

Pre-print version of the paper published in *Ethics and Information Technology* (2014), 16, 91-102.

A meta-ethical approach to single-player gamespace: Introducing constructive ecumenical expressivism as a means of explaining why moral consensus is not forthcoming

Garry Young

The morality of virtual representations and the enactment of prohibited activities within single-player gamespace (e.g., murder, rape, paedophilia) continues to be debated and, to date, a consensus is not forthcoming. Various moral arguments have been presented (e.g., virtue theory and utilitarianism) to support the moral prohibition of virtual enactments, but their applicability to gamespace is questioned. In this paper, I adopt a meta-ethical approach to moral utterances about virtual representations, and ask what it means when one declares that a virtual interaction ‘is morally wrong’. In response, I present constructive ecumenical expressivism to (i) explain what moral utterances should be taken to mean, (ii) argue that they mean the same when referring to virtual and non-virtual interactions and (iii), given (ii), explain why consensus with regard to virtual murder, rape and paedophilia is not forthcoming even though such consensus is readily found with regard to their non-virtual equivalents.

Keywords: Descriptivism; Moral realism; Moral supervenience; Moral perception; Virtual ethics

1. Introduction

The issue of whether virtual representation and enactment (hereafter, VR/Es) within single-player video games should be the subject of moral scrutiny and the corollary concerning whether they should be morally policed continues to be debated (see Young, 2013, for a detailed discussion).¹ Those positioned within the amoralist camp, as the name suggests, hold that VR/Es are not a morally matter. Among the reasons given for this is that a physical description reveals VR/Es to involve nothing more than the manipulation of pixels. Even if one moves beyond the physical description, so as to include reference to what the pixels are designed to represent – say, a brutal assault on an innocent victim – given that this representation is taking place within, and in fact should be *confined to*, the space of play, any putative moral concern is ultimately misplaced: for what is occurring is *just* a game. In contrast, there are those who, for various reasons, consider that such game content should be the subject of moral concern. Focusing on this group, which I will call *moralists*, how should we understand their declaration that it is legitimate for certain VR/Es (say, those involving murder or sexual assault or paedophilia) to be judged morally wrong? In other words, what should we take this claim to mean?

To be clear, I am not asking *why* moralists consider certain VR/Es to be morally wrong; nor am I concerned with whether statements pertaining to their moral wrongness (or rightness) are true or false. Instead, by adopting a meta-ethical approach, I seek to clarify what it means to say that some VR/Es are morally right or wrong, good or bad (although I accept that an

¹ By ‘single-player video games’ I mean games played on personal computers (PCs) or consoles such as X-Box 360/One, PlayStation 1-4, and Wii. Although such games may be available in formats that allow one player or a number of players (connected through the Internet) to play, I restrict discussion to situations in which a single player plays alongside and/or against non-player characters (NPCs).

answer to this question will help inform the ‘why’ question). After all, in order to continue the debate on gaming content, particularly content representing violent or otherwise taboo activities, we should be clear on what claims to moral wrongness or rightness mean within the context of single-player video games. A major aim of this paper, then, is to clarify the meaning of such moral utterances as a prerequisite to any further debate on the moral status of VR/Es.

To illustrate the potential for moral ambiguity, is the proposition “VR/E(x) is morally wrong” truth-apt, insofar as any claim it makes about moral wrongness (or rightness) is capable of being true or false irrespective of whether it actually is? Or is it that such a claim is incapable of this because the idea of truth or falsity is meaningless in the present context: in much the same way as trying to establish the truth or falsity of the proposition “Boo VR/E(x)” is meaningless and therefore misguided. If “VR/E(x) is morally wrong” is truth-apt, then what is involved in the proposition being true or false? Conversely, if it is not truth-apt, then what should we take such a proposition to mean?

In order to address the questions above – which specifically concern the meaning of moral claims about VR/Es – a useful strategy would be first to adopt a meta-ethical approach to more general moral statements. The same questions could therefore be directed towards the proposition “ x is morally wrong”. Given this fact, whether a different conclusion should be drawn in relation to VR/Es will depend on the more general conclusions drawn regarding the truth-aptness of “ x is morally wrong” and the applicability of these to propositions concerned with the moral status of VR/Es. It may be that any general conclusions are able to transcend the virtual and non-virtual spaces thus making them universally applicable and so able to

account for moral statements about VR/Es. Or it may be that the virtual and perhaps even playful nature of single-player gamespace prevents conclusions about the proposition “VR/E(x) is morally wrong” from being subsumed within the more general conclusions about x .²

A further aim of this paper is therefore to contrast the propositions “ x is morally wrong” and “VR/E(x) is morally wrong”, where VR/E(x) is a virtual representation/enactment of x . In the context of meta-ethics, do the two propositions convey the same moral meaning? If so, what does each statement mean, morally? If not, then why not? Is it, for example, that the amoralist is correct when it comes to moral judgements about single-player gamespace: namely, that there is no case to answer; meaning that, at best, the proposition is merely mimicking moral content? Alternatively, if they do convey the same meaning then what should we understand a moral statement to be and is such a statement able to convey truth-aptness in the case of murder *and* virtual murder, rape *and* virtual rape, paedophilia *and* virtual paedophilia (and so on), or is it that neither is truth-apt?

It is my contention that the propositions “ x is morally wrong” and “VR/E(x) is morally wrong” convey the same moral meaning insofar as each expresses the subject’s *attitude* towards x or VR/E(x), respectively. As such, I intend to argue against moral realism and endorse a form of expressivism which I am calling *constructive ecumenical expressivism*. Despite being anti-realist, this new approach is able to account for the high level of moral consensus (objectification) regarding views on the morality of murder or rape or paedophilia. It is able to explain why this is so in cases of non-virtual actions, but also why there is much

² Owing to the need for brevity, the paper does not concern itself with the issue of morality and play. For discussion on this, see Young and Whitty (2012).

less moral consensus when it comes to the virtual representation and enactment of the aforementioned activities, leading some to the view that, despite the moralists objections, single-player gamespace should be an amoral space.

Given that this is the first presentation of constructive ecumenical expressivism, it is important to spend some time preparing the way for this novel approach. In defence of constructive ecumenical expressivism, I begin by outlining problems with descriptivism and moral realism (Section 2). I then introduce moral supervenience (Section 3) and the constraint this imposes on one's moral attitude towards identical actions and representations. Following this, I introduce those existing moral positions that have most influenced constructive ecumenical expressivism. I begin by outlining traditional expressivism (Section 4), then a variation on this, ecumenical expressivism (Sections 5), before introducing constructive sentimentalism (Section 6). After discussing the respective strengths of ecumenical expressivism and constructive sentimentalism, which I have integrated into this new meta-ethical approach, I present the case for constructive ecumenical expressivism (Section 7). Finally, although the overarching aim of this paper is simply to give an account of the *meaning* of moral utterances within gamespace, towards the end of this section, I briefly discuss how this new meta-ethical approach might inform future debate on the morality of gaming content.

2. Descriptivism and moral realism

Should we understand the sentence “ x is morally wrong” 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? Descriptivism holds that moral language (such as the proposition “ x is

morally wrong”) and non-moral language (for example, “Mount Everest is the highest mountain in the world”) share the same syntactic structure and must adhere to the same rules of logic. Therefore, in the same way that it is contradictory for Mount Everest to be described as both the highest and not the highest mountain in the world, so x cannot be described as morally wrong and not morally wrong without violating the same law of contradiction. Likewise, just as the proposition “Mount Everest is the highest mountain in the world” predicates Mount Everest with being the highest mountain, so “ x is morally wrong” predicates x with having or being in a state of or constituting moral wrongness.

In the case of the description of Mount Everest, however, what this description is said to be picking out – the highest mountain in the world – exists independently of the description. It is therefore an unremarkable fact that the truth or falsity of the description of Mount Everest is dependent on the truth or falsity of those independent characteristics of Mount Everest the description depicts. In other words, the truth or falsity of the proposition “Mount Everest is the highest mountain in the world” is dependent on the truth or falsity of Mount Everest *being* the highest mountain in the world. For the same to be said of moral descriptions, there must exist independent morally good or bad things for the description to pick out in a manner that makes the description either true or false. Consequently, when describing subject S as morally good, the possibility of this statement being true is dependent on the independent truth of S being morally good or of possessing some attribute correctly described as moral goodness.

Descriptivism with regard to moral utterances is dependent on the truth of moral realism if it is to do anything more than postulate examples of correct syntax and non-contradictory propositions in the absence of an independently existing referent (such as “unicorns have a

horn and cannot be all black and all white” (see Shafer-Landau, 2003). The problem with descriptivism in the context of moral realism, however, is that moral utterances seem to be both descriptive and *evaluative*. They 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; that is, in terms of having truth-aptness.

To illustrate: the proposition “*x* is morally wrong” contains a nominal descriptive component, *x*, but the second component seems to go beyond mere description. To describe *x* as ‘being wrong’ does not seem to be on par with describing *x* as cylindrical or as made of wood (for example). If someone were to say that “*x* is morally wrong” then how should we understand this utterance? What function does the predicate ‘is morally wrong’ serve and is this also served by the predicate ‘is cylindrical’? Certainly both tell us something about *x*. But do the predicates ‘is cylindrical’ and ‘is morally wrong’ each pick out independent properties of *x*?

Where *x* is an action, the latter component (‘is morally wrong’) does not serve merely to describe action *x* in more detail. Uttering “shooting *S* is morally wrong”, adds more detail to *x*, insofar as we now know that it entails the shooting of *S*, but to then describe the action as morally wrong is to do more than simply add to the descriptive quality; rather, it functions to inform us that *x* (*qua* shooting *S*) is something that one *ought not* to do. What is not debated is that the descriptive component picks out some aspects of the world that is at least capable of being true or false: the occurrence of *x* or, more specifically, the shooting of *S*. Whether the same can be said of the latter component, however, is a more contentious issue (i.e., that the

act of shooting S has or is capable of having the property of ‘wrongness’ which is picked out within the subject’s description of x as being morally wrong).

In the case of “shooting S is morally wrong”, the descriptive and evaluative components are easily demarcated. On other occasions, however, although this difference is less apparent, it is still to be found. Take the proposition “A is a murderer”. Here, the description of A as ‘a murderer’ is truth-apt given that A engaged in the premeditated and unlawful shooting (and killing) of S. But describing A as a murderer is also evaluative insofar as it intimates that what A has done (committed murder) is wrong and so ought not to have been done. But does the assertion that it is wrong to murder have truth-aptness? For the moral realist, to have truth-aptness, wrongness must be a property of murder that can be picked out. There are lots of ways that S could have been murdered other than by being shot, of course, so perhaps the wrongness we are looking for is not an exclusive property of the act of shooting someone, but of what some shootings and all other acts of murder have in common: namely, that they are premeditated and unlawful killings. We can then claim that whatever fits this description possesses ‘wrongness’.

There are two issues with this move. First, it is not clear how one could describe ‘unlawful’ in a way that is not circular or culturally and historically dependent, thereby making the concept ‘unlawful’ (which is a key part of what it is to commit murder) *not* an independent property of the world. Of course, moral realists do not require *all* putatively moral claims to pick out correctly independent truth-apt moral properties of the world, only some; but, even if this is the case, murder would seem to be a paradigm example of moral wrongness and therefore something that we need to be able to proffer an independent description of.

Second, putting aside the seeming lack of independence of the concept ‘unlawful’ and therefore ‘murder’, perhaps more importantly, what remains unclear is the relationship between premeditated and unlawful killing and wrongness? It is not obvious that such a description *entails* ‘wrongness’; nor is the relationship analytic (such that an analysis of the concept of ‘premeditated, unlawful killing’ contains the concept of ‘wrongness’). I will return to this last point in Section 3 when discussing supervenience.

Of course, moral realists might object to the argument above. It may be claimed, for example, that a conservative theologian does not object to murder simply because it is unlawful but because murder being wrong is a theological fact.³ In response, one might enquire as to the basis of this ‘fact’. Is the alleged fact that murder is wrong a fact because God decrees that it is? If it is, it would appear that murder being wrong is contingent on the will of God. As such, if God willed that murder is good would it then be a theological fact that murder is good? Alternatively, is murder being wrong a moral fact independent of God’s will and therefore something God cannot will to be otherwise? If this is the case then one might question the omnipotence of God: for there would appear to be at least one moral fact that even God must subscribe to.

If the wrongness of murder is a moral fact (that even God must subscribe to) then, although it is not reduced to unlawfulness, it nevertheless remains unclear why it is wrong. The grounds for its status as a brute fact have yet to be established. If, on the other hand, murder is wrong because God says it is, and for this reason forms part of God’s law/moral code, then, to

³ I thank the anonymous reviewer for raising this point.

reiterate, such a ‘fact’ appears to be contingent on the will of God and is not *necessarily* the case. Consequently, it is still reducible to unlawfulness, only this time in relation to God’s law. A moral fact that is not necessarily, but only contingently, so undermines moral realism.

So far, I have argued that moral utterances have both a descriptive and evaluative component, even where the same word serves as both a description and an evaluation (e.g., murderer). How might the discussion on descriptivism and the argument against moral realism inform the issue concerning the respective meanings of “*x* is morally wrong” and “VR/E(*x*) is morally wrong”? Consider the following:

1. Shooting S is morally wrong
2. Virtually shooting S is morally wrong
3. A is a murderer
4. Virtual A is a virtual murderer

Each statement (1-4) has a descriptive component (‘shooting S’ and ‘virtually shooting S’, ‘A being a murderer’ and ‘virtual A being a virtual murderer’). In each case, the evaluative component is the same: it is ‘morally wrong’. This is either explicitly stated in the case of (1) and (2) or implied in (3) and (4). All propositions can therefore be said to conform to the same structure: namely, [description] + [evaluation], although in the case of (3) and (4), the descriptive term also serves as the evaluative term. The descriptive component, whether referring to the virtual or non-virtual, does not describe a moral property; it does not pick out that which is morally good or bad, right or wrong. The moral work is being done by the

evaluative component. In the next section, I consider the relationship between the description and the evaluation in more detail.

3. Moral supervenience

If each moral utterance contains two components which function in different ways – namely, descriptive and evaluative – such that, in the case of the latter evaluative component, the predicate “is morally wrong” not only describes that property (of ‘wrongness’) as moral realism claims but does so as a means of evaluating x or $VR/E(x)$, then what is the relationship between each component? Moral supervenience holds that a moral property M (the ‘wrongness’ of x , for example) is said to be *supervenient* on a physical (non-moral) property P such that there can be no change in M without a change in P (see Blackburn, 1971, 1984, 1985). This supervenient relation is famously illustrated by Hare (1952):

Suppose that we say ‘St. Francis was a good man’. It is logically impossible to say this and to maintain at the same time that there might have been another man placed exactly in the same circumstances as St. Francis, and who behaved in exactly the same way, but who differed from St. Francis in this respect only, that he was not a good man. (p.145)

The relationship between the moral property of ‘goodness’ and the non-moral or physical properties that constitute St Francis – his physical description – is such that there can be no moral change in the absence of a physical/behavioural change. This, of course, does not mean that *any* physical difference must entail a moral difference. If a physical duplicate of St Francis were suddenly to find that his hair had gone grey overnight or that he had suddenly

become short-sighted, then it is unlikely (or certainly would not necessitate) that this would have any impact (detrimental or otherwise) on his ‘goodness’. Any physical difference between the duplicate saints, if it is to be cited as the reason for a moral difference between the two men, must constitute *reasonable grounds* for this difference (Sidgwick, [1874] 1981); it must be a physical difference of the right kind.

It seems to be the case that if it is said of *a* that *a* is morally good then, *a priori*, where *a* is identical to *b*, this entails that it must likewise be said of *b* that *b* is morally good. A failure to conclude that *b* is morally good indicates a failure to understand moral concepts. A consequence of this *a priori* truth is that the moral property of ‘goodness’ – which in this case renders St Francis a good man – is dependent on some physical (non-moral) property of the right kind. But ‘moral goodness’ is not reducible to this physical property; it is not something that we can describe in purely physical terms. As such, in addition to there being no change in M without a change in P, moral supervenience holds that M is irreducible to P (Ridge, 2007). Neither is the relationship analytic whereby a physical description of *a* (or *b*, if identical) would reveal to us those components that make up moral goodness. In other words, the proposition ‘*a* is morally good’ is not analytic in the way ‘a triangle has three sides’ is. Furthermore, we do not directly detect, *a posteriori*, the actual instantiation of a moral property (say, wrongness); rather, we “conclude that particular acts are wrong in virtue of some empirically detectable feature; for example, because it *causes pain*, *involves deceit*, or *violates an agreement* (etc.)” (Coons, 2011, p.85; emphasis in original) and, in doing so “postulate the unobserved to explain the observed” (*ibid.*) (see Section 5).

3.1 Moral perception

McNaughton (1988), in contrast, argues that moral properties are directly observable: for if “we are prepared to allow that I can see that this cliff is dangerous, that Smith is worried... [etc.] then there seems no reason to be squeamish about letting in moral observation” (p.57). Thus, for McNaughton, observing that *a* is good (or bad) is compatible with perceiving that the cliff is dangerous (see also Chappell, 2008). Do we perceive ‘dangerousness’ directly as McNaughton would have us believe?

Gibson (1979) proposed that potential for action could be directly perceived in the form of affordances. An affordance property should be understood as a relational property that reflects characteristics of the object and the action capabilities of the subject. Typically, a path through the jungle affords traversability (for example), or a rock affords grasping and throwing. Each of these properties can be directly perceived. A steep cliff face may afford a lack of traversability to the untrained eye (that is, someone who is not a trained climber and/or does not have the necessary equipment). It may also afford rapid descent in the guise of falling, but falling is not the same as dangerousness; although it is likely to lead someone to conclude that the cliff face is dangerous for this reason. Certainly, one can connect the judgement that it is dangerous to some physical characteristic of the cliff, in conjunction with physical characteristics of the subject (i.e., the potential for rapid descent *for me*), but this is not the same as saying that dangerousness is a property directly observable. Instead, the dangerousness of the cliff face is a conclusion one draws about this object in virtue of the fact that it affords (among other things) limited traversability and rapid descent; in much the same way that the goodness of *a* is a conclusion one draws about *a* given the fact (let us say) that *a* increases overall pleasure to me but not at the cost of increased displeasure for anyone else (again, see Section 5). To say “I can see that the cliff face is dangerous” is not to see ‘dangerousness’ as an object of perception; rather, it is to see that the cliff face has limited

traversability, etc. (in virtue of what it affords) and so warrants the judgement “It is dangerous (for me)”.

3.2 Attitudinal supervenience

When spoken intelligently and sincerely, the respective propositions “It is dangerous (for me)” and “*a* is morally good” are equivalent to asserting “I believe it is dangerous (for me)” or “I believe that *a* is morally good”. In light of this, and focusing again on moral properties, given that the relationship between *a* and moral goodness is not analytic, nor is moral goodness directly detectable *a posteriori*, it is left unclear how our moral beliefs are justified. A way to overcome this problem is to hold that the proposition “*a* is morally good” does not express a belief about *a* in terms of some independent moral fact – that can either be true or false – which one is obliged to justify in light of the possibility of it being false. Instead, it is an expression of one’s attitude towards *a* that, broadly construed, amounts either to one’s approval or disapproval of *a*.⁴ Importantly, the supervenient constraint should still hold in the case of one’s moral attitude: for, where *a* is identical to *b*, if one morally approves of *a* then one should morally approve of *b* also.

Contra moral realism, the evaluative component of the proposition “*x* is morally wrong” should not be thought of as a description of some mind-independent moral property which functions in an evaluative way. Instead, it should be thought of as a direct *expression* of one’s attitude to *x* which functions in an evaluative way. A supervenient relation therefore holds between *x* and one’s attitude to *x*. As such, given the supervenient relation between *x* and ‘moral wrongness’ (*qua* one’s attitude to *x*), where *y* is identical to *x* then *y* should also be

⁴ For the sake of simplicity, I have discounted a morally neutral attitude.

morally wrong. In the case of x and $VR/E(x)$, however, these two properties are not identical. It does not follow from our discussion on supervenience that $VR/E(x)$ is morally wrong because x is morally wrong, although this does not negate $VR/E(x)$'s moral wrongness for another reason (Young, 2013). Moreover, in light of the irreducible nature of moral properties to non-moral (physical) properties, and the argument that moral utterances are not referring to mind-independent properties of the world but to one's attitude towards x or $VR/E(x)$, it seems that what is being conveyed by either statement – “ x is morally...” and “ $VR/E(x)$ is morally...” – is simply one's attitude. In the next section, I explore this idea further.

4. Expressivism

Expressivism (see Blackburn, 1985; Gibbard, 1990; Horgan & Timmons, 2006) argues that a sentence like “ x is morally wrong” is not describing some mind-independent property of x that we happened to have detected or otherwise stumbled upon (namely, its wrongness). Instead, it indicates our reaction to x which we express in the form of a negative attitude (in this case). As an aside, expressivism should be viewed as a continuation of/improvement on emotivism (Ayer, 1952; Stevenson, 1937) and the sentimentalist tradition of Hume (see Section 6) in which moral utterances are expressions of one's attitude (expressivism) not simply one's emotions (emotivism). It is therefore because the proposition ‘ x is morally wrong’ is not describing some independent fact about a representation or act within the world that it cannot be true or false. At best, if uttered sincerely, it simply expresses the subject's disapproval of x , much like declaring “Boo x !” Talk of moral properties should not, then, be talk of properties realized by certain non-moral (physical) properties (of the right kind). Instead, they should be thought of as the projection⁵ of one's attitude onto certain physical properties of the world.

⁵ For this reason, this approach is sometimes referred to as *projectivism*.

The truth of expressivism therefore entails that there is no moral reality independent of our attitudinal projections; but this fact (the negation of moral realism) does not *entail* expressivism (Ross, 2009). So why does it come about? Perhaps what expressivism reveals is a contingent psychological truth about us.

In the case of, say, a positive moral utterance, expressivism postulates that one is simply expressing a preference or an idea of what one wants or wishes to affirm, or perceives to be an “ok thing to do” and nothing more (Ross, 2009). To illustrate further, Ross invites us to:

Imagine an artist who feels a strong preference for abstract expressionism over figurative painting – this value utterly structures his life. But no one, including the artist himself, thinks of the sentence “Abstract expressionism is in all ways superior to figurative painting” as expressing a fact in the world he was lucky enough to detect. (*ibid.*, p.48)

Analogous to the claim of the artist, and what it does and does not mean, expressivism holds that a moral claim is a description articulating how one *feels* about x – one’s attitude towards x , whether positive or negative, which will likely shape one’s actions – and not a way of demonstrating one’s ability to detect some independent property of x .

If we accept that the proposition “ x is morally wrong” is equivalent to “ x is something I disapprove of” then, given the previous discussion, it is also equivalent to “I believe that x is something I disapprove of”. Yet, crucially, where moral wrongness or rightness is said to be nothing more than an expression of moral attitude, rather than part of the actual make-up of x , one’s *belief* (that x is something I disapprove of) becomes self-evident – and so justified – every time one sincerely and intelligently asserts “ x is morally wrong”.

When thought of in this way, the supervenience relation ceases to be a relation between two sets of real properties – one non-moral and the other moral – and so ceases to describe a *metaphysical* relation. Instead, it becomes a way of governing our moral attitudes and the moral language designed to express them. As Meyers (2012) explains:

Instead of saying that every token of the same natural type must *possess* the same moral *properties*, we can say that every token of the same natural type must be *ascribed* the same moral *predicates*. If one attaches a moral predicate (such as “is morally wrong”) to some token action, then one must attach that same moral predicate to any other token action of the same type. (p.19)

Even in the absence of truth-aptness, one’s moral attitude is still required to be consistent. Therefore, the predicates one ascribes must show this to be the case. To develop a point made in Section 3.2, if one has a negative attitude towards a token of x (x_{token1}) then one should also have a negative attitude towards a different token of x (x_{token2}) and this should accord with the predicates ascribed.

By ascribing the predicate “is morally wrong” to x , one is *expressing* a particular state of mind (an attitude) towards x . To say “I disapprove of x ”, on the other hand, is to *report* directly one’s attitude (the same attitude) expressed by “ x is morally wrong”. For expressivists, the propositions “ x is morally wrong” and “I disapprove of x ” amount to descriptions of the same state of mind; they differ only insofar as the former is an expression of that which the latter reports explicitly.

4.1 Expressivism and gamespace

If applied to VR/Es, expressivism does not appear to undermine the amoralist argument that there is no case to answer. If moral utterances are simply expressions of one's attitude towards VR/E(x) then such utterances have no truth-aptness and so cannot be said to be true or false. Thus, the proposition "virtual murder (or rape, paedophilia, etc.) is wrong" is equivalent to "Boo virtual murder (rape, paedophilia, etc.)" which itself carries as much moral weight as "Hurray virtual murder (rape, paedophilia, etc.)". In neither case, has one detected some independent property of moral rightness or wrongness with regard to VR/E(x).

But if expressivism is true then it is not only true for proposition "VR/E(x) is morally..." but also " x is morally..." In other words, it is not only true when expressing one's attitude about virtual representations and enactments but also when expressing one's attitude towards non-virtual representations and actions. If the claim that "murder is wrong" is equivalent to "Boo murder" which, in terms of moral weight, is on par with "Hurray murder" then even if everyone in the world intelligently and sincerely expresses the view that murder is wrong (they are of one mind on this matter, let us allow), given that virtual murder is not equivalent to actual murder, this does not entail that everyone will express the view that virtual murder is wrong (likewise with virtual rape, paedophilia, etc.). Moreover, in the scenario just described, under expressivism, it cannot be said to be *true* that murder is wrong, only that everyone is of the opinion that it is wrong. Where opinions differ (as is the case with virtual murder, it is fair to say), there seems to be no place within expressivism for claims that one view is right and the other view is wrong because neither is picking out some independent moral truth.

The expressivist may argue that this is indeed the case, and not just with VR/Es but with any act or representation that is subject to moral scrutiny. A consequence of this is that moral policing (prohibition and censorship) becomes a matter of consensus and majority rule

without legitimate reference to right and wrong beyond the ascription of predicates reflecting one's moral attitude (See Section 6 for a return to this point). In the case of moral violations mentioned previously (murder, rape, paedophilia), each of these can be morally prohibited (even legally so) and even be consistently referred to as things which are morally wrong without it being the case that it is *true* that they are moral wrongs insofar as the propositions pick out some property of wrongness that exists independently of one's negative attitude.

4.2 Problems with expressivism

We have seen that expressivism holds that the propositions “ x is morally wrong” and “I disapprove of x ” describe the same negative attitude (the same state of mind of the subject), only the former is said to express it while the latter reports it. The predicate “is morally wrong” is therefore merely a way of expressing subject S 's disapproval of x . The truth of x 's wrongness equates to (is nothing but) the truth of S 's disapproval of x . Given this, it seems that the following somewhat counter-intuitive position holds: If I did not disapprove of x then x would not be morally wrong. Conversely, if I approved of x then x would be morally right; or even, it is the fact that I approve of x that makes x morally right. Of course, what these statements actually articulate is the following: If I intelligently and sincerely report my attitude to x in a certain way (say, “I disapprove of x ”) then it follows that I would express this attitude in a manner congruent with the way I report it (that “ x is morally wrong”, for example). The relationship between approval/disapproval and rightness/wrongness seems reasonable if one is expressing one's attitude towards, say, a particular form of music (e.g., heavy metal, jazz or rap) or art (cubism or surrealism, perhaps) or some social convention such as queuing. Thus, given that I disapprove of queuing, it makes sense that I would express this disapproving attitude in the following way: “Queuing is wrong”.⁶ It also seems

⁶ Perhaps in the case of music or art, one would be more inclined to say bad rather than wrong.

reasonable to accept the following (or at least accept that it should be seen as less of a concern given what it is about): I express my attitude about queuing – that it is wrong – if and only if I disapprove of queuing. Here, one may acknowledge that there is a rational relation embedded within the statement (between moral utterance and attitude) irrespective of whether one shares the attitude towards queuing.

Consider, then, the following:

1. I express my attitude that murder is morally right if I approve of murder.

Irrespective of whether or not one approves of murder, one should attest to the coherence of this proposition: that it necessarily follows that *if* I approve of murder then I should assert that it is morally right.⁷ Now consider the following:

2. I express my attitude that murder is morally right if and only if I approve of murder.

Proposition (2) is, again, logically coherent. But do we want to allow, as (2) indicates, that the *only* reason I assert that murder is morally right is because I approve of it: to the point of claiming that one's approval is both necessary and sufficient for the claim that murder is morally right? If this is the case, then one need look no further than one's positive attitude to find a reason to claim that murder is a morally good thing to do. The next example, illustrates well the point I am making:

3. I express my attitude that murder is morally wrong if and only if I disapprove of murder.

⁷ For the sake of brevity, I am discounting that one's approval of murder co-occurs with one's amoralist view.

Do we really wish to accept that the *only* reason for me to express the proposition “murder is wrong” is that “I disapprove of it”? For some, this may seem counter-intuitive; yet this is precisely the position expressivism adopts. Perhaps what we are looking for, instead, is the idea that one’s attitude coincides with some independent moral truth and that this independent moral truth is the reason why I have the corresponding attitude I do. If this is so, then such a position is not to be found within expressivism.

4.3 Moral utterances as more than expressions of attitude

It also seems to be the case that we sometimes devise moral statements about x that do not express our attitude towards x . This is often referred to as the Frege-Geach problem for expressivism (see Geach, 1965). To illustrate: if it is morally wrong to murder then it follows that it is morally wrong to persuade one’s friend to murder. This statement makes a moral claim in its conclusion through its adherence to a logical structure (*modus ponens*) but it does so without expressing the attitude of the subject: for I can accept that the conclusion necessarily follows from the premise without agreeing with the claim that murder is morally wrong. In other words, I can accept that *if* it were the case that murder is morally wrong then it necessarily follows that the conclusion – it is morally wrong to persuade one’s friend to murder – is true. As it is irrational to agree with this premise while denying the conclusion, I can accept the necessity of the moral claim given the moral premise; but I am able to do all of this without any part of the aforementioned statement expressing my own attitude towards murder.

The Frege-Geach problem illustrates that moral utterances function in ways other than to express one’s attitude or even one’s belief about one’s attitude; rather, they constitute beliefs

about moral rights and wrongs that allow for logical coherence independent (and therefore irrespective) of one's attitude. In other words, such utterances are sufficiently independent of one's attitude to form compositionally stable content which can be embedded in logically complex contexts (Chrisman, 2011). As a means of overcoming this problem, let us consider a recent development within expressivism.

5. Ecumenical Expressivism

According to ecumenical expressivism (Ridge, 2006), when uttered, the proposition “*x* is wrong” reveals two interrelated facts about the mental states of the subject. The first concerns an attitude. The second relates to a particular belief which is said to make anaphoric reference to this attitude.⁸ Importantly, though, for proponents of ecumenical expressivism, the subject does not hold a negative attitude (in this case) towards a particular act. Instead, in the case of ‘*x* is wrong’, the subject disapproves of some property – call it *P* – and believes that *x* realizes *P*. Thus, in declaring that theft (which in descriptive terms equates to taking something without authorization/permission) is wrong, the subject holds a negative attitude towards *P* and believes that an act of theft (taking something without authorization/permission) realizes *P*. Moreover, in stating that theft is wrong, the subject is not (should not be) declaring only that *this* (token) act of stealing is wrong but, rather, that the *type* of action of which *this* particular act is a token – in virtue of realizing property *P* – is wrong. 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 theft (in virtue of property *P*) then one should also disapprove of kidnapping.

⁸ 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’ clearly refers to Fred.

Now, what counts as property P may differ from individual to individual. Some may view P in terms of negative utility (the realizing of more displeasure than pleasure, say; which is *a posteriori* discoverable), while others may hold it to be a violation of God's law, or characteristic of a failure in one's duty to others, or a vice rather than a virtue, and so on. Declaring "x is wrong" – where x equates to theft – reveals the following:

- (i) Subject S disapproves of P and believes that x realizes P (thus making anaphoric reference to that which S disapproves of).

Given that what counts as P can vary from subject to subject, in the case of negative utility we get:

- (ii) Subject S disapproves of P (where P equates to increased negative utility) and believes that x realizes P (thus making anaphoric reference to that which S disapproves of).

Or where P equates to a violation of God's law:

- (iii) Subject S disapproves of P (where P equates to violating God's law) and believes that x realizes P (thus making anaphoric reference to that which S disapproves of).

In accordance with a meta-ethical approach, ecumenical expressivism does not hold that x is morally wrong *because* it increases negative utility or violates God's law. It holds that what it means to say that x is morally wrong is that x (whatever x describes) realizes some property P (whatever P is) and the subject (a) disapproves of P and (b) believes that x realizes P . Ecumenical expressivism thus provides the framework for explaining the meaning of “ x is wrong” by explaining what it means for x to be wrong irrespective of what x is or the particular reason for one's disapproval of some property P that S believes x realizes.

One's disapproval of P is not truth-apt but whether or not x realizes P *is* truth-apt. The independent truth-aptness of P and its relation to x therefore allows moral utterances to express more than one's attitude and therefore form the stable content of logical structures irrespective of one's attitude, as noted by the Frege-Geach problem. Recall the following example: “If it is morally wrong to murder then it follows that it is morally wrong to persuade one's friend to murder”. This statement is problematic for traditional expressivism because moral wrongness is playing a role that is more than the expression of a negative attitude; it is adhering to a logical structure (modus ponens). However, if one holds the following: If x (murder) is wrong because it realizes P (a truth-apt condition), which *I happen to disapprove of*, then persuading one's friend to commit x (murder) would also realize P (which we know *I happen to disapprove of*) which entails that this is also wrong. In addition to one's attitude towards P , and therefore x in virtue of its act of realizing P , the independent truth-aptness of P (independent of one's attitude) allows moral utterances about x to form the content of logical structures. It is in virtue of x 's realizing P , which occurs both in the case of one's x -ing and persuading one's friend to x , that enables compliance with modus ponens. This, one can acknowledge irrespective of whether one approves of P or not.

One might object that all ecumenical expressivism does is replace moral claims about x with moral claims about P . In other words, proponents of ecumenical expressivism have simply added an extra component to the moral debate – namely, some property P – and so switched the focus away from the particular event x to a property (P) realized by x . If one disapproves of x in virtue of its realized property P then what remains unclear is why one should disapprove of P in the first place.⁹

We know that, in keeping with expressivism, ecumenical expressivism endorses the view that any moral claim about P is simply an expression of one's attitude. Therefore, the only truth-aptness regarding negative assertions about x is that the utterer disapproves of P and believes that P is realized by x . This disapproval finds expression in the form “ x is morally wrong”. All that ecumenical expressivism explains, then, is what the subject *means* when he/she says that x is morally wrong. How ecumenical expressivism differs from expressivism is that, with the inclusion of property P , the content of one's disapproval and with it the reason for one's disapproval (of x) can vary from subject to subject. Importantly, what ecumenical expressivism does not provide is a means of identifying the ‘correct’ (*qua* most justified) moral position regarding P and subsequently x , or indeed any objective justification at all for the disapproval of P or x .

It is not within the remit of ecumenical expressivism to justify the approval or disapproval of some property (P) realized by x . What ecumenical expressivism *entails* is the negation of moral realism; what it *suggests* is that established or alleged objective moral thinking is

⁹ Again, I thank the anonymous reviewer for drawing my attention to this point.

nothing but the elevation of one particular attitude over contrasting others. This may be the case; but this is not to deny *some form* of moral reality as I intend to show in Section 6: for if constructive ecumenical expressivism is to prove useful in the continuing debate over the moral status of video game content then it must offer more than a moral position consisting of a plethora of different attitudes with no way of differentiating between them in terms of moral weight. Before discussing how moral consensus might be achieved, however, let us first apply ecumenical expressivism to gamespace.

5.1 Ecumenical Expressivism and gamespace

Recall that the overarching aim of this paper is to explain the meaning of moral utterances of the kind “ x is morally wrong” and “VR/E(x) is morally wrong”. The extent to which some property (P) realized by either x or VR/E(x) is the same is the extent to which supervenience determines that one’s attitude to this property and ultimately to x and VR/E(x) should be the same. To illustrate, it may be that:

- (a) Subject S disapproves of P (where P equates to increased negative utility) and believes that x and VR/E(x) each realize P (thus making anaphoric reference to that which S disapproves of).

Of course, it may be that x and VR/(x) do not realize the same property P but in fact realize different properties – say p_1 and p_2 which S also disapproves of. Thus, it may be that:

(b) Subject S disapproves of p_1 and believes that x realizes p_1 (thus making anaphoric reference to that which S disapproves of).

And

(c) Subject S disapproves of p_2 and believes that $VR/E(x)$ realizes p_2 (thus making anaphoric reference to that which S disapproves of).

Why p_1 and p_2 are disapproved of may be for the same reason (say, increasing negative utility) or for different reasons (increased negative utility and violating God's law, respectively). It may of course be that one is approved of and the other not. After all, it would hardly strike the reader as incoherent that the vast majority of gamers who play violent video games have an approving attitude toward virtual murder but not actual murder. Neither is it contentious to say that, where disapproval of $VR/E(x)$ exists, it is likely to be as a result of its relationship to x (i.e., that it is virtually representing/enacting x) and the fact that, first and foremost, one disapproves of x . But this likelihood does not *entail* that if one disapproves of x one must also disapprove of $VR/E(x)$. After all, as we have established, the relation between x and $VR/E(x)$ is not one of identity. But, to reiterate, the *reason* for one's approval or disapproval is not the issue here; what is of concern is making sense of one's moral utterances.

6. Constructive sentimentalism

Prinz (2007) presents what he calls *constructive sentimentalism*. This is the view that although moral judgements stem from sentiment, or what Hume ([1739] 1978) refers to as feelings of disapprobation or approbation (recall the brief mention of emotivism in Section 4; see also Greene, Sommerville, Nystrom, Darley & Cohen 2001; Nichols & Mallon 2005), these sentiments nevertheless form the basis for rules which have their own objective status within the socially constructed space they occupy (Copp, 2011, likens this view to what he calls *realist expressivism*; see also Copp, 2001). As Prinz states: “Things that we construct or build come from us, but, once there, they are real entities that we perceive” (2007, p.168). Prinz’s position is compatible with the neo-sentimentalism proffered by Gibbard (1990), who argues that wrongful acts are judged to be so, not simply because one has a negative feeling towards the act but because such a feeling is appropriate. The addition of this normative element – that the negative feeling of guilt (for example) is not simply something we happen to feel but what we *should* feel – means, for Nichols (2008), that “even if one has lost any disposition to feel guilt about a certain action, one can still think that feeling guilty is *warranted*” (p.258; emphasis in original). Nichols goes on to argue that the emotions we feel in relation to a given action have helped shape our cultural norms, by determining which are sustained and which are lost. His *affective resonance hypothesis* essentially states that: “Norms that prohibit actions to which we are predisposed to be emotionally averse will enjoy enhanced cultural fitness over other norms” (*ibid.*, p.269).

For the new generation of sentimentalists like Prinz and Nichols, whether an action is deemed good or bad, or right or wrong, is in no small part dependent on the emotion we feel towards that action. However, these theorists have built on more traditional Humean sentimentalism by arguing for a degree of objectification within a given space. If at a given time I experience a sense of disapprobation towards an individual’s actions, this does not *as a general rule* –

based on how I feel *right now* – make my disapproval grounds for the judgement that what this person is doing is bad; nor does it make my disapproval a good thing, unless the disapproval is warranted.

7. Constructive ecumenical expressivism

In the same way that expressivism is seen as a progression from emotivism (feelings of disapproval in the form of disgust, for example) – by positing moral utterances as expressions of attitude rather than sentiment – it is my contention that one can borrow from neo-sentimentalists like Prinz and argue for constructive ecumenical expressivism. Such a position is not incompatible with constructive sentimentalism; it does however have a wider scope owing to the fact that it does not describe moral utterances exclusively in terms of sentiment but, rather, in terms of one's attitude: although I accept that this may in part be shaped by some form of association with sentiment. As such, it may be that irrespective of whether, as an individual, one approves or disapproves of x , if it is the case that x realizes property P then it may be that one *ought* to disapprove of x in virtue of this fact and the fact that it is an established cultural norm (objectification of attitudes towards P) that property P is disapproved of and therefore wrong.

Constructive ecumenical expressivism proffers the freedom for different subgroups to disapprove of P for different reasons, each objectified within their sub-culture. Subgroups 1-3 may each disapprove of x in virtue of x realizing P, but recognize P as being a different property in each case (as ecumenical expressivism allows). Group 1 may hold that P equates to increased negative utility, group 2 to the violation of God's law and group 3 to a vice rather than a virtue. In each case, a meta-ethical approach maintains that the sub-groups *mean* the

same thing when they utter “ x is morally wrong” despite having different reasons for holding the attitude they do.¹⁰

In the case of the moral utterance “VR/E(x) is morally wrong”, one’s disapproval of VR/E(x) may occur because one holds that x is morally wrong (to some, this may have strong intuitive appeal in cases of virtual rape or paedophilia). What has not been established in the case of VR/E(x), however, is any sense in which one’s disapproval is *warranted* and therefore has been established as a cultural or even sub-cultural norm. Has it been established that x and VR/E(x) realize the same property P which is morally proscribed? Or is it that, in light of the fact that $x \neq$ VR/E(x), each event realizes a different property and only the property realized by the former event has achieved some form of moral consensus? Therefore, while the structure of the two propositions may be the same from a meta-ethical point of view, in the absence of an established cultural norm, unlike “ x is morally wrong” (in virtue of realizing some property P) – where x equates to murder or rape or paedophilia – the proposition “VR/E(x) is morally wrong” (*qua* virtual murder/rape/paedophilia) does not reflect any established moral objectification within either the virtual gaming space to which the utterance is directed or the non-virtual social community looking on. From a meta-ethical perspective which endorses constructive ecumenical expressivism, the meaning of “ x is morally wrong” and “VR/E(x) is morally wrong” is equivalent. Where they depart is that, given the need for moral authority based on objectified moral norms rather than moral realism, the former utterance is deemed to be warranted whereas, at present, the latter is not; or at least cannot be said to be warranted owing to the lack of such norms with regard to VR/Es.

¹⁰ I am not discounting the possibility that an individual or group may recognize that x realizes each of these things: negative utility, violation of God’s law, vice. Nevertheless, it is likely that one will be held as a more fundamental reason for disapproval than the others.

In short, constructive ecumenical expressivism explains the consensus found regarding the immorality of murder, rape and paedophilia in terms of a shared negative attitude towards x (the act of murder, rape, paedophilia), not because of the act itself but some property (P) realized by the act (x). Different reasons may therefore be proffered for the negative attitude, but given that these different reasons are seen to be justifying the agreed moral condemnation of x , a state of objectification can be established whereby this condemnation is judged to be warranted. Because $VR/E(x) \neq x$, however, whatever shared condemnation of x exists does not necessitate the shared condemnation and therefore the objectification of a negative moral attitude towards $VR/E(x)$. It has not been established that the moral condemnation of $VR/E(x)$ – that is, virtual murder/rape/paedophilia – is warranted given that it has not been established that each event (x and $VR/E(x)$, respectively) realize the same property P. To reiterate, *a priori*, no argument to the contrary has been established (Luck, 2009; Young, 2013) and, *a posteriori*, evidence is at best mixed that $VR/E(x)$ results in some form of increased negative utility (Schulzke, 2010) and, in some cases, simply unavailable.

So if it is the case that utterances about the morality of gaming content are merely expressions of attitude that lack the same level of consensus and therefore objectification as that found in non-virtual examples, then how is constructive ecumenical expressivism to inform the debate over $VR/E(x)$ – say, the virtual torture and rape of a non-player character (NPC) – and whether it is morally problematic or whether, as the amoralist attests, it is just part of a game and therefore not a moral matter? Constructive ecumenical expressivism may help us understand the nature of moral utterances and why no moral consensus is forthcoming regarding VR/Es , but can it help us move towards a consensus, and what might that consensus

be? The present answer may be unsatisfactory. It is my contention that, *a priori*, there is no means of identifying which, if any, property P, realized by VR/E(x), is more justified as a measure of moral praise or condemnation. Elsewhere, I have argued that decisions about which gaming content (if any) should be prohibited, and which should not, should be based, not on any particular approach to morality but on the extent to which, psychologically, we can cope in our interacting with them (Young & Whitty, 2012). At present, given the absence of psychological research detailing how well we are able to cope with such symbolically violent and/or otherwise taboo interactions, all we can say for now is that if a consensus is to be arrived at about which content should be prohibited then it is likely that it will be based on a shared condemnation of VR/E(x) in light of some property P that VR/E(x) realizes. This does not mean or even require, however, that there is a consensus on what constitutes P, although I proffer that it should be based on the psychological effects on those engaged in VR/E(x); rather, it requires only that P (which will amount to different things for different people) is realized by VR/E(x): a position constitutive of constructive ecumenical expressivism.

As a final point, it may be objected that seeking *a posteriori* evidence for negative psychological effects is simply a form of moral consequentialism. To be clear, psychological evidence should be used to inform decisions about which content (if any) to prohibit, but any decision to prohibit should not be a *moral* decision; it should not be based on the view that it is morally wrong or bad to engage in V/E(x), only that evidence indicates that it is psychologically damaging (for some or many/all) to do so. Any putative moral claim based on psychological evidence is not picking out some independent moral truth. Instead, as constructive ecumenical expressivism explains, it is nothing more than an attitude expressing disapproval couched in moral terms. Denying moral realism does not negate the possibility that a form of moral consensus may emerge, as discussed above, and this may be based on

future psychological evidence (or should be, I contend). If such a moral consensus is forthcoming at some future time regarding VR/Es, then what I hope to have provided here is a description of what that moral consensus amounts to: in effect what it means to declare “VR/E(x) is morally wrong”.

8. Conclusion

It has been the aim of this paper to proffer a meta-ethical approach to propositions which make reference to virtual and non-virtual interactions in order that we might better understand what is meant by utterances which exclaim that x or VR/E(x) “is morally wrong (or right, good or bad)”. I have rejected moral realism and proposed instead a form of expressivism that provides a moral framework applicable to both virtual and non-virtual engagements. Constructive ecumenical expressivism posits a form of objectified moral reality based on social norms such that one’s expression of negative attitude towards x or VR/E(x) can be judged appropriate or not depending on what is believed to be warranted. In the case of paradigm examples of morally prohibited acts – such as murder, rape and paedophilia – social norms and therefore consensus hold these to be things one should morally disapprove of. Why this is the case may vary, but constructive ecumenical expressivism is able to incorporate such variation into its approach. In the case, of virtual representation and interactions, however, despite the meaning of “VR/E(x) is morally wrong” being equivalent to “ x is morally wrong” in terms of its meta-ethical structure, the lack of moral authority with regard to whether such an utterance is appropriate or not (that is, in keeping with what is warranted) can be explained, at present, by a lack of a prevailing social norm and majority consensus regarding those VR/Es that depict non-virtual, morally proscribed activities.

References

Ayer, A. J. (1952). *Language, Truth, and Logic*. New York: Dover Publications.

Blackburn, S. (1971). Moral realism. In J. Casey (Ed.), *Morality and moral reasoning* (pp.101-124). London: Methuen.

Blackburn, S. (1984). *Spreading the Word: Groundings in the Philosophy of Language*. Oxford: Oxford University Press.

Blackburn, S. (1985). Supervenience revisited. In I. Hacking (Ed.), *Exercises in analysis: Essays in honour of Casimir Lewy* (pp.47-67). Cambridge: Cambridge University Press.

Chappell, T. (2008). Moral Perception. *Philosophy*, 83, 421-437.

Chrisman M. (2011). Ethical Expressivism. In C. Miller (Ed.) *The Continuum Companion to Ethics*. London: Continuum.

Coons, C. (2011). How to prove that some acts are wrong (without using substantive moral premises) *Philosophical Studies*, 155, 83-98.

Copp, D. (2001). Realist-Expressivism: A Neglected Option for Moral Realism. *Social Philosophy and Policy*, 18, 1-43.

Copp, D. (2011). Jesse Prinz, *The Emotional Construction of Morals* (Oxford: Oxford University Press, 2007): Prinz's Subjectivist Moral Realism. *Noûs*, 45(3), 577-594.

Geach, P. (1965). Assertion. *The Philosophical Review*, 74(4), 449-465.

Gibbard, A. (1990). *Wise choices, apt feelings: A theory of normative judgment*. Cambridge: Harvard University Press.

Gibson, J.J. (1979). *The Ecological Approach to Visual Perception*. Boston: Houghton-Mifflin.

Greene, J.D., Sommerville, R.B., Nystrom, L.E., Darley, J.M., & Cohen, J. (2001). An fMRI investigation of emotional engagement in moral judgment. *Science* 293(5537), 2105-2108.

Hare, R.M. (1952). *The Language of Morals*. Oxford: Clarendon Press.

Horgan, T., & Timmons, M. (2006). Morality without Moral Facts In J. Dreier (Ed.), *Contemporary Debates in Moral Theory* (pp. 220-238). Oxford: Blackwell.

Hume, D. (1978). *A Treatise of Human Nature*, L.A. Selby-Bigge (Ed.), 2nd edn. Oxford: Oxford University Press. (Originally published 1739)

Luck, M. (2009). The Gamer's Dilemma: An analysis of the arguments for the moral distinction between virtual murder and virtual paedophilia. *Ethics and Information Technology*, 11, 31-36.

McNaughton, M. (1988). *Moral Vision*. Oxford: Blackwell.

Meyers, C. (2012). Expressivism, constructivism, and the supervenience of moral properties. *Philosophical Explorations*, 15(1), 17-31.

Nichols, S. (2008). Sentimentalism Naturalized. In W. Sinnott-Armstrong (Ed.) *Moral Psychology: The Evolution of Morality, Volume 2* (pp.255-274). Cambridge, Mass.: MIT Press.

Nicols, S. & Mallon, R. (2005). Moral Dilemmas and moral rules. *Cognition* 100(3): 530-542.

Prinz, J.J. (2007). *The emotional construction of morals*. Oxford: Oxford University Press.

Ridge, M. (2006). Ecumenical Expressivism: Finessing Frege. *Ethics*, 116(2), 302-336.

Ridge, M. (2007). Anti-reductionism and Supervenience. *Journal of Moral Philosophy*, 4(3), 330-348.

Ross, S. (2009). The End of Moral Realism? *Acta Analytica*, 24(1), 43-61.

Schulzke, M. (2010). Defending the morality of violent video games. *Ethics in Information Technology*, 12(2), 127-138.

Shafer-Landau, R. (2003). *Moral Realism: A Defence*. Oxford: OUP.

Sidgwick, H. (1981). *The Methods of Ethics*. Cambridge: Hackett Publishing Company.
(Originally published 1874)

Stevenson, C.L. (1937). The Emotive Meaning of Ethical Terms. *Mind*, 46(181), 14-31.

Young, G. (2013). *Ethics in the Virtual World: The morality and psychology of gaming*.
Durham: Acumen.

Young, G. & Whitty, M.T. (2012). *Transcending Taboos: A moral and psychological examination of cyberspace*. London: Routledge.