

A New Approach to Resolving the Gamer's Dilemma: Applying Constructive Ecumenical Expressivism

Abstract In this final chapter, a new approach to understanding the gamer's dilemma is presented which seeks not to identify a single morally relevant factors which differentiates virtual murder from virtual paedophilia but, rather, aims to articulate the means by which (a) we acquire the moral attitude we do and (b) how this attitude is elevated to the status of a social norm. Constructive ecumenical expressivism is posited as the means of accounting for this and therefore explaining the intuition that is said to form the basis for the gamer's dilemma. The new approach's ability to resist objections raised against an appeal to social convention is also discussed, as is the form a normative ethic would take if one were to endorse constructive ecumenical expressivism.

Keywords Moral attitude · Moral realism · Anti-realism · Meta-ethical approach · *De re* and *de dicto* attitude

So far, I have presented a critical review of the various attempts proffered to resolve the gamer's dilemma. To be fair, a number of these attempts were rejected by the author at the time of their original publication (e.g. Luck) or have been challenged by other author's since (e.g. Luck & Ellerby's and Patridge's responses to Bartel's proposed resolution). Some recent attempts offer promise, however (e.g. Patridge and Ali),

although for reasons discussed fail to resolve the dilemma and lack the resources to be co-opted as a normative ethic.

In this final chapter, I present my own thoughts on how the gamer's dilemma might be resolved. My strategy is to consider the nature of the moral discrepancy on which the dilemma is built: the intuition that virtual murder is permissible and virtual paedophilia is not. What does it mean to morally approve or disapprove of something? My answer is that it means that one has either a positive or negative attitude towards the target of one's moral inquiry. Given this, how does one arrive at this attitude? Rather than simply falling back on the idea of an intuition, I intend to examine the relationship between the object of moral concern, our moral attitude towards it and the process by which we arrive at this attitude. Once this has been established, I intend to apply this approach to resolving the gamer's dilemma, with a view to broadening its application to virtual gaming content more generally.

6.1 NOW *THAT* IS IMMORAL, ISN'T IT?

There is no objective understanding of what an image might represent – it is in the mind of the viewer.

(Simpson 2009, p. 260).

Suppose I agree to play *Child Sexual Assault* with a friend. Shortly after commencing the game, my friend turns to me and pointing at something within gameplay says: "*That* is immoral". For my part, I disagree and tell him that *that* is not immoral. What has brought about this moral disagreement? Before responding, a passage from Patridge should prove informative:

[I]t seems that the gamer who cannot help but see *Child Sexual Assault* as a reflection of or extension of our moral reality should be unable to find this content enjoyable (or, again, should find it very difficult to do so). This is so because the object of her amusement is a different object altogether from the object as interpreted by the gamer who sees it only as a bit of harmless fun. The objects are interpreted quite differently and as a result the instances of amusement involved have different intentional objects. (Patridge 2013b, p. 32)

Here, Patridge is contrasting the two gamers (above) based on how amusing they find aspects of the gameplay. Given the discussion on

Patridge in [Chapter 5](#), the difference in their amusement is meant to reflect a difference in their moral attitude towards certain enactments: moral disapproval in the case of not being amused and, at the very least, failure to disapprove in the case of amusement.¹ Returning to my example, if we take the demonstrative pronoun ('that') to be referring to an instance of virtual paedophilia then, following Patridge, the moral disagreement my friend and I express is not symptomatic of a difference in moral attitude towards *actual* child rape (or not necessarily so); instead, it is that in the case of *this* token enactment, my friend and I have interpreted the intentional object differently. Again, following Patridge, the enactment is seen as either an extension of (in my friend's case), or a departure from (in my case), our moral reality.

This difference in interpretation is important and deserving of further attention. Patridge's comments are in fact in keeping with a meta-ethical position I have previously discussed called *constructive ecumenical expressivism* (CEE) (Young [2014](#), [2015b](#)). CEE offers a new way of thinking about the gamer's dilemma. It does this not by identifying what the morally relevant difference is between virtual murder and virtual paedophilia, in some moral-realist sense, but by explaining why a difference in moral *attitude* occurs. But more than this, CEE provides the means of understanding what is required for a normative ethic to be established; not only in regard to those virtual enactments involved in the gamer's dilemma, but all video game content.

6.2 CONSTRUCTIVE ECUMENICAL EXPRESSIVISM

CEE, like its forerunner – ecumenical expressivism (Ridge [2006](#)) – holds that, when uttered, the proposition "*That* is immoral" reveals two inter-related facts about the mental states of the subject. The first concerns an attitude. The second relates to a particular belief that is said to make anaphoric reference to this attitude.² Importantly, though, the subject does not hold (in this case) a negative attitude towards the particular act referred to by the demonstrative pronoun. Instead, in the case of "*That* is immoral", the subject disapproves of some property – call it P – and believes that *x* (which represents what 'that' refers to) realizes P. Thus, in declaring that murder is immoral, the subject holds a negative attitude towards P (some property yet to be described) and believes that an act of murder realizes P. Moreover, in stating that murder is immoral, the subject is not (should not be) declaring only that *this* token act of murder

is immoral but, rather, that the *type* of action, of which *this* particular act is a token, – in virtue of realizing property P – is immoral. Where a different type of act realizes the same property, then one should disapprove of any token act of this action type. If kidnapping, for example, realizes property P then, if one disapproves of murder (in virtue of property P), one should disapprove of kidnapping, also.

So what is property P? Property P can and does amount to different things for different people. S₁ may view P in terms of negative utility – for example, the realizing of more displeasure than pleasure (say, in the form of increased harm) – while S₂ may hold it to be a violation of God's law, or constitutive of a failure in one's duty to others. S₃, in turn, may characterize P as a vice rather than a virtue, and so on. Declaring that "*x* is immoral" – where *x* equates to murder – reveals the following:

(CEE) S disapproves of P and believes that *x* realizes P (thus making anaphoric reference to that of which S disapproves).

To state that *x* is immoral does not denote the truth of the proposition "*x* is immoral"; rather, it denotes the truth of (a) S's disapproval of P, and (b) S's belief that *x* realizes P.

In the context of *Child Sexual Assault*, where A (*qua* my friend) declares "*That* is immoral" and B (*qua* myself) denies this, both A and B are expressing a moral attitude. A disapproves of *that* (whatever 'that' happens to be) and B does not. More specifically, A disapproves of *p* and believes that *that* (whatever 'that' happens to be) realizes *p* (thus making anaphoric reference to that of which A disapproves). Before moving on, note that I have used the lowercase *p* in italics to denote the specific property of which A in particular disapproves. This should be contrasted with the uppercase P used earlier, which refers to some unspecified property of which a generalized subject disapproves.

Why does B not disapprove of *that*? It could be that both A and B disapprove of property P (where property P refers to the same thing). Nevertheless, it could also be that they are employing different interpretations of *x*: the event within the game. How each interprets the virtual event will shape whether they come to believe that *x* realizes P, and based on this belief whether they disapprove of *x*. Alternatively, it could be that both A and B consider *x* to be immoral. In other words, both agree that "*That* is immoral" in relation to the same virtual event. How might this be achieved?

In this case, moral agreement (or shared moral attitude) is achieved because both believe that x realizes some property (P) of which they each disapprove. It could be that P is the same for both players. However, it could also be that what counts as P is different in each case. Suppose A believes that x realizes p : where p equates to commending one to delight in that which is immoral (actual paedophilia, in this case). B, on the other hand, does not believe x realizes the property just described but still considers x to be immoral because B interprets x as realizing q , where q equates to an increase in harm, either to oneself and/or others (based on increased negative affect/attitude/behaviour), which is something B disapproves of.

According to this explanation, both A and B consider x to be immoral but for different reasons. It is not that A does not disapprove of something which causes increased harm, or that B does not disapprove of something which commends us to delight in the immoral; rather, it is that A does not believe that increased harm is a property realized by x , or does not prioritize it above a different property (commending one to delight in the immoral) which is held to be the main reason for A's moral disapproval. *Mutatis mutandis*, the same goes for B. After all, in the case of B, it is possible to believe that x does not commend one to delight in the immoral but still believe it can lead to increased harm. Conversely, one can believe that x commends us to delight in the immoral, and that this is reason enough for disapproval, irrespective of whether it causes any increased harm. In short, in this scenario, both A and B express a negative attitude towards x but for different reasons: that is, in virtue of believing that x realizes some property (P) which equates to something different in each case (p or q), but nevertheless serves the same reason-giving function. To illustrate, in the case of " x (*qua* virtual paedophilia) is morally wrong":

- (CEE a) A disapproves of p (where p equates to commending one to delight in the immoral) and believes that x realizes p (thus making anaphoric reference to that of which A disapproves).
- (CEE b) B disapproves of q (where q equates to increased harm) and believes that x realizes q (thus making anaphoric reference to that of which B disapproves).

In each case, the moral attitude towards x is the same: namely, "it is wrong". This is because some property (P), of which A (*qua* property p) and B (*qua* property q) disapprove, is believed by A and B, respectively, to be realized by

x. There can be any number of reasons for one's attitude towards something. What CEE teaches us is that moral attitude is no different.

6.2.1 *Constructing a Moral Norm*

Where a shared moral attitude occurs with regard to some object or event, as a society we are able to create or *construct* a social norm that then acquires its own objectified moral standard. As Prinz (2007) states, "Things that we construct or build come from us, but, once there, they are real entities that we perceive" (p. 168). With the force of social consensus, and the moral norm this creates, we can adopt a normative position whereby a particular (agreed) attitude is the one we *ought* to have, at least with regard to *this* object of moral inquiry. Copp (2011) likens this to what he calls *realist expressivism* (see also Copp 2001). Where S shares this attitude, we can commend her for doing so. Where S does not, it is appropriate (given the constructed moral norm's objectified status) to rebuke her for her alternative (some might even say deviant) moral attitude. This is because both the rebuke and a change of attitude on the part of S are deemed to be warranted (Nichols 2008).

In the case of "murder is immoral", I may share this attitude with a neighbour. Our shared attitude in turn aligns with the objectified moral norm of our society. Unlike my neighbour, though, I am not interested in what constitutes a violation of God's law and so cannot be said to have a negative moral attitude towards murder *because* it violates God's law. Despite these differences, we (my neighbour, myself and wider society) still express a shared negative attitude towards the act. These similarities and differences CEE is able to accommodate and explain. In this instance, how (for example) my neighbour and I interpret the act (the object of our moral inquiry) leads us to draw the same conclusion about whether we approve or disapprove of it. But, importantly, not because of the fact that we agree on what the act is (a token example of a type of intentional, illegal killing: namely, murder), nor because we agree on what properties the act realizes (because in this instance we do not, at least not completely) but, importantly, because we each disapprove of at least one property we believe the act realizes, even though this property is different for each of us. In other words, even though we both interpret the act in the same way (as murder), we nevertheless differ in terms of our beliefs about a certain property or properties it realizes, or how we prioritize these properties (i.e. a violation of God's law or something else: say, violating Kant's categorical

imperative or increased negative utility, a vice rather than a virtue, and so on). But the fact that we each believe that murder realizes *some* property of which we disapprove, even where we fail to agree on what this property is, means that our attitude towards murder is the same. Where enough people share the same attitude, but not necessarily for the same reason (*qua* belief about a property or properties realized), a social norm is constructed resulting in an objectified moral standard.

With run-of-the-mill first-person shooter games, however, in which one can enact random murder, while moral consensus is forthcoming in the case of actual murder – because such an event realizes a selection of properties at least one of which, but quite possibly more, the majority of people are willing to condemn (e.g. negative utility, a violation of God's law or Kant categorical imperative, a vice rather than a virtue) – this is less the case with regard to virtual murder. In essence, with immoral acts involving actual persons (i.e. any kind of sexual assault and murder), if one endorses CEE, such is the array of properties to disapprove of that it is simply a case of taking one's pick. In contrast, it is far less clear which properties are realized by the virtual enactment of an immoral act because this depends much more on how one *interprets* the event which, in turn, affects one's belief about the properties it realizes. Recall, for example, my discussion in [Section 4.3](#) on whether virtual paedophilia is child pornography. Therefore, whether I interpret a token act of virtual paedophilia as *child* pornography or just pornography, or neither, will likely affect the properties I believe the virtual event realizes and, depending on whether I approve or disapprove of these properties, my moral attitude towards this type of enactment. Conversely, it may be that how I interpret and therefore categorize a virtual event depends on the properties I believe it realizes. Again, using the example of virtual paedophilia, I may not categorize this type of enactment as child pornography because I do not believe it realizes the property of child abuse, and therefore do not have a negative attitude towards it for that reason at least.

When observing or interacting with a virtual event, if one cannot agree on what *that* is – and by this I mean what the virtual enactment is meant to represent and/or one's belief about the properties it is said to realize – then it becomes easier to see why we might find it harder to agree on whether *that* should be judged morally good or bad. Having said that, where different people do hold different beliefs about which properties a virtual event realizes then as long as these different properties are disapproved of (by each respective person holding the belief about their

realization), there is a good chance that moral consensus will be achieved and a shared moral attitude established. What I hope to show in the discussion to follow, which amounts to the application of CEE to the gamer's dilemma, is why there is a greater consensus of negative moral attitude towards virtual paedophilia than virtual murder. In addressing the 'why' question I hope to establish a normative ethic that is able to justify the selective prohibition of virtual content not only in relation to the gamer's dilemma but also beyond.

6.3 SOCIAL CONVENTION BY ANOTHER NAME?

One anticipated objection to CEE is that, in reality, it amounts to the same, albeit slightly more sophisticated, argument based on social convention rejected earlier (see [Chapter 2](#)). In response to this anticipated objection, I would say that it is precisely this added sophistication that enables CEE to overcome the problems raised against the social convention argument. Each, it must be said, adopts an anti-realist approach to moral utterances. While it is not my intention to defend my anti-realist stance in detail here (for further discussion, see [Young 2014](#)), a cursory exposition of my reasoning is required in order to show how CEE can overcome the challenges levelled at an appeal to social convention and therefore moral subjectivism more broadly construed.

When considering the proposition " x is morally wrong", one should not understand the sentence to be a *description* of some state of the world which captures a moral reality: namely, that x , so described, is literally picking out a moral wrongdoing. The problem with descriptivism in the context of moral realism is that moral utterances seem to be both descriptive and *evaluative*. To illustrate, the proposition "S is a paedophile" purports to describe some fact about S (that he is sexually attracted to children). This statement is truth-apt and is therefore either true or false. In addition, there is also an implied evaluative component: that being a paedophile is something one ought not to be because it is morally wrong. Moral utterances therefore contain both an 'is' (descriptive) component and an 'ought' (prescriptive) component, with the former being on a more secure metaphysical footing than the latter in terms of purporting facts about the world. Consequently, while it is true that S being a paedophile is either true or false based on how one defines paedophile and the sexual preferences of S, what is less clear is what makes it *true* (as in some independent fact about the world) that being a paedophile is morally

wrong? Recall, for example, the Ancient Greek custom of *paiderastia* (meaning *boy love*), mentioned in [Section 5.3.3](#). Certainly, in its day, this practice was not judged to be immoral. Contra moral realism, then, the evaluative component of the proposition “*x* is morally wrong” should not be thought of as a description of some mind-independent moral property; instead, it should be thought of as a direct *expression* of one’s attitude to *x* which functions in an evaluative way. In Ancient Greece, the attitude to adult men having sex with adolescent boys was certainly different to our *conventional* moral attitude; but in neither case – that is, either then or now – could the utterance “paiderastia is morally wrong” be said to be true or false: for it lacks truth-aptness. What it amounts to is either something one agrees with (*qua* approves of) or does not.

Embracing anti-realism with regard to moral utterances does not, in and of itself, help us differentiate between social convention and CEE. As we saw in [Section 6.2](#), the constructive component of this form of expressivism describes the social elevation of a shared attitude to the point where it achieves, through consensus, an objectified moral status; at which point it acts as an independent (of any individual) measure of morality and, in doing so, delineates what is morally warranted within a given society (again, independent of any individual moral attitude), and therefore what constitutes *that society’s* moral reality. How this differs from the social convention argument presented in [Chapter 2](#) is outlined below:

To state that A and B have a shared attitude towards *x*, such that they both hold that *x* is immoral, is to declare that they have the same *de re* attitude. When considering the act that A and B’s attitude is directed towards (the intentional object), their attitude towards that act (the thing in itself) is the same. But this shared *de re* attitude exists in virtue of the belief that *x* realizes some property (P) which they both disapprove of, but which can be (and is) different for A and B: A believes that *x* realizes *p* and B believes it realizes *q*. Their differing belief about which property is realized by *x* means that they have different reasons for their shared *de re* attitude. One could say that they have different *de dicto* attitudes regarding *x*... (namely, different beliefs about why it is immoral). (Young 2015b, pp. 317–318)

What A and B have in common is their negative attitude towards *x*. However, this singular attitude (it is singular because it is held by both A and B: hence, *de re*) is adopted by A and B for different reasons (they have different *de dicto* attitudes). It is therefore my contention that, in the

absence of the truth of moral-realism (which is the position I am endorsing here), a *de re* attitude shared by the majority of people within a given society (such that it becomes the constructed moral norm or, if you like, convention of that society) is more robust if it is the product of a number of different *de dicto* attitudes (i.e. if it is based on a number of different reasons for having the moral attitude). This should not be taken as evidence of inconsistency, and therefore as a reason to undermine the normative authority of the moral attitude; rather, and to reiterate, it should be taken as evidence of its robustness, insofar as there are purportedly many reasons for why *this* (whatever ‘this’ happens to be) is morally wrong. It just so happens that different people have different views on what these reasons are or how they prioritize them. Therefore, to undermine the moral (*de re*) attitude, one would have to undermine the various reasons (*de dicto* attitudes) justifying its (objectified) normative status. Such a position does not rule out a change of *de re* attitude – there is therefore a degree of fluidity inherent within the position – but it does make any such change less capricious.

An appeal to social convention is therefore vulnerable to capriciousness in a way that CEE is not: because the appeal to social convention presented in [Chapter 2](#) was said to be intertwined with the more elusive idea of moral intuition which is not the case with CEE. But what if one particular society has a shared negative attitude towards people of a particular race or sexual orientation or who hold certain religious beliefs? Given the lack of truth-aptness in regard to moral utterances, is the proposition “*These* people are morally inferior” just as valid as the converse utterance? In response, I would say that while the proposition itself lacks truth-aptness, what is capable of being true (or false) is (i) whether S *believes* that a particular racial type (for example) realizes a certain property, (ii) whether S disapproves of this property, and (iii) whether this particular racial type actually realize this property, thereby confirming or disconfirming the belief held by S. After all, it is possible for S and S’s society to ground their moral attitude on a false belief. Given this, one could challenge the reason (*de dicto* attitude) for the negative (*de re*) attitude should one believe that this reason stems from a false belief.

But suppose S believes that Afro-Caribbeans realize the property of darker skin (darker than S’s pale skin, at least) and disapproves of this property. S’s negative moral attitude towards Afro-Caribbeans is based on a belief that is in fact true and so S’s attitude cannot be challenged, in this instance, for being based on a false belief. Of course, one would simply ask

why S disapproves of darker skin – for example, what belief is the disapproval of darker skin based on? (etc.) – until, ideally, one arrives at a demonstrably false belief. I say ‘ideally’ because the (alleged) truth or falsity of certain beliefs may depend on how one interprets the evidence available. Some beliefs may therefore lack any definitive verification (i.e. one may fall back on a belief in the word of God, or what is written in some (held to be) sacred text). Such a situation makes it likely (inevitable, some might say) that moral disagreements will arise. This is indeed the state of our moral reality today. I see no reason to advocate moral realism because of this, however.

CEE therefore provides a more fine-grained explanation for why social conventions differ across different cultures/societies (e.g. Japanese attitudes to what in the West we would judge to be virtual child pornography; see [Section 2.1.1](#)). In the case of differences across societies or even disagreements within one’s own society, given that a society’s shared *de re* attitude is likely to be based on different reasons (*de dicto* attitudes), where these reasons are few(er) or perhaps less entrenched, there is more chance that one could challenge the social norm (*de re* attitude) by undermining the different beliefs (*de dicto* attitudes) held by different members of that society (assuming one believes the *de re* attitude to be grounded, at some point, on a false belief). If, for example, the only reason for a society’s negative moral attitude towards murder is that it violates God’s law then, if one wished, one could challenge the merits of this reason by trying to undermine the belief in the existence of God and, with it, the putative authority of God’s law. Fortunately (as already noted), there are many reasons (*de dicto* attitudes) for why one should have a negative (*de re*) attitude towards murder such that undermining them all would prove difficult: thereby accounting for why a negative attitude towards murder is universally expressed and entrenched within different cultures/societies.

In the case of virtual enactments, establishing whether a certain property is realized by a particular enactment is harder to do because of the greater scope for different interpretations compared to actual events. Consequently, it is harder to challenge the *belief* on which the moral attitude is grounded (whether it is the belief that P is realized by *x* or the belief that P is not realized by *x*). Does virtual murder realize the property of eliciting delight in the idea of actual murder, for example? Likewise, does it elicit the property of being a vice or negative utility, and so on? It is also the case that actual events will realize different properties compared to virtual enactment. These former properties are easier to

establish in terms of the truth of their occurrence (i.e. the negative utility of actual murder compared to virtual murder), thereby making it likewise easier to validate a *belief* in their realization (or refute it, depending on the belief).

6.4 APPLYING CONSTRUCTIVE ECUMENICAL EXPRESSIVISM TO ALI'S ACCOUNT

Pace Ali (whose argument was discussed in [Chapter 5](#)), I do not consider it controversial to accept the claim that gamers (or even the wider community) typically hold virtual murder to be less objectionable than virtual paedophilia. As we have seen, there have been a number of suggestions/arguments proposed to account for why this is. Each has posited a single factor (a) in the role of the 'morally relevant means of differentiating between virtual murder and virtual paedophilia' and (b) given (a), as the basis for a normative ethic that guides the selective prohibition of wider video game content. While some have shown more promise than others, in my view, each has failed to deliver on points (a) and (b). CEE, in comparison, does not seek to identify a single morally relevant factor. In accounting for the difference in moral attitude between virtual murder and virtual paedophilia, CEE is able to accommodate a number of different reasons based on differing *de dicto* attitudes (*qua* beliefs) about properties realized by the respective virtual event. Why is it more likely that a gamer will find virtual paedophilia morally objectionable? Because the gamer holds the belief that at least one property of which they disapprove is realized by this type of enactment (the converse epistemic relationship typically occurring in the case of virtual murder). These may include, but are not restricted to, prurient appeal (i.e. delighting in the idea of actual paedophilia), the likelihood of harm (whether to children or women), the targeting of non-morally relevant characteristics and so on. Each one of these beliefs (and others) has been challenged and shown to be problematic with regard to (a) and (b). This is partly because of the different ways the virtual event can be interpreted, which also relates to player motivation ([Young 2015b](#)). Importantly, though, the same ambiguity that works against satisfying (a) and (b) helps maintain those beliefs the gamer uses to ground their moral attitude (recall, the belief can function in this role even if it is false or not verified). It also accounts for why, in accordance

with Ali (2015), certain gamers may hold that some token acts of virtual paedophilia are not morally objectionable (or any more so than certain token acts of virtual murder). This is because, according to CEE (which is not incompatible with Ali's view), the gamers do not believe that these tokens virtual events, in the context in which they occur, realize the properties of which they disapprove. This does not rule out the possibility that other token acts of virtual paedophilia will amount to tokens of a certain type that *do* (they believe) realize these properties owing, say, to the context in which they occur (again, as Ali attests). Similarly, where virtual murder may appear to be gratuitous and/or targeted (e.g. racist) and/or excessively violent, in the absence of a morally mitigating context, gamers may likewise believe that such enactments realize the same or similar properties to those realized by certain enactments of paedophilia, meaning that these should be disapproved of, also.

Gamers (and the wider community) typically hold a more negative moral attitude to virtual paedophilia because, collectively, there are believed to be more *de dicto* reasons for disapproving of this type of enactment compared to virtual murder (irrespective of whether these *de dicto* reasons *qua* beliefs are true). Here, I think Ali's point about the default position of gamers who are asked about virtual paedophilia is germane. Because there are no commercially available games that enact the types of hypothetical scenarios gamers are often asked to envisage, it is plausible that they struggle to contextualize virtual paedophilia within a mitigating narrative (i.e. there are no existing examples for them to draw on, unlike virtual murder). Such a mitigating narrative could negate the belief that a token act of virtual paedophilia realizes a certain property of which the gamer would otherwise disapprove. In the absence of such a belief, CEE teaches us, the gamer has no reason to adopt a negative moral attitude. But this also means that in the context of a (hypothetical) simulation game, perhaps along the lines of *Child Sexual Assault*, should a gamer have no reason to believe that a certain property they disapprove of is realized through enacting a token paedophilic event (even 'for its own sake', although it remains unclear what Ali means by this) then that gamer would, and indeed should, have no reason to disapprove of the enactment. Should this apply to the wider community of gamers and beyond then CEE would predict the establishment of an objectified moral norm that would permit virtual paedophilia, at least of the type alluded to here.

6.5 APPLYING CONSTRUCTIVE ECUMENICAL EXPRESSIVISM TO PATRIDGE'S ACCOUNT

In the context of *Child Sexual Assault*, recall how Patridge holds that virtual paedophilia is an extension of our lived morality because the enactment is a further example of targeting for harm (albeit virtual harm) based on non-morally relevant personal characteristics (in this case, being a child). One is actually targeting something (an avatar) based on characteristics the avatar is represented as having (child-like features). Even targeting based on *representations* of non-morally relevant personal characteristics is, for Patridge, immoral: because the act of targeting for the reasons described is performed elsewhere and on other occasions (i.e. actual child abuse) as part of our lived morality, and it is immoral to do so. *Child Sexual Assault* is therefore just an extension of that immoral activity under the guise of a game.

Interpreting Patridge through the theoretical lens of CEE, we could say that she disapproves of p (targeting for harm based on a non-morally relevant personal criterion) and believes that x (a token act of virtual paedophilia) realizes p . Therefore, she has a negative moral attitude towards x in virtue of her belief that x realizes p and her disapproval of p . Let us allow that my friend (the one I introduced earlier in this chapter) has the same negative moral attitude towards x for the same reason. I, on the other hand, do not believe that x realizes p . Moreover, I do not believe that x realizes any property I disapprove of. Consequently, I do not have a negative moral attitude towards x .

Why do I not believe that x realizes p ? Perhaps it is because I do not believe that targeting a *virtual* character, based on *representations* of non-morally relevant characteristics, in order to engage in *simulated* immoral activity is equivalent, morally, to targeting an *actual* person (individually or as part of a minority group) based on *actual* non-morally relevant characteristics for an *actual* immoral activity (e.g. harming them). Importantly, the representation I have in mind is akin to Baudrillard's (1983) notion of a *simulacrum* (mentioned in Section 4.3; meaning a copy or a representation of a thing that has no original but goes beyond itself) and so is not meant to include an actual photograph or similar recording. Let us say that the reason I do not believe they are equivalent is because although the putative non-morally relevant characteristics of the representation are *characteristics*, they are characteristics of a *representation*, thereby making the characteristics of the representation at the same

time representations of characteristics of an actual person (but not a specific person, in accordance with Baudrillard's simulacrum). As such, whether targeting these characteristics is of concern in any *moral* sense is debatable. In short, what I am contesting is whether representations of non-morally relevant characteristics are (and indeed can be) non-morally relevant characteristics of a representation. I am contesting this insofar as I am contesting whether all representations of this kind have any moral relevance at all.

In making this point, I am reminded of Edward Castrovona's notion of a *closed world*. According to Castrovona (2004), the border between our real world and a closed virtual world is impermeable. Consequently, while the virtual world may represent aspects of the real world, the features of each world are incommensurate. For Castrovona, closed worlds are spaces in which nothing matters; where assets have no real value and therefore losses are unimportant. In such a world, what does it matter if one murders someone for no other reason than one can, or abuses a child for the same reason? After all, in this world, nothing matters. Dunn (2012), however, casts doubt on the plausibility of a world in which nothing matters; where everything is without meaning. Why, Dunn asks, would anyone want to inhabit such a world? This is a pertinent question. If one's enactment was completely devoid of meaning then what would be the attraction? Therefore, for Dunn, the border between the two worlds is necessarily *not* impermeable. The meaning of the enactment is derived from what the enactment represents. Consequently, real-world meaning necessarily transcends the two worlds. If I seek to satisfy vicariously my desire to carry out the act for real then the enactment is meaningful as a substitute for that action. On the other hand, if I seek to enact murder or paedophilia because what I am enacting is a taboo in the real world then, similarly, the enactment's meaningfulness as an act I desire to perform is parasitic on what the enactment is meant to represent: namely a *taboo*. Such *unidirectional* transcendence is necessary for the virtual act to be imbued with any meaning *as an enactment*. What is contested, however, is whether the meaning of the virtual act itself transcends worlds so that it has real-world meaning, particularly moral significance. If the meaning of a virtual act does not transcend gamespace, thereby conveying *real-world meaning* and subsequently moral import, or if only certain actions do or should be said to, then the virtual world could be thought of as *partially closed* (Dunn 2012).³

It is the partially closed nature of the virtual world that gives me reason not to believe (contra Patridge and my friend) that x realizes p . Which of us is correct? That is, does x realize p or not? As I have argued previously, in the context of virtual enactments, this is difficult to verify because it is dependent on how one interprets x (x being much more open to interpretation when it constitutes a virtual act within gamespace). Such dependence on interpretation does not prevent the formation of *beliefs* about the realization of p , however. Instead, and to reiterate, it makes it difficult to *verify* the belief.

What CEE allows (although it does not require this) is that my friend and I can have a shared negative moral attitude towards actual paedophilia (that it is morally repugnant) while agreeing that what we are enacting within *Child Sexual Assault* is virtual paedophilia. We are simulating the very thing we disapprove of in the real world. Yet it does not require that this real-world accord about actual paedophilia and our agreement on what is being simulated carries with it a further *moral* agreement about the rights and wrongs of what we are doing when enacting virtual paedophilia. As I have outlined, CEE accounts for different moral attitudes, as well as shared attitudes for different reasons. The only consistency required is between one's attitude and one's belief(s). Where I disapprove of p and believe that x realizes p then I should (in both a rational and moral sense) disapprove of x (*qua* possess a negative moral attitude) in virtue of my belief that x realizes p .

If I believe that the virtual targeted murder of homosexuals within a video game (say, when playing *R.A.C.I.S.T.*) is morally wrong then, according to CEE, it is because I believe that this virtual act realizes some property (P) that I disapprove of. If I do not have a negative moral attitude towards *S.H.: Random Attack* then it is because I do not believe that a property of which I disapprove is realized by any of the token virtual engagements within this game. Likewise, I may hold that the (fictitious) video game *Sexual Assault* is morally wrong for similar reason to those given when discussing *R.A.C.I.S.T.*, or hold a negative attitude towards only certain token enactments within this game: say those involving the sexual assault of minors (again, for similar reasons to those just discussed). CEE also accounts for why some gamers may consider non-visual or implied acts of virtual paedophilia, or virtual grooming (etc.) to be morally wrong and, in doing so, is able to accommodate a broader approach to the gamer's dilemma, as suggested by Luck and Ellerby 2013; see [Section 4.1](#)).

6.6 ESTABLISHING A NORMATIVE ETHIC

CEE is fundamentally a meta-ethical approach concerned with understanding the nature of moral utterances. As part of this understanding, its focus is primarily on the process by which an individual comes to have the moral attitude they do, followed by how different individuals, possibly with different beliefs, can nevertheless have the same moral attitude. What CEE posits is that a shared *de re* attitude need not be the result of a shared reason (*de dicto* attitude); rather, it is simply the product of a consistency between belief (*de dicto* attitude) and *de re* moral attitude: the *de re* attitude being the appropriate (*qua* rational) expression of one's *de dicto* attitude. What CEE accepts is that one's moral expression is the product of a belief, not necessarily a fact, and so is not itself a moral truth. Given this, is CEE robust enough to establish a normative ethic? In order to address this question, let us consider what CEE as a normative ethic would look like.

According to CEE, what we ought to do is determined by our moral attitude; and when I say 'our' I mean, of course, the attitude indicative of the objectified moral norm within our society. Such a norm is constructed; it is the product of the moral attitude shared by the majority (but, to reiterate, not necessarily for the same reasons). In [Section 6.3](#), I argued for the robustness of this objectified moral norm compared to the type of social convention discussed in [Chapter 2](#). I also described how CEE permits moral change – and so is fluid – without being capricious. What CEE advocates, then, is a morality and, I would argue, a normative ethic based on the shared moral attitude of a given society (or the majority within that society) which evolves into an objectified moral norm. Where an individual's moral attitude deviates from this norm then their attitude deviates from what it *ought to be*. In the case of my own (fictitious) moral attitude towards enactments of virtual paedophilia within *Child Sexual Assault*, I argued that this was based on a particular belief I held regarding *x* (a token enactment of paedophilia) and a particular property I believed it did not realize, thereby giving me no reason to disapprove of it. If, however, the majority of gamers (and even wider society) believe differently, such that their belief(s) lead them to share the same negative moral attitude towards virtual paedophilia, thereby making it an objectified moral norm, then my attitude is off-kilter with theirs. A change in my attitude is therefore warranted.

What I could try to do is challenge this moral norm by challenging the beliefs on which it is founded; by trying to show either, *a priori*, that the

respective beliefs of those who have contributed to the moral norm are inconsistent or conceptually flawed, or that, *a posteriori*, they are not supported, or perhaps are even refuted, by empirical evidence (much as I have been trying to do throughout this book). Again, recall that a *de re* attitude can be grounded on any number of beliefs *qua de dicto* attitudes and therefore a number of different arguments may need to be forwarded to quash different beliefs. Should I prove to be successful at challenging the beliefs on which a particular moral attitude is based, resulting in the formation of new beliefs, alongside a corresponding change of attitude, then this change of attitude will become the new social norm and therefore the new proscribed way of thinking about *x*. As an aside, I have left unexplored the question of whether one has the *legal* right to have an alternative moral attitude and therefore whether we should defend this right.

6.7 CONCLUDING REMARKS

Where does this leave us with regard to the gamer's dilemma? With qualification, I accept that most gamers within the West and also our wider (Western) society find the idea of virtual paedophilia more morally objectionable than virtual murder. (I say 'with qualification' because some gamer's and non-gaming members of society may differ in their moral attitude depending on the *context* in which the virtual paedophilia and the virtual murder occur, as discussed in relation to Ali, and as allowed by CEE.) When attempting to resolve the dilemma, it is first necessary to understand that the shared moral attitude on which the dilemma is premised is likely to be grounded on different beliefs about virtual paedophilia which are different to corresponding beliefs about virtual murder. This means that there is not just one reason for why a difference in moral attitude exists but potentially many. Consequently, looking for a single morally relevant factor for why the dilemma exists is futile, as is challenging any attempt at resolving the dilemma based on the same single factor. Therefore, in trying to resolve the gamer's dilemma, I have presented CEE as a means of accounting for why this shared moral attitude towards virtual paedophilia exists and why this differs (typically) from our shared moral attitude towards virtual murder.

In conclusion, according to CEE, ultimately, the premise on which the gamer's dilemma is built is the product of different beliefs and not differences in some mind-independent moral fact that differentiate virtual paedophilia

from virtual murder. Beliefs, as has been discussed, do not have to be true; but, once held, one's moral attitude needs to be consistent with them. Attempts at resolving the gamer's dilemma have previously failed because they have targeted single factors (*one* morally relevant difference). What I have argued is that the premise on which the dilemma is built stems from a difference in attitude which itself is not based on a single factor or a single morally relevant difference. To resolve the dilemma, one would need to undermine each or a large number of the different beliefs which ground the single moral attitude (objectified social norm). This, in itself, may be difficult to achieve given that the beliefs are likely to be based on a particular interpretation of the intentional object which, for some, may be less amenable to reinterpretation, particularly in the absence of contradictory evidence.

As a normative ethic, CEE posits an objectified moral norm that is constructed within a given society in virtue of a (majority) shared attitude. What is morally acceptable is therefore based on what the majority consider to be morally acceptable in virtue of their approving moral attitude. A *de re* attitude may be challenged and change over time, but only when a sufficient number of beliefs on which the *de re* attitude is grounded change. Where a number of different beliefs contribute to a particular attitude and where some/all of these are difficult to undermine, the attitude and hence the social norm will be maintained. Of course, this potentially limits the universality of the normative ethic, given that different societies may have different beliefs and therefore attitudes – certainly in the case of virtual enactments – but this closely matches our moral reality and is not therefore reason enough to dismiss CEE.

NOTES

1. I appreciate that there may be occasions when one is both amused and morally disapproves: say, when finding a joke amusing despite disapproving, morally, of the inherent sexism.
2. An anaphoric reference occurs when a word in a text refers to a previous idea in the text for its meaning. In the sentence "Fred always looked unkempt but this never seemed to bother him", the word 'him' makes anaphoric reference to Fred.
3. My use of Dunn's term is slightly different to his original usage.