples where X and Y come apart.¹⁶

We can illustrate this by Lippert-Rasmussen's 'four eyes'-case. Since the forcible extraction of someone's eye in this world requires 'unwanted bodily incursions' as well as an infringement of the right to self-ownership, our intuition that forced eye redistribution is wrong can equally well be explained by a ban on unwanted bodily incursions as by a right to self-ownership.¹⁷

Next, X might be morally relevant only given the presence of certain enablers,¹⁸ which are present in our world, but not in all outlandish worlds (and, similarly, given the absence of disablers). By presenting an outlandish case, where (one or more of) the relevant enablers present in this world are lacking, we see that our moral principle P is at best imprecise, and that a fully precise description of P would include a description of all the relevant enablers.¹⁹ For example, the fact that, unlike in Lippert-Rasmussen's 'two hundred legs and arms'-case, we are naturally able to move our bodies around is not in itself the reason why we have a right to self-ownership of our body; nevertheless, without this fact, a right to self-ownership would not enhance autonomy, and hence (if we believe our right to self-ownership can be derived from a right to autonomy) we would not have a right to self-ownership.²⁰

While these explanations show how a principle valid for this world might fail to explicate intuitions about other worlds, they do not necessarily tell us how to proceed once our principle P has thus been invalidated. Indeed, when we have an intuition which is not explicated by P, we do not necessarily know, by that same intuition, *why* P is not valid in that case. Furthermore, different versions of the general intuition-based moral method will give different advice as to how to deal with recalcitrant cases. While we ideally want a principle covering all of our intuitions, most forms of the method of reflective equilibrium (RE) allow us to reject certain of our intuitions in order to save an otherwise valid moral principle. It is thus perfectly admissible to reject the outlandish intuition, in order to keep P, if it is not possible to find a new principle covering all of our intuitions (call it P*). It follows that when intuitions about outlandish cases are to be used to reject P, the critic of P working within the standard RE-approach must either find such a covering principle (P*), or argue that we should reject some of our intuitions about *this* world, in order to save another principle P**, which covers the outlandish intuitions, and (presumably) some, but not all of our intuitions about this world. Simply noting that P has a counter-intuitive consequence is not sufficient for rejecting P, on this approach.

By contrast, according to extreme variants of the intuition-based method,²¹ we may never abandon any of our intuitions: we must find principles which match all of them. Two differences from the traditional RE-method follow. First, in order to refute P, all we need is the counter-example; we do not need to find an alternative principle P*. Next, the search for a better principle can never take the form of P**, which is only valid if we reject some of our intuitions about this world. The critic of P who has a positive, as well as a negative aim, thus has two ways in which he can identify a valid principle to replace P (corresponding to the two explanations proposed above).

- a. Either by replacing the reference to one morally salient feature Y with a reference to another feature X, which explains our intuitions in all the proposed cases.
- b. or by specifying P, mentioning all the necessary enablers, without mention of which P would not be valid in all outlandish worlds. (In this case, P will need to be completed by a vast number of other principles, covering the cases where one or the other enablers of P are absent.)

5. The Scope of Moral Principles

I now consider the first line of defence of realism, which consists in saying that moral principles have a limited scope. The conceivabilist acts as if our moral principles should have universal scope, but given that the aim of morality is essentially practical and action-guiding, we only need principles which are valid for *this* world. Strictly speaking, we should thus always precede our proposed moral principles with a scope-specifier, such as ‘In the world in which we live’, but this is pragmatically unnecessary as long as it is implicitly understood that ethical principles are only meant for this world.²² This is, I believe, the most common approach among the critics of outlandish cases in ethics.²³

Building on the framework developed by G. A. Cohen,²⁴ I will discuss two ways of arguing for the limited scope of moral principles, and hence for the exclusion of outlandish intuitions.²⁵ These two lines of argument correspond to two distinct ways in which we can derive a principle of (relatively) narrow scope from a principle of wider scope.²⁶ On the one hand, we can derive principles of the form P ‘We should respect the right to self-ownership’ from wider principles such as P1 ‘We should respect the right to autonomy’, combined with facts such as F1 ‘Control of one’s life is necessary for autonomy’, F2 ‘Self-ownership is required for controlling one’s life if one can move one’s body around’ and F3 ‘We can move our bodies around’.²⁷ P is, in G. A. Cohen’s vocabulary, ‘fact-dependent’, since its validity depends on facts F1–F3. For the same reason, its scope is limited to the world in which these facts obtain: fact-dependence and scope-restriction go hand in hand.²⁸ P1, by contrast, also covers worlds where these factual claims are not true. Importantly, both P and P1 are *true* moral principles, expressing what we, following Jeff McMahan, can call ‘deep morality’.²⁹

There is a second type of principle which we can derive from wide-scope principles of deep morality, viz. what G. A. Cohen calls ‘rules of regulation’. Cohen describes the relation between rules of regulation and deep morality as follows:

A rule of regulation is ‘a device for having certain effects’, which we adopt or not, in the light of an evaluation of its likely effects, and, therefore, in the light of an understanding of the facts. And we evaluate those effects, and thereby decide which fact-bound principles to adopt, by reference to principles that are not devices for achieving effects but statements of our ultimate and fact-free convictions.³⁰

It is these ultimate moral convictions which constitute deep morality. Rules of regulation thus differ from the principles of ‘deep morality’ in two ways: First, while the moral principles of deep morality are principles we discover, the rules of regulation are principles which we try to inculcate in ourselves and others. Next, the success criterion is different: in our search for the principles of deep morality, we seek principles which express the moral truth, whereas in our search for the best rules of regulation, we seek the principles which will have the best consequences.³¹ The distinction between deep morality and rules of regulation provides us with two ways of justifying the limited scope of our moral principles: We can argue that since our moral principles are to be understood as rules of regulation, and are only supposed to regulate *this* world, we need not care about outlandish cases; or we can argue that even if our moral principles are principles of deep morality which purport to express the moral truth, and not just to

regulate behaviour, we are only interested in the moral truths relevant for our world. I now go on to examine these two forms of justification.

5.1. Principles as Rules of Regulation

R. M. Hare is one of the most prominent critics of outlandish cases. On Hare's view, the role of what he calls *prima facie* principles³² is not to express the truth about morality, but to have desirable consequences. *Prima facie* principles are thus akin to Cohen's 'rules of regulation'. Hare writes: 'To generalize: if we are criticizing *prima facie* principles, we have to look at the consequences of inculcating them in ourselves and others; and, in examining these consequences, we have to balance the size of the good and bad effects in cases which we consider against the probability or improbability of such cases occurring in our actual experience.'³³ From this view of moral principles, it follows clearly that no outlandish intuitions can invalidate them.³⁴ It is irrelevant, when assessing a principle P, whether P would give the wrong advice in a situation which would never occur in real life, since the criterion of validity of P is whether its inculcation will have good consequences *in real life*.

The problem with this approach is that it makes even this-worldly intuitions irrelevant for assessing P. For a two-level theorist like Hare (or Cohen), the *only* thing which matters when assessing rules of regulation is whether their consequences are good, and, as Cohen notes, this is a question best answered by social science (to predict the consequences) combined with a theory of value (to evaluate the consequences).³⁵ It is true that often a rule of regulation which is at odds with deep morality will tend to have undesirable consequences, and for this reason a match with our intuitions (assuming that these reflect deep morality) is an indication that the rule of regulation will have undesirable consequences. But it is nothing more than an indication, and in various circumstances, well studied by consequentialists over the years, it would be sub-optimal in terms of consequences to have rules of regulation matching deep morality. It follows that the Hare-style argument against outlandish intuitions is neither here nor there when we ask *which* cases we can use when we use intuitions to support moral principles. This question only makes sense given the general methodology sketched in Section 2, and Hare does not share this methodology.

5.2. Principles as Fact-Dependent Principles of Deep Morality

Unlike Hare's approach, the second line of argument is only concerned with principles which express the truths of deep morality. But, the argument goes, we do not need to identify all the truths of deep morality, since our concern with ethics is fundamentally of a practical nature. David Miller, for example, writes: 'I start from the assumption that political philosophy is a branch of practical reason — it is thought whose final aim is to guide action, as opposed to having a merely speculative purpose.'³⁶ Given the practical purpose of ethics, we need only identify the true principles which cover *this* world, even if these principles are derived from principles of a wider scope which cover other worlds as well. This argument does not deny that there are also true claims to be made about the wider-scope principles or that someone who is interested in truth for its own sake might be interested in these claims; it simply says that we do not need to worry about these claims if we see ethics as essentially practical. This distinction between the wide-scope

principles we seek for the sake of truth and the narrow-scope principles which we seek (also) for practical reasons is made several times by Cohen (who himself sees ethics as part of the search for truth for its own sake and not as essentially practical). For example, he writes:

It is true that we don't need to know what [the ultimate warrant of our norms] is, for practical purposes, but (in my view) philosophy's role is not to tell us what we need to know (in that sense), but what we want or ought (for non-practical reasons) to know. Not all will agree. Some might attack my philosophical presuppositions as 'philosophist'.³⁷

This distinction opens for saying that there are simply two separate projects running parallel in ethics and political philosophy: one project seeks to identify *all* true moral principles, simply for the sake of truth, while another project only seeks to identify the moral principles we need for guidance in the world we live in.³⁸ The realist's defence against outlandish intuitions will then simply consist in saying that while these intuitions might be relevant for a different philosophical project — that of identifying the *whole* moral truth — they are irrelevant for the realist's project, which is limited to what is needed for practical purposes. This approach would allow one way of specifying which cases should be rejected as too outlandish: a case should be rejected as outlandish if it does not represent a practical ethical challenge for us. Possible future cases are for that reason not outlandish, at least if they take place in the foreseeable future which we can act upon.

This approach is, however, somewhat defeatist: It abandons the implicit supposition in much philosophical work that philosophers discussing the same topic (such as just distribution of resources, the right to self-defence, etc.) are all engaged in the same enterprise, and that when a new article is published on this topic, it is relevant for all those who work on this topic. Rather, this approach suggests that all publications in ethics should somehow be marked as either belonging to the one or the other category, so that those only concerned with principles for this world could know that they did not have to read the other. While this defeatist picture might have some descriptive adequacy, it is, I believe, unfortunate as a normative proposal for philosophical methodology.

There is also a more substantial reason why this line of defence should be abandoned, even if one starts within the practical-oriented approach and is only concerned with identifying the principles we need in our world. I take it that the proponents of this scope-restricted approach to ethics are interested in identifying the *true* moral principles for this world: the pragmatic attitude which they have towards the scope of ethical principles does not extend to the content of these principles, even if certain false principles would do just as good a job as the true principles of guiding our actions so that we act morally. If this is our goal, outlandish cases might be necessary for identifying the true principles, for the reason described above (Section 4): if two features X and Y always co-exist in this world, we need outlandish examples which separate X and Y in order to see which feature is morally relevant.³⁹ The realist might reply by saying that she is perfectly happy to accept the false principle which states that X, rather than Y, is the relevant right-making factor; after all, it makes no practical difference, in this world, as long as both principles explicate all of our intuitions about this world.

There are three problems with the realist's reply. One is that many moral theories require us not only to do the right thing, but also to do the right thing *for the right reason*.

If we act on a principle according to which X is the right-making factor, when in fact Y is the right-making factor, we do not act on the right reason.⁴⁰ Of course it is possible to reject the claim that we should act for the right reason; but this is a substantial claim in moral theory, and I take it to be a desideratum of methodological positions in ethics that they do not exclude any substantial moral positions.⁴¹ Next, if one does not care about the truth of one's principles, but only about the fact that they make us perform the right action, it would be better to not search for principles of deep morality at all, but rather for principles of regulation. Indeed, it might be the case that there is a principle of regulation which, while it does not even explicate all of our non-outlandish intuitions, does a better job than any principle which does explicate these intuitions at making us act rightly. In that case, if one does not care about truth, why not adopt this principle of regulation? However, once we have adopted this logic for choosing principles, intuitions are no longer relevant as data, as shown above. Finally, most intuition-based methods of theory-building do not content themselves with simply finding principles which explicate all of our intuitions; in order for us to accept as valid the principles which explicate our intuitions, we must also examine them for their own sake, to see if they are plausible.⁴² Given this methodology, it is clear that it matters that our principles identify the right- and wrong-making factors in a correct manner.

To conclude: if we do wish to reject some cases as too outlandish, it cannot be because such cases are of no practical interest.

6. The Limits of our Intuitive Competence

The second main line of defence of realism consists in saying that we cannot trust our intuitions about outlandish cases. So even if we should need to find moral principles which cover certain outlandish cases, the right way to find these principles is not by considering our intuitions about these cases. This line of defence builds on two ideas: First, we have a basic intuitive moral competence, which we apply when we consider various cases. Next, this competence is not of universal scope: for certain hard cases, we fail to intuitively see what is right. Indeed, this is why we need the moral principles developed through intuitions about easier cases.

If we grant that our intuitive competence does not cover all cases, what are the grounds for saying that (some) outlandish cases are among the ones not covered? I will consider two lines of argument. The first line of argument states that our moral competence — the capacity we have for, once we have a full understanding of the facts of a case, applying our moral intuition to see what should be done — is limited to realistic cases. An explanation of this claim could be found in the origins of our moral competence. Here, different stories might be told. For example, our moral competence might be biologically grounded, as the result of evolutionary pressures.⁴³ Outlandish cases were never encountered in our evolutionary past, and hence we have not evolved the relevant intuitions. This story has some obvious problems. One is that it would exclude many present-day cases as well, since they were not present in the environment of evolutionary adaptedness.⁴⁴ But more importantly, the supposition that the intuitions given us by evolution because they serve our genetic survival just happen to also be the morally correct intuitions has come under serious criticism.⁴⁵ A similar story could be told about the social origin of our intuitive competence, and a similar debunking approach is available for this story too.⁴⁶

These problems point to a general worry for this first line of defence. All attempts at showing why our moral competence does not cover outlandish cases must say something substantial about the nature of our moral competence — how it works, why we normally can trust it, etc. However, as mentioned in Section 2, there is currently a vogue of sceptical argument questioning the claim that we should have faith in our intuitions, and there is no commonly accepted account of how moral intuitions work which will uncontroversially serve to meet this scepticism. Of course, such an account might come. But in the meantime (and given how hard it is, in general, to reply to sceptical worries, the wait might be long), philosophers who are primarily interested in methodology, not in meeting the sceptical challenge, would do well to adopt some form of ‘method of avoidance’ in this regard. We would not use the intuition-based method if we did not have a general faith in our moral competence, but as long as we lack an account of *how* our competence works, we just have to take its working as a given fact. The price of this method of avoidance, however, is that we cannot appeal to any story about how our moral competence works in order to exclude outlandish cases.⁴⁷

It might be objected that this strategy is too pessimistic, and that while we may not yet have a full account of how our moral intuitions work, we might have, or be able to develop, at least a rudimentary sketch of what such an account could look like, and such a sketch might suffice to at least partially explain why we do not have valid intuitions about outlandish cases.⁴⁸ In reply, I would say that I do not think there is today a consensus concerning even what the beginnings of an account of our moral intuitions should look like, and so it is unclear *which* account-in-progress one should look to for excluding outlandish intuitions. I do, however, see two ways of pursuing this strategy, even at the present unfinished state of theories of moral intuition. First, one might show that it follows from all the plausible sketches of a theory of intuition that outlandish cases should be excluded. Next, one might identify some constraints which any theory of moral intuitions has to satisfy, and show that it follows from these constraints that outlandish cases should be excluded. In either case, one could use fledgling theories of moral intuitions to rule out outlandish cases without having to decide which of these fledgling theories of intuitions has the most plausibility.

While I admit that these are possible strategies, I will here rather pursue a second and, in my view, more promising line of argument, which states that in order for us to apply our moral competence to a given case, we need to fully understand the case, so that we can identify all the morally relevant factors involved.⁴⁹ And, this is the proposed claim, in essentially outlandish cases, we fail to fully understand the situation described. I will discuss three reasons in defence of this claim.

First, as Derek Parfit has suggested, we might sometimes be unable to imagine sufficiently what the outlandish element *itself* is like.⁵⁰ As an example, Parfit suggests that it is impossible for us to imagine what Nozick’s utility monster is like. The two-hundred arms people might be a similar case: I cannot imagine what it is like to be them sufficiently well to know whether for them it is very important to decide which of their hundred pair of arms to move at a given time. While such cases are important, I will move on to a second type of argument, according to which even if we might very well imagine the outlandish element itself, we might fail to imagine how that element changes the totality of the case described.

This second type of argument is one which has been developed in detail both by Kathleen Wilkes and by Sören Häggqvist.⁵¹ Its point of departure is the observation that

thought experiments in general are underdescribed, and that outlandish cases are radically underdescribed — only the features differing from our world are mentioned, but not how these features change other features which might be relevant.⁵² Take ‘four eyes’: we do not know how this strange situation affects the social structure and the psychological make-up of the four-eyed people. Lippert-Rasmussen (unusually, for such cases) does mention one psychological effect of this situation: people feel alienated from their extra eyes. But this effect is compatible with many other things being the case. For example, the extra pair of eyes might become invested with various rich social meanings, as deep markers of social identity; or the psychological struggle to overcome one’s eye-alienation might become an important rite of passage among the four-eyed. Despite these possibilities, it might still be acceptable to take the eyes forcibly, but that is not the point here. The point is that, in order to have intuitions about whether this is the case, I would need to know all these details. The case about which we have intuitions is never just the vignette presented to us: it consists of the details given in this vignette plus all the other (relevant) facts about the world, which we somehow have to infer.⁵³

This is so just as much in the case of this-worldly imaginary cases: the cases would rarely make sense if we could not bring to them our background assumptions about what the world is like. For example, the trolley problem would be pointless if we were not allowed to fill in all the blanks, such as ‘The world and all its inhabitants will not suddenly cease to exist once the trolley has run its course.’ For this reason, any methodology for the use of imaginary cases must tell us how we should fill in the facts that are not explicitly stated, but without which the case would nevertheless be meaningless. A fruitful suggestion, made by Jonathan Ichikawa and Benjamin Jarvis, is that we should use the notion of ‘truth in fiction’ to explain how we interpret thought experiments.⁵⁴ Not all of the things which are true in a work of fiction are explicitly mentioned in the text presenting that fiction, and so we need an account of what other things are true in the fiction. While this idea can be cashed out in various ways, it will do for our purposes to follow David Lewis in saying that what is true in the fiction is constituted by, in addition to the text itself, the background beliefs of the community in which the text was written.⁵⁵ If we add that we typically are part of the community in which the ethical thought experiments that we read are written, we can say that the standard methodology for filling out the necessary blanks in a thought experiment is just to add our background beliefs about the world. However, this method, which works well for cases from our world, fails for outlandish cases, since we do not have background beliefs about outlandish worlds, and since our background beliefs about *this* world are not necessarily true of the outlandish world. Indeed, it is implausible that the radically different facts stipulated by outlandish cases should create no other morally relevant factual differences between our world and the outlandish world. As long as we have no other way of filling out the details of the case, we should not trust our capacity to imagine in full the outlandish cases.⁵⁶

Might not the case-designer simply stipulate that all our background beliefs remain true? One reply is that even if he did, we might be unable to refrain from more or less unconsciously filling in these differences (as, for example, we do when we read science-fiction), and thus have intuitions about a different case than the one proposed to us.⁵⁷ A stronger reply, stressed both by Wilkes and by Häggqvist, is that it is sometimes impossible that all of our background beliefs remain true. For example, in order for the people in ‘two-hundred legs and arms’ to have lives otherwise just like ours, they would have to

have a psychological make-up very unlike ours, since anyone with our psychological make-up would lead different lives from ours in such conditions. So one of two things would necessarily change in this case: either people's lives or their general psychological make-up.

As a solution to this second challenge for imagining outlandish cases, Sören Häggqvist has proposed a criterion for successful thought experiments.⁵⁸ In brief, Häggqvist argues that we must find the closest possible world where the supposition made in the thought experiment is true. Since the description of this world cannot contain any inconsistencies, it must include all the further changes that the outlandish example commits us to accepting. Häggqvist calls this process 'accommodating' the case. Given this requirement, attempts at refuting a theory by way of thought experiments can fail in three ways: because there is no way to accommodate the thought experiment, because once we have accommodated the thought experiment, the intuition we have is now in accordance with our theory, or because the world thus described falls outside of the scope of the theory the experiment is meant to test.⁵⁹ Häggqvist notes that the latter form of failure 'sometimes seems to be the case for instance with moral principles',⁶⁰ thus suggesting a defence of realism of the type discussed in Section 5. Häggqvist's proposal is relevant for the debate between conceivabilism and realism in several ways. First, it provides a useful reminder that *if* we are to use outlandish cases, we must proceed with care, as it is easy to overlook a failed accommodation.⁶¹ Next, it explains how the use of outlandish cases can be allowed at all, thus providing some ammunition for the conceivabilist. But Häggqvist also provides arguments for the realist who wants to exclude outlandish cases, since each of the three ways in which a thought experiment can fail might be appealed to by the realist in order to reject outlandish cases. Both the conceivabilist and the realist can therefore apparently profit from Häggqvist's criterion.

Nevertheless, while Häggqvist's proposal is fruitful in many ways, I believe that there are limits to its use in the case of ethics, for three reasons. First, as I argued in Section 5, we cannot appeal to the limited scope of moral principles in order to exclude the use of outlandish intuitions — the third form of failure of thought experiments can thus not be used to defend realism in ethics. Next, two problems face the conceivabilist who wants to draw on Häggqvist to defend outlandish cases. One is that once we have accommodated an outlandish case, a large number of factors will differ between that case and the actual world. This means that we cannot use the technique of equalizing cases in order to identify precisely which factor accounts for the difference in intuition between the this-worldly case and the outlandish case, something which is essential if we want to use intuitions for theory-building in ethics. The second problem is that Häggqvist's solution does not allow us to deal with the two other difficulties involved in imagining outlandish cases — that the outlandish element in itself might be impossible to imagine, and that the moral relevance of various factors might change with the presence of an outlandish element. It is to this last problem that I now turn.

My third argument against the possibility of fully understanding outlandish cases is that even if no other feature of the world changes, the moral import of these features might change: The effect of removing factor X is not just that X is absent, but that Y might matter in a different way than with X present, as Jonathan Dancy has argued at length.⁶² I want to suggest that we cannot trust our ability to perceive these changes in the moral relevance of Y just from a short description of the case, a description where Y might not even be mentioned. My point of departure is the idea that our moral

competence is first and foremost used when we encounter morally challenging situations in the real world. What we do when we consider imaginary cases, is to apply this same competence ‘offline’ to the case presented to us.⁶³ And since we are used to the ways in which the removal or addition of one factor X will *in this world* influence the relevance of other factors, we know which factors to take into consideration when we assess an imaginary case which could occur in this world. If, for example, we are presented with a this-worldly case where one of a seeing person’s eyes was forcibly taken away to help a blind person see, and where nothing else than this is said, we immediately (and often unconsciously) take into consideration a certain number of facts about the world which we from experience know matter in such cases: the fact that eyes allow us to see, the sense of loss typically resulting from the violent loss of a body part, the danger and pain associated with the operation of eye removal, the (small) problems due to lack of stereoscopic vision, the value associated with gaining sight, etc. Only by taking these facts into consideration can we have a valid intuition about the acceptability of the action described. When, by contrast, we are presented with the ‘four eyes’-case, we do not know which further facts we should take into consideration (even supposing all facts are as in this world), for we do not know what would be relevant in such a case. And as we do not know which features of the case to take into account, we cannot properly apply our intuitive competence.

At this point, the conceivabilist might reply that *he* at least does imagine outlandish cases well and has intuitions about them.⁶⁴ Indeed, even their realist opponents probably have intuitions about the cases, or they would not be troubled by attempts to refute P by way of outlandish intuitions. However, the fact that one has intuitions about a case C is no proof that one has imagined C sufficiently well to have valid intuitions. At the same time, I readily admit to having no proof that these cases cannot be imagined. I nevertheless hope that the difficulty of imagining such cases should provide some pause to conceivabilists. Further pursuit of these lines of inquiry will be necessary in order to convince conceivabilists to abandon their ways.

Finally, let me note that the conceivabilist who accepts my argument, but nevertheless wants to keep using outlandish cases, might try to provide a richer description of these cases by drawing on well-known fictional universes, such as those of *The Lord of the Rings*, or *Star Trek*. I do believe that such an approach would help us imagine the cases: not only would the relevant blanks be filled in by the rich fictional context, but the reader’s familiarity with these fictions, and her habit of thinking about them in moral terms (‘Boromir was weak-willed, but not fundamentally evil’) allows her to know which kinds of fact are morally relevant in that strange context. And a richer description can help imagine the outlandish elements themselves as well.⁶⁵ However, this approach shares one of the problems of Häggqvist’s solution, which is that it would not allow us to use thought experiments as they typically are in moral philosophy, using the ‘technique of equalizing cases’ to identify exactly which factor is morally relevant. If we, by way of intuitions, learn that one has no right self-ownership in the Lord of the Rings-universe, there are far too many factors which could potentially explain this fact for the intuition to be very useful.

7. Conclusion: A New Road Towards Scepticism?

An advantage with using the argument from intuitive competence rather than the argument from the restricted scope of principles to exclude outlandish cases, is that we

get a better account of the need for moral principles. Implicit in the scope-restriction approach was the idea that the domain of cases for which we need moral principles and the domain of cases from which we can draw intuitions to test principles are identical. This assumption does not give principles any role in discovering what we should do (at best in explaining why we should do what we should do), something which seems at odds with actual moral reasoning. By rejecting this assumption, we can say that we use principles to know what to do in cases which we fail to have trustworthy intuitions about. This would be of little practical use if only otherworldly cases were in that category. But the proposed account suggests that it is not so: many possible future cases, and perhaps even some cases which we are presently confronted with, will be so different from all we know that we fail to imagine them fully: while possible, they are still in one sense outlandish.⁶⁶ This is perhaps the case for some of the possibilities created by modern biotechnology, such as human reproductive cloning. On the competence-based account of outlandishness, the cases which we are the most apt at applying our moral competence to are those which resemble the world in which we have lived, and it is by drawing on these cases that we develop principles for harder cases.

One worry remains, however: The conjunction of the argument from competence and the rejection of the argument based on scope-restriction creates a new source of moral scepticism.⁶⁷ I have argued that we do sometimes need to consider outlandish cases in order to arrive at the true moral principles which we need for *this* world. But if we are, as I have suggested, unable to imagine these outlandish cases, we might never be able to identify the true moral principles we need. So even if we adopt a ‘method of avoidance’, bracketing general scepticism for the purpose of developing a sound ethical methodology, a new form of scepticism reappears as the result of this methodological discussion. How this new scepticism can be dealt with is an issue for further investigation.

Jakob Elster, Department of Philosophy, Classics, History of Art and Ideas/ IFIKK, P.O. Box 1020 Blindern, 0315 Oslo, Norway. jakob.elster@ifikk.uio.no

Acknowledgements

Earlier versions of this article were presented at the conference ‘Political Ethics — its Nature and Methods’, Uppsala, 16–18 June, 2010, and at the Society for Applied Philosophy Annual Conference 2010, 2–4 July, St Anne’s College, Oxford. I am grateful to the participants at these events for their comments, and in particular to Kim Angell, Lene Bomann-Larsen, Paula Casal, Thomas Christiano, David Copp, Beate Elvebakk, Eli Feiring, Chloe Fitzgerald, Hans-Johann Glock, Robert Huseby, Cathrine Holst, Raino Malnes, Shaun Nichols, Kasper Lippert-Rasmussen, Folke Tersmann, Frej Klem Thomsen and Wybo Wiersma. I am also grateful to an anonymous referee and an associate editor for the *Journal of Applied Philosophy* for helpful suggestions.

NOTES

- 1 I am only concerned with the use of outlandish cases in ethics, as described in Sections 2–4 below, and therefore I do not take a stand on whether my arguments in this article can be extended to the use of outlandish cases in other areas of philosophy. The use of outlandish cases in philosophical thought experi-

- ments in general has in recent years been treated in depth in several works, in particular Kathleen V. Wilkes, *Real People: Personal Identity Without Thought Experiments* (Oxford: Clarendon Press, 1988), chapter 1; Roy Sorensen, *Thought Experiments* (Oxford: Oxford University Press, 1992), in particular pp. 277–285; Sören Häggqvist, *Thought Experiments in Philosophy* (Stockholm, Almqvist & Wiksell International, 1996); and Tamar Szabó Gendler, *Thought Experiment: On the Powers and Limits of Imaginary Cases* (New York: Garland Publishing, 2000). While these discussions are of course largely relevant for the use of outlandish cases in ethics, I am not able to engage in detail with them within the scope of this article. Particularly relevant for the argument in this article, however, are Wilkes and Häggqvist, who have developed one of the lines of criticism against outlandish cases discussed in Section 6; I discuss Häggqvist's solution to that criticism below.
- 2 Among the moral philosophers I take to be conceivabilists (as identified mainly by the kinds of case they use) are Kasper Lippert-Rasmussen, Frances Kamm and Derek Parfit. Parfit explicitly defends conceivabilism in Derek Parfit, *Reasons and Persons* (Oxford: Clarendon Press, 1987), e.g. p. 200 and pp. 388–9. We might also place G. A. Cohen in the conceivabilist camp. While Cohen does not primarily discuss the use of intuitions about outlandish cases to test principles, he provides support for conceivabilism by defending the view that we need to find moral principles which also cover outlandish cases, and his arguments have provided the framework for several discussions of outlandish cases in ethics in recent years (including in this article). Cf. G. A. Cohen, *Rescuing Justice and Equality* (Cambridge, MA: Harvard University Press, 2008), chapter 6.
 - 3 In the realist camp are e.g. Thomas Pogge, R. M. Hare, and David Miller. Cf. Thomas Pogge, 'The effects of prevalent moral conceptions', *Social Research* 57,3 (1990): 649–63, especially p. 660; Thomas Pogge, 'Cohen to the rescue!', *Ratio* XXI,4 (2008): 454–75, e.g. p. 475; R. M. Hare, *Moral Thinking: Its Levels, Method, and Point* (Oxford: Clarendon Press, 1981), e.g. p. 47; David Miller, 'Political philosophy for earthlings', in D. Leopold & M. Stears (eds) *Political Theory — Methods and Approaches* (Oxford: Oxford University Press, 2008), pp. 29–48. Pogge, for example, writes: 'I regard it as a key weakness of certain analytic moral philosophers that they do not distinguish between these two dimensions: A hypothetical case that can arise in *our* life context may be very important, even if it is quite far-fetched, while a hypothetical case that can arise only in very remote life contexts is likely to be irrelevant to *our* moral reflection' (Pogge 1990 op.cit., p. 660.)
 - 4 John Rawls belongs to this camp when he requires that our 'considered judgments', which are the input to a process of reflective equilibrium, be about a case 'in which there is an actual conflict of interest. Thus, all judgments of hypothetical cases are excluded' (John Rawls, 'Outline of a decision procedure for ethics', *The Philosophical Review*, 60, 2 (1951): 177–197 at p. 182).
 - 5 F. M. Kamm, *Intricate Ethics — Rights, Responsibilities, and Permissible Harm* (Oxford: Oxford University Press, 2007), pp. 427–8. Kamm describes the technique as follows: '[W]hat some philosophers who argue for a moral distinction between harming and not-aiding have done is construct cases that are alike in all respects other than that one is a case of harming and the other a case of not-aiding. It is only if all other factors are equal that we can be sure that people's responses to cases speak to the question of whether harming and non-aiding per se make a moral difference. It can be very difficult to construct equalized cases, and often philosophers think that they succeed when they do not' (p. 428).
 - 6 Kamm makes this point in F. M. Kamm, *Morality, Mortality, Volume I: Death and Whom to Save from It* (Oxford: Oxford University Press, 1993), p. 7.
 - 7 These two types of argument are briefly distinguished in Parfit op. cit., pp. 388–9. A number of other arguments against outlandish thought experiments (which he labels 'far out') are discussed, and mostly rejected, by Roy Sorensen: Sorensen op. cit., chapter 10, especially pp. 277–285. I briefly discuss the relation between my argument and what I take to be Sorensen's main argument against outlandish cases in endnote 43. (It would go beyond the scope of this article to more substantially engage with the details of Sorensen's rich account of most aspects of thought experiments.)
 - 8 Although I have illustrated this intuition-based testing with principles concerning an action's deontic status, the method is equally used for other moral principles, such as principles stating what rights people have.
 - 9 As noted in endnote 4, Rawls only accepted judgements about actual cases as input. Although the method of reflective equilibrium originated with Rawls, it has by now become common philosophical property, and exists in many variations (even within Rawls' work), depending on the kind of cases which are used as input, on the kinds of principles one relates these cases to, on one's willingness to abandon recalcitrant intuitions, etc. (Kamm, for instance, describes her approach as a variant of the method of reflective equilibrium. Kamm 2007, op. cit., p. 5.)

- 10 For example, Peter Singer, 'Ethics and intuitions', *The Journal of Ethics* 9 (2005): 331–52, Walter Sinnott-Armstrong, 'Framing moral intuitions', in W. Sinnott-Armstrong (ed.) *Moral Psychology Volume 2: The Cognitive Science of Morality: Intuition and Diversity* (Cambridge, MA: MIT Press, 2008), pp. 47–76.
- 11 Cf. Sorensen op. cit., pp. 288–9 for a related discussion (and a related analogy).
- 12 Kasper Lippert-Rasmussen, 'Against self-ownership: There are no fact-insensitive ownership rights over one's body', *Philosophy and Public Affairs* 36, 1 (2008): 86–118. This article is a cornucopia of outlandish cases, and indeed the author's explicit strategy depends on the use of such cases.
- 13 Robert Nozick, *Anarchy, State, and Utopia* (Oxford: Blackwell Publishing, 1974), p. 34.
- 14 Admittedly, Lippert-Rasmussen suggests the contrary (Lippert-Rasmussen op. cit., p. 99).
- 15 Cf. Sorensen op. cit., p. 280.
- 16 Cf. Lippert-Rasmussen op. cit., p. 94; Sorensen op. cit., p. 91.
- 17 Cf. Lippert-Rasmussen op. cit., pp. 97–98.
- 18 To use Jonathan Dancy's term. Enablers, or 'enabling conditions', are features of a situation 'which do not favour [one action or another], but whose presence or absence is required for others to do the favouring job'. Jonathan Dancy, *Ethics Without Principles* (Oxford: Oxford University Press, 2004), p. 52. In terms of G. A. Cohen's framework, which I present in Section 5, saying that a principle is only valid given the presence of an enabler amounts to saying that the principle is fact-dependent, the fact upon which it depends being the presence of the enabler. (See also the discussion of such principles in Miller op. cit.)
- 19 On the general idea of specifying principles by incorporating enablers, cf. Brad Hooker, 'Moral particularism: Wrong and bad' in B. Hooker & M. Little (eds) *Moral Particularism* (Oxford: Oxford University Press 2000), pp. 1–22: 7–11. Incorporating enablers into our principles corresponds to what Pogge (in his discussion of Cohen) describes as the internalization of fact-sensitivity, whereby we identify principles which are valid for all worlds but which specify in detail the conditions under which they become relevant (Pogge 2008 op. cit., pp. 459–460).
- 20 Lippert-Rasmussen op. cit., p. 109. See also p. 90 for this general argumentative strategy.
- 21 Such as perhaps that of Frances Kamm, cf. Kamm 2007 op. cit., p. 5.
- 22 Cf. Miller op. cit. p. 38. For a parallel point, see Hooker op. cit., p. 14.
- 23 For examples, see the references in endnote 3. Pogge 2008 op. cit. and Miller op. cit. both describe scope-restriction in terms of G. A. Cohen's framework, in similar ways to what I do here, and criticize Cohen's defence of wide-scope principles.
- 24 Cohen, op. cit. especially chapter 6. While Cohen himself is a prominent defender of wide-scope principles, his framework can be used to defend narrow principles as well.
- 25 Thomas Pogge has proposed what I take to be a third approach to scope-restriction (Pogge 1990 op. cit., pp. 659–663; Pogge 2008 op. cit.). The arguments I consider in the main text do not deny that there are valid principles which cover outlandish worlds, but simply say that we do not need to care about these worlds, only about our world. However, Pogge, as I read him, defends realism by (based on a pragmatist understanding of ethics) denying that there are any valid principles (or intuitions) about outlandish cases, and that this is the reason why outlandish cases are irrelevant. I do not find this line of defence promising, but a full assessment of it would go far beyond the scope of this paper. Pogge might, however, also be read as presenting a form of the argument I discuss in Section 5.2. (Some of Pogge's remarks — e.g. Pogge 2008 op. cit. p. 467 — might also be taken to indicate the kind of view I discuss in section 6.)
- 26 Cf. Cohen op. cit., chapter 6.
- 27 Cf. Lippert-Rasmussen op. cit., pp. 108–109.
- 28 Cf. Pogge's reading of Cohen in Pogge 2008 op. cit., p. 463.
- 29 Jeff McMahan, 'The ethics of killing in war', *Ethics* 114 (2004): 693–733. Cohen (somewhat misleadingly) calls the principles of deep morality 'fundamental principles'.
- 30 Cohen op. cit., p. 265.
- 31 Cohen op. cit., p. 263–268.
- 32 Hare treats intuitions about cases and *prima facie* principles as entities of the same kind; presumably case-intuitions are more specific versions of the more general *prima facie* principles, and thus purportedly allow us to test the latter.
- 33 Hare op. cit., p. 48.
- 34 As Hare and Cohen both make clear. Cf. Hare op. cit., p. 47; Cohen op. cit., pp. 264–266.
- 35 Cf. Cohen op. cit., p. 266.
- 36 Miller op. cit., p. 44. Note that, as I read him, Miller does not primarily base his rejection of the need for principles covering outlandish worlds, worlds where basic facts about human nature do not obtain, on this

- view of political philosophy as essentially practical, but on a rejection of Cohen's basic argument (Miller op. cit., pp. 32–38). Rather, the view of ethics as essentially practical is used to discuss a *further* restriction of the scope of political principles in terms of factual assumptions which are true of some societies, but not of others. However, the fundamentally practical nature of ethics can also be used to reject principles covering outlandish worlds, and it is this use that I discuss in the main text. (This use might also be implicit in Miller's article, e.g. on p. 45.) As I note in endnote 25, Pogge might also be read as supporting a position of the kind I discuss here. (Cf. e.g. Pogge 1990 op. cit., p. 660.)
- 37 Cohen op. cit., p. 247. For this distinction, cf. also Miller op. cit., p. 45; Sorensen op. cit., p. 279.
- 38 For a similar point, cf. Sorensen op. cit., pp. 278–280.
- 39 The realist here acts as if the only way in which outlandish examples could refute P was by showing that certain enabling conditions are required for P to be valid. But as shown in Section 4, that is not the case.
- 40 A strict interpretation of this requirement would say that we need a full understanding of the reason for which we act, in which case we can never act on a derived principle without knowing the principles from which it was derived (for the reasons given by Cohen). In that case, the practical-oriented defence of scope-restriction would never be valid.
- 41 This desideratum might be outweighed by other desiderata on methodological views, but it still provides a *pro tanto* argument against the realist's reply.
- 42 For example, Kamm 2007 op. cit., p. 5.
- 43 For a discussion of an evolutionary account of thought experiments in general, see Sorensen op. cit., especially chapters 3, 4 and 10. Sorensen uses this evolutionary account to suggest a 'biological baseline', which indicates which kinds of case work best in thought experiments (Sorensen op. cit., pp. 252–4), and he argues that thought experiments work better with cases which are familiar (pp. 65–67) and not too exotic (p. 280). The limits I propose for the use of outlandish cases fits well with the limits motivated by such an evolutionary account (although it would require further examination to see exactly to what degree these limits coincide). But since we cannot (as I argue in the text) use an evolutionary account of our moral competence to explain why we should not trust our intuitions about outlandish cases, another kind of argument is needed in the case of ethics. (Sorensen also discusses the difficulties of an evolutionary account of ethical thought experiments, p. 69–70.) One possible connection between Sorensen's evolutionary account of thought experiments and the argument I propose for moral intuitions is that the evolutionary account can help explain the difficulty of imagining outlandish cases well enough to apply our moral competence to them. Developing this possibility would go beyond the scope of this paper, however.
- 44 Singer op. cit.
- 45 For example, Singer op. cit.; Sharon Street, 'A Darwinian dilemma for realist theories of value' *Philosophical Studies* 127 (2006): 109–66.
- 46 Cf. Singer op. cit.
- 47 See, however, Gendler op. cit., chapter 4, for what might be interpreted as an approach of this kind.
- 48 I owe this point to one of the associate editors of the *Journal of Applied Philosophy*.
- 49 By 'full understanding', I mean a sufficiently good understanding of the case to identify its morally relevant factors, not necessarily an understanding of *all* the facts involved in the case. Since, especially in outlandish cases, it is not given in advance which factors are morally relevant (as argued below), we cannot in advance give precise criteria for what it takes to fully understand a case.
- 50 Parfit op. cit., pp. 388–9. Rachel Cooper also has a discussion of thought experiments failing because we cannot imagine the cases correctly. As an example, she criticizes 'Bernard Williams's thought experiment of people that split like amoebas' as a failed thought experiment, because we are unable to tell how this splitting takes place' (Rachel Cooper, 'Thought experiments', *Metaphilosophy* 36, 3 (2005): 328–347, at p. 342). This criticism might be read as an example of the inability to imagine what the outlandish element itself is like, but it might also be seen as an instance of the criticism that outlandish cases are underdescribed, as discussed in the next paragraph.
- 51 Wilkes op. cit., chapter 1; Häggqvist op. cit. in particular chapter 8. (For related discussions, see also Cooper op. cit.; Sorensen op. cit., e.g. pp. 258–9, 268–9.) I became aware of their writings only after I had developed a version of this argument independently.
- 52 This observation has been made about thought experiments in general by Jonathan Dancy and by Marilyn Friedman (Jonathan Dancy, 'The role of imaginary cases in ethics', *Pacific Philosophical Quarterly* 66 (1985): 141–153 and Marilyn Friedman, 'Care and context in moral reasoning' in E. F. Kittay & D. T. Meyers (eds) *Women and Moral Theory* (Totowa, NJ: Rowman & Littlefield, 1987), pp. 190–204) and about outlandish cases more specifically by Wilkes and by Häggqvist. For discussion, see also Sorensen op. cit., pp. 246–8.

- 53 Cf. Jonathan Ichikawa & Benjamin Jarvis, 'Thought-experiment intuitions and truth in fiction', *Philosophical Studies* 142, 2 (2007): 221–246.
- 54 Ichikawa & Jarvis op. cit.
- 55 David Lewis, 'Truth in fiction', in David Lewis, *Philosophical Papers Volume I* (New York: Oxford University Press, 1983).
- 56 This argument against the use of outlandish cases is thus not just a repetition of Dancy's and Friedman's criticism of thought experiments, a criticism which applies to realistic cases as well; rather the argument makes the point that while there is a possible reply to their criticism in the case of realistic cases, this reply does not work in the case of outlandish cases.
- 57 Cf. Sorensen op. cit., pp. 258–9.
- 58 Häggqvist op. cit., chapter 8.
- 59 Häggqvist op. cit., p. 152.
- 60 Häggqvist op. cit., p. 152.
- 61 Häggqvist op. cit., p. 158.
- 62 For example, Dancy 1985 op. cit. In this paragraph I draw on (but adapt) Dancy's account of learning from experience in Dancy 1985 op.cit. and Jonathan Dancy, 'Moral particularism', in E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy* (Spring 2009 edn.); available at <http://plato.stanford.edu/archives/spr2009/entries/moral-particularism/>. There is a danger in using particularistic arguments in this context, since they typically involve completely rejecting the use of cases to identify valid principles which we can then apply to new cases. A possible reply (to be developed further) is that in this world there is enough regularity in the relations between different moral factors for moral principles to be found (cf. e.g. Hooker op. cit., pp. 7–11; for a related point, cf. Sorensen op. cit., p. 273); while there might be such regularities in outlandish worlds as well, we do not have enough experience with these worlds to identify them.
- 63 See Cooper op. cit., Ichikawa & Jarvis op. cit. and Timothy Williamson, *The Philosophy of Philosophy* (Oxford: Blackwell 2007) for this general approach to thought experiments. Ichikawa and Jarvis use their account to propose a moderate defence of outlandish cases (pp. 237–239), which I nevertheless believe is compatible with my argument (as I believe it only defends superficially outlandish cases).
- 64 Cf. Kamm 2007, op. cit., p. 359.
- 65 This solution differs from Häggqvist's in that the description of the world of the thought experiment is not meant to be conservative in relation to our world; while this aspect might have its disadvantages, a possible advantage is that it might allow us to deal better with the first and third problems I have raised for imagining outlandish cases.
- 66 Cf. Sorensen's observation that 'familiarity is not the same as realism' (Sorensen op. cit. p. 66). Note that there might of course be other reasons than outlandishness which cause us to fail to have intuitions about a case.
- 67 I owe this point to Thomas Christiano.