



# The conceptual exportation question: conceptual engineering and the normativity of virtual worlds

Thomas Montefiore<sup>1</sup> · Paul-Mikhail Catapang Podosky<sup>1</sup>

Accepted: 8 December 2023 / Published online: 12 January 2024  
© The Author(s) 2024

## Abstract

Debate over the normativity of virtual phenomena is now widespread in the philosophical literature, taking place in roughly two distinct but related camps. The first considers the relevant problems to be within the scope of applied ethics, where the general methodological program is to square the intuitive (im)permissibility of virtual wrongdoings with moral accounts that justify their (im)permissibility. The second camp approaches the normativity of virtual wrongdoings as a metaphysical debate. This is done by disambiguating the ‘virtual’ character of ‘virtual wrongdoings’. Doing so is supposed to provide illuminating ontological distinctions that inform ethical aspects of the debate. We argue that each approach faces its own set of issues, and as a result, motivates consideration of an alternative. The alternative we suggest turns inquiry concerning the normativity of virtual wrongdoings into a distinctively conceptual question. Rather than asking whether some action is right or wrong, or whether some virtual phenomenon counts as a particular action at all, we argue that research into the normativity of virtual wrongdoings may be guided by reflecting on whether a concept that originated and developed within a non-virtual context should be exported into a foreign virtual domain. We consider this approach and several objections.

**Keywords** Conceptual engineering · Conceptual ethics · Virtual action · Videogame ethics · Virtual ethics · Virtual worlds · Virtual reality · Applied ethics · Gamer’s dilemma

Debate over the normativity of virtual phenomena is now widespread in the philosophical literature, taking place in roughly two distinct but related camps. The first considers the relevant problems to be within the scope of applied ethics, where the general methodological program is to square the intuitive (im)permissibility of virtual depictions of immorality (‘virtual wrongdoings’ henceforth) with moral accounts that justify their (im)permissibility (see Luck, 2009a, 2009b, 2022; Bartel, 2012; Tillson, 2018; McCormick, 2001; Patridge, 2011, Wolfendale, 2007; Dunn, 2012). The second camp considers the first to have somewhat jumped the gun, and approaches the normativity of virtual wrongdoings as a metaphysical debate. Rather than granting a ‘virtual action’ to be a kind of action, the virtual kind, and then identifying what, if anything, would make it justifiably

immoral, it is argued that clarity must be reached concerning what kind of action a virtual wrongdoing is, if it is any kind of action at all. This is done by disambiguating the ‘virtual’ character of ‘virtual wrongdoings’. Doing so is supposed to provide illuminating ontological distinctions that inform ethical aspects of the debate (see Davnall, 2020; Heinrichs, 2020; Seddon, 2013; Brey, 2003).

Despite this shift in focus away from applied ethics, concentrating on the metaphysical status of virtual wrongdoings raises its own set of issues which appear unlikely to be resolved in the near future—and perhaps may not be resolvable at all. This presents a significant practical issue: if metaphysics or applied ethics cannot help us map the normative contours of particular virtual phenomena, then there are serious epistemic limitations regarding how we ought to respond to such phenomena, which includes not only the range of fitting reactive attitudes but also whether forms of social sanction or technological regulation are appropriate. To be clear, we still maintain that metaphysics and applied ethics can provide a certain level of clarity on the normativity of virtual wrongdoings, however we argue that alternative directions are also worthy of consideration.

✉ Thomas Montefiore  
Thomas.montefiore@mq.edu.au

✉ Paul-Mikhail Catapang Podosky  
paul.podosky@mq.edu.au

<sup>1</sup> Department of Philosophy, Macquarie University, Sydney, NSW, Australia

The alternative we suggest turns inquiry concerning the normativity of virtual wrongdoings into a distinctively *conceptual* question. Rather than asking whether some action is right or wrong, or whether some virtual phenomenon counts as a particular action at all, we argue that research into the normativity of virtual wrongdoings can be guided by reflecting on what we will call the *conceptual exportation question* (CEQ):

CEQ: Should a concept that originated and developed within a non-virtual context be exported into a foreign virtual domain?

For example, take the concept *murder*. So, we can ask: Should *murder*, which has its origins and development in a non-virtual context, be exported into the virtual world? We will want to impose some constraints on answering this question, such as narrowing the scope of its application to particular aspects of the virtual world (i.e., video games). Nevertheless, the question does appear to have traction in contrast to other ways of reflecting on the relationship between murder, virtual reality, and wrongdoing. We do not need to settle the question of whether a virtual ‘action’ counts as an action to entertain CEQ. It is instead asking a normative conceptual question, one that is sensitive to several different practical factors.

Taking a Nietzschean genealogical approach, as articulated by Matthieu Queloz (2022; forthcoming), we consider the needs that create and sustain concepts, and as a result, endow them with particular points. Because concepts emerge in response to domain-specific needs, and their point is to serve such needs, we argue that this is a useful way to reflect on CEQ. If a concept that originated and developed in non-virtual reality is exported into virtual reality, will it respond to the same needs? Will the concept maintain its point? Should our concepts be aiming to serve needs or preserve their points? Answering these questions gives us some idea of how we might want to answer CEQ, and highlights the pragmatic spirit of our approach, which is currently an under explored direction for investigating the normativity of virtual phenomena.

Before we move on, we want to highlight the novelty of this paper. First, we are advancing a new methodology for theoretical inquiry into the normativity of virtual wrongdoings, one that is conceptual rather than just within the scope of applied ethics or metaphysics. Second, this approach is a way of systematising relatively new philosophical research programs called *conceptual ethics* and *conceptual engineering*. Instead of asking whether any particular concept should be removed, revised, or replaced, we are introducing a methodological schema of conceptual ethics/engineering that applies to a whole domain of inquiry. This is done with the intention of adding new resources to the debate on virtual ethics, rather

than closing off other avenues through which this debate may be explored.

## The ethical approach

The favoured approach to investigating the ethics of virtual wrongdoings is, unsurprisingly, to consider it a debate in virtual ethics, indirectly responding to what has been called the ‘amoralist challenge’ for virtual actions (Ostritsch, 2017; Patridge, 2011; Young, 2017a, 2017b). The amoralist challenge appeals to the ‘unreal’ and ludological character of virtual depictions of immorality (primarily in video games) to argue for their moral harmlessness. In other words, the amoralist argues that concepts, such as ‘wrongness’, when applied to non-virtual wrongdoings like murder, fail to maintain a moral symmetry with virtual wrongdoings like virtual murder. The amoralist challenge therefore asks what morally relevant feature(s) of a virtual wrongdoing would justify its moral impermissibility.

The amoralist embraces the spirit of Huizinga’s conception of actions in games being shielded from moral criticism via a demarcating “magic circle” inside which moral norms do not apply (Huizinga, 1944). The notion of a magic circle has been defended, criticised, and revised in its application to virtual actions and specifically video game actions (being a subset of virtual actions) (see Brey, 1999; Juul, 2008; Salen & Zimmerman, 2003). However, the main thrust of Huizinga’s amoral conception of game actions remains somewhat applicable to video game actions and virtual actions generally, insofar as talk of virtual game worlds can sensibly track on to the intuitively forceful amoral aphorisms: “It’s only a game” and “It’s not real” as *prima facie* justifying one’s virtual conduct as morally permissible (see Formosa et al., forthcoming). For example, when one runs over an innocent pedestrian in the video game *GTA IV* (Rockstar Games, 2008) (being a canonical example in the literature (see Luck, 2009a, 2009b; Seddon, 2013)), while they may be enacting a ‘virtual murder’, the virtual wrongdoings moral wrongness is diffused, or at least complicated, by its being either part of a game and/or its being virtual.

Virtual wrongdoings being within a game make them likely to be playful (see Stenros, 2012; Seddon, 2013) or be consented to (in multiplayer games) via rules (Huizinga, 1944; Howe, 2008, p. 7; Young, 2013, 14), which arguably transforms their morally normative character (Hurd, 1996). Virtual wrongdoings being virtual make them unreal in respect to ‘real-world’ entities, and have been accordingly treated as analogous to how one appraises the existence of fictional actions, whereby it is only fictionally the case that a wrongdoing has occurred (see Mildenerger, 2017; Luck, 2022; Friend, 2008; Walton, 1990, p. 70; 2013; McDonnell & Wildman, 2019, p. 391). Further, the fictional character

of video game wrongdoings raises the issue of interpretive flexibility and the difficulty of treating a video game wrongdoing as really a depiction of wrongdoing. Is a virtual murder that occurs off-screen a virtual wrongdoing? (see Ali, 2015) Can a virtual wrongdoing be interpreted as morally objectionable regardless of its place within the context of the work? (see Patridge, 2011; Gribble, 1983; Seddon, 2012). For example, the literature on the paradoxes of horror (Carroll, 1990; Woodcock, 2013) and tragedy (Feagin, 1983), highlight the ambiguities over why one may wish to engage with different kinds of fictions, including fictional depictions of immorality.

The amoralist challenge is responded to by the moralist, who argues that the virtual and ludological elements of virtual wrongdoings cannot exhaustively insulate *all* moral wrongdoings from ethical criticism. For example, an action being playful or consented to via rules while perhaps somewhat morally insulating *some* virtual wrongdoings in some cases, does not morally insulate *all* virtual wrongdoings in *all* cases (Dunn, 2012; Wolfendale, 2007). As Dunn puts it, “a rule that says you can hit me in the face can make hitting me in the face no longer morally wrong. But it is consistent with this that mutual consent to rules that say you are allowed to eat me does not make it the case that eating me is morally acceptable” (2012, p. 11). Similarly, the fictionality of a virtual wrongdoing does not *necessarily* insulate it from moral criticism. Building off Bartel (2012) and Patridge (2011) for example, one may argue that *some* virtual wrongdoings such as the depiction of child sexual assault in the video game *RapeLay* (2006), may lack the interpretive flexibility to be treated as anything other than an instance of child pornography, which fictional or not, is morally wrong on the grounds that it constitutes pornography that eroticises inequality (Levy, 2002; Patridge, 2011), or simply in virtue of its being an instance of child pornography (Bartel, 2012).

However, if fictional and ludological features can justifiably insulate some moral wrongdoings from moral criticism, such as in the case of virtual murder, ethical dilemmas may arise. This is investigated by the Gamer’s Dilemma literature, which argues that while amoralism is intuitive insofar as it captures the unreality of virtual wrongdoings, and accordingly the intuitive moral permissibility of many virtual wrongdoings such as virtual murder, it also generates counterintuitive consequences by allowing for the moral permissibility of *all* virtual wrongdoings, such as intuitively morally repugnant virtual wrongdoings like virtual sexual assault. The Gamer’s Dilemma literature is devoted to either explaining what makes some virtual wrongdoings, like virtual sexual assault, morally impermissible, and others, like virtual murder, morally permissible, or grounding the source of our moral intuitions to other morally relevant factors such as the agency of the user (Ali, 2015), the realism of the virtual experience (Ramirez, 2020), or on other non-moral

grounds such as taste or convention (Montefiore & Formosa, 2022; Young, 2017b, 2019).

The Gamer’s Dilemma is not considered a legitimate dilemma in all cases. For example, Ramirez’ (2020) attempts to dissolve the dilemma by arguing that our intuitions regarding the moral permissibility of virtual actions do not come into conflict in highly realistic contexts. The conflicting intuitions of the dilemma do hold, however, in a contextually equal range of cases which fall short of this (Montefiore & Formosa, 2022) or in particular kinds of gaming contexts, such as in competitive shooter games (Nader, 2020; Luck, 2018), and therefore leaves an active, albeit limited Gamer’s Dilemma, and with it a generalisable ethics of virtual actions out of reach (see Formosa et al., 2023 for an empirical investigation of this claim).

While it is not the aim of this paper to demonstrate that no ethical approach is able to generate justified moral norms that can guide our behaviour in virtual worlds (as it certainly can in many cases), this section should raise scepticism regarding this approach being the sole avenue of further research and should motivate alternative additional approaches such as the conceptual approach we suggest in “[The conceptual approach](#)” section. In the next section, we will expand on the scope of disagreement that emerges when generalising virtual wrongdoings across all virtual worlds. This will be done by focusing on the metaphysical aspect of virtual wrongdoings that questions the nature of the ‘virtual’ and the ontological status of actions within virtual contexts. We will conclude that attempts at clarifying conceptions of virtual wrongdoings, while adding a needed nuance to the debate, insofar as they provide clarity as to what talk of virtual wrongdoings might amount to under different conditions, faces similar issues to the ethical approaches. That is, metaphysical approaches are heavily contested, aren’t easily generalised across virtual worlds, and face intuitively unappealing results.

## The metaphysical approach

Rather than approaching the ethics of virtual wrongdoings directly by asking what it is about certain virtual wrongdoings that would make them wrong, many consider the upstream task of determining the ontology of virtual wrongdoings to be a more central concern, which when clarified, may yield more promising ethical approaches (see Davnall, 2020; Heinrichs, 2020; Seddon, 2013; Brey, 2003). In other words, it is important that we are clear as to what virtual wrongdoings consist of if we wish to investigate their ethical character. Video game worlds, let alone virtual worlds, are highly diverse in character and like games, evade strict definition (see Aarseth et al., 2014; Seddon, 2013). As a result, the amoralist challenge may apply to *some* virtual worlds,

like in *some* video games, but may not get off the ground in spaces that are neither fictional nor part of a game, yet are nonetheless sensibly said to be virtual. A focus on the metaphysical status of actions in virtual worlds, as either ‘real’ actions that just happen to take place in virtual worlds, or ‘virtual’ actions that take place in virtual worlds, is therefore considered a central concern that some argue must preface any moral consideration of actions in these worlds.

Craft (2007), for example, considering the ethics of the infamous World of Warcraft funeral crash (see Luck, 2009b) writes:

We can only undertake a moral analysis of the Guiding Hand’s actions after we have understood the metaphysical nature of the environment in which they acted, and the nature of virtual worlds is a paradoxical one. Virtual worlds have both representational and actual elements, and the intermingling of the two is cause for ambiguity and misunderstanding. However, it is the dependent relationship of these elements that has allowed us to judge certain behaviours within virtual worlds as immoral. (Craft, 2007, p. 215)

Virtual intellectual property, such as a visual graphic image created entirely with digital assets via digital software, for example, appears to be ontologically equivalent to its non-virtual counterpart; both are visual graphic images which can be owned or valued. It may follow that digital intellectual property theft is morally equivalent to non-digital property theft insofar as they share a certain kind of ontological equivalence (see Brey, 2003; Craft, 2007). A similar kind of analysis seems available for other virtual wrongdoings, such as lying in multi-user non-game environments, insofar as the ontological equivalence of virtual and non-virtual lying (virtual lying and non-virtual lying are just both instances of lying) appears to entail a moral equivalence (see Brey, 2014; Mildenerger, 2017). In other words, wrongdoings may extend to virtual worlds beyond fictional and ludological spaces, and accordingly, amoralist arguments that defend the amorality of virtual wrongdoings on fictional or ludological grounds cannot be generalised across the gamut of virtual worlds. This is a kind of metaphysical approach to the ethics of virtual wrongdoings, as it makes a claim as to the ontological character of virtual entities and then draws ethical conclusions based on this ontological description.

Brey (2003, 2014) provides a metaphysical approach which illustrates an attempt to provide a morally relevant ontological distinction between virtual entities.<sup>1</sup> This

<sup>1</sup> ‘Virtual entities’ is being used here to indicate both virtual actions and objects for simplicity. Brey (2014) distinguishes virtual actions from virtual objects insofar as virtual actions are defined over virtual objects (virtual actions necessarily involve virtual objects and not vice versa) and can produce intravirtual and extravirtual effects. A virtual actions intravirtual effects being effects that solely take place within a virtual world and extravirtual effects being effects which take

analysis, it is hoped, will account for why some virtual wrongdoings, such as lying and theft, are sometimes morally equivalent to non-virtual wrongdoings, while others, such as virtual murder, are not. Let’s examine this approach to see the value metaphysical accounts can offer, as well as the potential issues. Brey distinguishes between simulated virtual entities and ontologically reproduced virtual entities, or ‘ontological reproductions.’ Ontological reproductions, according to Brey, are analogous to Searle’s (1995) ontological category of institutional social entities. What makes institutional social entities what they are, are *status functions* authorised by those who engage with them and their capacity to be recognised as the entities that perform these status functions. A pawn in chess, for example, is a pawn as long as it is recognisable as a pawn by those engaging with it, and performs the function that pawns perform in chess (moving forward one space at a time and so on). A pawn could therefore be ‘transmedial’ (Juul, 2005, p. 49), insofar as the concept of a pawn could be moved into a virtual world—an online chess game—and still count as a pawn because it is able to be recognised and perform its status function virtually. Similarly, lying and theft may be ontological reproductions, and accordingly transmedial, as telling a lie in a non-game virtual world, say over tele-conferencing software like zoom, would still clearly constitute a lie; the virtual medium does not bear on the nature of the lie.

The ontology of simulated virtual entities, on the other hand, is not only determined by status functions but also the capacity of their real-world counterparts to fulfil their status functions. Searle refers to these entities as ‘ordinary social entities.’ A hammer, for example, cannot be a donut by being recognised and treated as such, because a hammer is unable to perform donut-like functions. For the same reason, argues Brey, certain entities are not transmedial to virtual worlds; they cannot perform their functions virtually. A virtual tennis video game for example is not an instance of tennis because one is not actually playing tennis. Instead, the virtual tennis entities (tennis balls, the tennis net and so on) are simulations of real-world tennis entities which perform different status functions which only simulate tennis. This is why Federer, for example, would be able to play tennis comparatively well on any tennis court besides a virtual one, because playing tennis is engaging in a different activity to playing tennis, in a way that virtual chess is not a different activity to playing a physical board game instance of chess.

Footnote 1 (continued)

place externally to the virtual environment, in the real-world. Virtual actions described in reference to extravirtual effects are, for Brey, real actions, whereas virtual actions (possibly the same actions) described in terms of their intravirtual effects, are not (Brey, 2014).

Similar to the case of tennis, virtual murder may be a simulated virtual entity. When one intentionally runs over an innocent NPC (non-player character) pedestrian in *GTA IV* (2008), killing them, their action is just a simulation of murder, rather than an ontological reproduction. Brey argues that this demarcation between a virtual entity being a simulation or an ontological reproduction will determine the entities 'reality status,' and it is this distinction that will bear on the ethics of a virtual action. Brey, however, argues that determining the reality status of a virtual entity may be an unclear process as it is an analysis that is malleable by being determined by a virtual community and how they treat entities within virtual worlds (Brey, 2003, p. 281).

Brey's analysis provides some clarity to the ethical landscape. For example, virtual wrongdoings which are ontological reproductions are often also intuitively morally impermissible, and simulative virtual wrongdoings are often also intuitively morally permissible. For example, theft in an appropriately engaged virtual role-playing game is simulative and intuitively morally permissible (see Seddon, 2013; Craft, 2007; Mildenerger, 2017) whereas lying over a zoom call is an ontological reproduction and intuitively morally impermissible. However, Brey's analysis is not without its issues. We want to bring out three. First, distinguishing between simulations and ontological reproductions cannot exhaustively establish a morally relevant distinction for the ethics of virtual entities, at least not across all token instances of particular action types. For example, following Brey's framework it appears as though many simulative entities which are intuitively impermissible, will be rendered morally permissible to the extent that they lack a reality status. Many cases of virtual sexual assault, for example, while perhaps not an instance of sexual assault *qua* ontological reproduction, and therefore simulative, may still be plausibly intuited as morally impermissible (see Luck, 2009a, 2009b).

Second, with the emergence of mixed-reality technologies it is not clear that a binary distinction can be drawn between simulated entities and ontologically reproduced entities. To illustrate by expanding on the above tennis video game example, while Federer may be performing an activity distinct to tennis when playing virtual tennis, it is not clear he is performing a distinct activity to virtual tennis when it is expressed via a virtual reality simulator where he can hit an actual tennis ball at a screen, for example. To argue that virtual activities are fundamentally distinct from their real-world counterparts is to make the error of seeing virtual technology as static and only expressible via current technological hardware capacities, such as through two-dimensional monitors, controllers and so on. As Aarseth points out, in reference to the difficulty of formally defining video game media via their material base (but can be extended for our purposes to include virtual entities generally), that "the rapid evolution in games and game technology makes our

assumptions about their media formats a highly unreliable factor to base a theory on. We simply cannot assume that the parameters of interface, medium structure, and use will provide a materially stable base for our observations, the way the codex paperback has remained the material frame for students of literature for more than five hundred years." (Aarseth et al., 2014, p. 487) Further, if we were to grant simulative entities to be clearly distinct from ontological reproductions, the wide variety of virtual worlds makes the character of simulative entities difficult to generalise. Again in reference to games but for our purposes extended to virtual entities, Aarseth observes that: "(e)ven within the narrower domain of games in virtual environments, there are tens, maybe hundreds, of thousands of games that are somehow formally different from each other. A game such as Tetris (Alexej Pajitnov, 1985) has almost nothing in common with World of Warcraft (Blizzard Entertainment, 2004), or with Super Mario Sunshine (Nintendo, 2002)." (Aarseth et al., 2014, Ibid).

Finally, building off Davnall (2020), it is unclear whether virtual actions, simulation or ontological reproduction, are kinds of actions at all, and therefore ethically evaluable on the same grounds as non-virtual actions. Davnall argues that one can either inflate the metaphysical status of virtual actions to a kind of action, calling this an 'inflationary approach', or deflate their metaphysical status to images, calling this a 'deflationary approach'. However, either metaphysical approach will lead to its own set of issues, where a deflationary approach fails to capture the interactive nature of virtual worlds, and the inflationary approach fails to capture the representational nature of virtual worlds (Davnall, 2020, p. 226). Instead, according to Davnall, virtual actions are closer to a kind of performance than a genuine action or mere representation.

Alternatively, rather than holding the reality status of virtual wrongdoings to be the key factor in determining their ethical salience, one might instead argue that all virtual wrongdoings are real on the grounds that they are ultimately realised by a physically digital base and therefore robustly exist. This approach, known as virtual digitalism (see Chalmers, 2018, 2022) depends on accepting an "it-from-bit-from-it" thesis, where virtual entities (its) are realised via computational processes (bits) which are themselves realised by (encoded onto) digital hardware like computer chips and so on (its) (See McDonnell & Wildman, 2019 for a response to this view). The ethics of virtual wrongdoings will then not come down to whether virtual wrongdoings are real or virtual, as they are all both real *and* virtual, but by assessing who is engaging with them. For those who exist and have always existed in virtual worlds, virtual wrongdoings would simply be wrongdoings. Chalmers appeals to this intuition via the simulation hypothesis, whereby 'real-world' wrongdoings may turn out to be simulated if we in-fact

unknowingly live in a virtual world, and would remain ethically salient if we were to discover this fact (Chalmers, 2022). However, this line of argument is not applicable to the set of virtual wrongdoings that we are considering, as we are considering the ethics of virtual wrongdoings *which we take to be virtual*. Therefore, the question can be recast back to Brey's initial framing (and the problems that it poses) being how we ought to engage in virtual wrongdoings which we take to exist virtually.

To summarise, while we can perhaps identify particular instances of ontological equivalence across virtual and non-virtual worlds, which might shed some light on the moral contours of virtual phenomena, these attempts have so far failed to be applicable to virtual actions generally. This may be because 'virtual' like 'game' is a difficult ontological concept to pin down, where virtual worlds may only have shared family resemblances, rather than a single set of features that is common to all (see Aarseth et al., 2014). This is in part due to the diversity of virtual worlds and modes of engagement within those worlds, from offline single-player video game worlds to online multi-user non-game social worlds (See Heinrichs, 2020), as well as the degree and scope of virtual affordances and rules within those worlds, from the ludic constraint on one's agency when choosing to eat a ghost in pac-man, to choosing to crash a funeral in World of Warcraft (see Luck, 2009b; Ali, 2015). Further, the wide range of degrees of abstraction, context realism and perspective fidelity that realise these worlds for the user, from arcade style console video games to highly immersive virtual reality technology simulating deepfaked persons (see Ramirez, 2020; Ohman, 2020). Each of these considerations bears on the meaning of one's virtual actions and shapes the way in which they can be morally evaluated. Further still, it is not clear that one ought to even consider virtual actions as a kind of action at all, or as mere representational images (Davnull, 2020). This section should therefore, as with the previous section, while not seeking to challenge metaphysical approaches in identifying the character of virtual entities, provoke a sense of scepticism towards the degree of clarity that metaphysical approaches can offer towards the normativity of virtual actions.

## The conceptual approach

So far we have made explicit grounds for scepticism about ethical and metaphysical approaches in being able to produce frameworks for navigating the normativity of virtual phenomena. Each domain of analysis appears to fall short of offering a general explanation of what constitutes permissible behaviour across the gamut of possible virtual worlds. This presents serious practical issues. Virtual phenomena exist, being meaningfully incorporated into our lives, and

will continue to evolve with advancements in technology and in response to demand. Therefore, there are present day problems that must be addressed going forward. It is imperative that we get clear on the normative contours of virtual phenomena so that we can think more productively about how it ought to be responded to, how it ought to be shaped, and how it ought to be regulated (or whether it ought to be regulated at all).

For the remainder of the paper, we will outline a new methodological approach to reflecting on the normativity of virtual wrongdoings. Rather than asking, say, whether virtual murder *is* murder, or assuming that it is and explicating its relevant moral similarities, our account shifts the focus in a more pragmatic way. We ask whether virtual murder *should be thought of* as murder. This move to thought, rather than a strict metaphysical or ethical analysis, such as ontological or moral equivalence, allows us to gain better purchase on normative questions about the virtual world. That is, practical guidance can be found in answering the latter question about thought *even if* the former questions about metaphysics and ethics are unanswerable.

We intend this new focus on thought to be distinctly *conceptual*. We are not investigating the metaphysics and/or ethics of action, at least not directly. Instead, we are simply asking whether our 'real-world' concepts, that is, those concepts that have their origins and development within non-virtual contexts, such as *murder*, should be exported into the virtual realm. Rephrased as a question: Should wrong-tracking concepts with histories of development and articulation in the non-virtual world be extended to apply to aspects of the virtual world? What this highlights is that we are specifically concerned with the normativity of concepts, or to be more precise, we are aiming to derive answers about the normative contours of virtual phenomena by asking whether our non-virtual concepts should expand their range of application. This approach is characteristic of the fast-growing areas of philosophy called *conceptual ethics* and *conceptual engineering*. Before we lay out the details of our approach, we will give an explanation of these areas of philosophical inquiry and why they have become a popular methodological program.

## What is conceptual ethics and conceptual engineering?

In brief, conceptual ethics involves the critical evaluation of concepts to determine whether they ought to be removed, revised, or replaced (Burgess & Plunkett, 2013a, 2013b; Cappelen & Plunkett, 2020). The aim is, ideally, to discern which concepts we have most reason to use, or perhaps less ambitiously, though equally important, to identify the concepts we have good reason to abandon. Once a judgement is made as to whether there is need for a concept to

be reconsidered, or indeed, for a new concept to be introduced to plug-in gaps within an existing set of hermeneutical resources, then conceptual engineering can step onto the plate. This form of analysis involves designing conceptual innovations, and discerning implementation strategies to make them practically worthwhile, relative to one's aims and goals (Cappelen, 2018; Burgess et al., 2020).

Exactly what's involved in the proper methodology of conceptual ethics and conceptual engineering is contested. Questions that occupy philosophers in these fields tend to be the descriptive limits of topic-continuity, the bounds of normativity, the nature of concepts, and to what extent implementation is possible. We needn't get into the weeds of these debates. Suffice it to say, there is a lot of optimism that conceptual ethics and conceptual engineering can deliver meaningful conclusions about many pressing social problems. Gender, race, disability, sexual orientation, and a variety of other contemporary focal points of political contestation have been put through this mill of analysis, asking whether the operational concepts in these discourses are indeed the right or best ones, and if not, which are the right or best ones to use. This is most clearly seen in recent feminist philosophy that considers whether liberation projects would be better served with a trans inclusive concept of woman (Barnes, 2020; Haslanger, 2000a, 2000b; Jenkins, 2016), and in debates in the philosophy of race about whether the concept of race is in fact empty, in roughly the same way as the concept of witch, and therefore ought to be jettisoned for practical, theoretical, and political reasons (Appiah, 1996; Hochman, 2017; Zack, 1993).

Whole debates within philosophy now centre on which concepts we have most reason to use, and whether certain concepts should be used at all. We don't think it's premature to say that conceptual ethics and conceptual engineering are now a part of the staple diet of philosophers; they are primary methodological approaches to arriving productively at answers to core philosophical questions, especially contemporary social and political problems.

The question that is perhaps most pressing is this: how do we know which concepts we should be using? There are many thoughts on this, but one is particularly relevant to us. We will take a *genealogical* approach (see Plunkett (2016), Catarina Dutilh Novaes (2020), and Amie Thomasson (2020)). Specifically, we intend to shift the typical perspective on normative theorising about virtual phenomena to one that engages a particular question: Should a concept that has its *origins* and *development* within non-virtual reality be exported into the virtual world? We call this the *conceptual exportation question* (CEQ). Here our concern is whether the genealogical history behind a concept, which exists as a response to domain-specific needs, should unfold into foreign contexts. Our goal is not to provide a universal answer to this question. Rather it is to make the point that we

cannot take for granted that for any given concept that has its origins and development in one context (i.e., non-virtual reality), it will be rightly and appropriately extended into another (i.e., virtual reality). That is, whether a non-virtual concept should expand its scope of application into virtual domains must be examined on a case-by-case basis. Before we draw these points out in more detail, we'll look at the relationship between the genealogy of concepts and conceptual ethics/engineering.

## Genealogical approach to conceptual engineering

A still emerging area of conceptual ethics and conceptual engineering is one that is particularly attentive to the genealogy of concepts. Roughly, genealogy is a "developmental narrative describing how a cultural phenomenon... could have come about", such as a concept (Queloz, 2022, p. 435). And genealogies have different aims depending on the kind of genealogy it is. Some are vindicatory (e.g., Williams, 1985), some are debunking (e.g., Nietzsche, 1887); some problematise (e.g., Foucault, 1961), some expand our understanding of what's possible moving forward (e.g., Lorenzini, 2020). Whatever the case, there is a sense in which each project is unified by exploring the origins and development of a cultural product, and this exploration is sensitive to the historical needs of actual or imagined communities (e.g., a state of nature).

Matthieu Queloz, a theorist who takes a genealogical approach to concepts, seeks to make sense of Nietzsche's version of conceptual ethics, though Nietzsche himself never used that descriptor. According to Queloz, Nietzsche's conceptual ethics involves evaluating concepts by 'working backward', that is, by examining first from the needs of concept-users to which the concept answers, and then to the circumstances and character of those who use them. What needs are concepts responsive to?

If we take the need that immediately underlies a concept to be an *instrumental* need... this need must itself be understood as a product of the way the *characteristics* of certain concept-users—their inner needs, their drives and affects, their strengths and weaknesses, their abilities and limitations—combine with their *circumstances*—their natural and geographical environment, their material and technological affordances, their social structures and institutions, their position in society—to render that concept, or some broader class of concepts of which it is an instance, *needful*. (Queloz forthcoming, 14)

So, concepts have histories of originating and developing in response to instrumental needs; concepts have *needfulness conditions*. When concepts meet such conditions, then they acquire a *point*. In other words, "[e]very concept... comes

with extraconceptual presuppositions that have to be realised for the concept to be *pointful* by meeting an instrumental end” (Queloz forthcoming, 14). This is all very abstract, so it’s worth exploring an example. Edward Craig, for instance, tells us that the concept *knowledge*, or *know-that* (in contrast to *know-how*), developed in response to the needs of humans in a state of nature to ‘flag approved sources of information’ through a socially-determined process of establishing a set of reliable signals that express trustworthiness and competence over a domain of facts (Craig, 1990, p. 11). That is, *knowledge* originated and developed in response to the need to identify dependable testifiers; and identifying dependable testifiers is the point of the concept *knowledge*.

Whether one believes this story about *knowledge* is neither here nor there. The point that we want to draw out is that concepts are contingent things, they are socially constructed insofar as they are created in response to the instrumental needs brought about by a situation or environment of particular concept-users. Once a concept meets an instrumental need, it can be ‘pointfully’ used, as Queloz puts it. That is, the concept has a point; it is *pointful* because it is instrumentally needful. So concepts have both needfulness conditions and pointfulness conditions: the former concerns the instrumental needs the concept presupposes, and the latter concerns the point of the concept in serving such needs. The needfulness of a concept is tied intimately with the *right concept-users* and the *right circumstances*.

How does this relate to conceptual ethics and conceptual engineering? For Nietzsche, concepts are evaluated by asking for whom they have a point, and “he approaches that question through a backward inference from a concept via the instrumental need it fills to the conditions that engender that need and thereby render the concept *pointful*.” (Queloz forthcoming, 25) And, importantly, the pointfulness of a concept is a matter of how it serves the concerns of concept-users through its effects—how it impacts the lives of those users who live by the concept.

After evaluating a concept, and discerning for whom it has a point, we can start to see how genealogy can be useful at the engineering stage, that is, the stage at which we reflect on how we might want to move forward. Because genealogy reveals a concept to serve a need, we can reflect on whether the need itself could be responded to in a way that is better or more up for the task—say, with a revised concept, or with a totally new concept, or by eliminating the concept, etc. Put differently, conceptual genealogy “tells us what work the concept can do for us, and this covers not just the work it *already* does for us insofar as it functions well, but also the work it *could* do for us if it functioned better...” (Queloz, 2022, p. 446). Importantly, genealogy allows us to reflect on what we want our concepts to do for us, with an understanding of how they have been serving our needs in the past. We may not always come to the position that we have the

right concepts, even if we revise them. This may call for a complete conceptual innovation, a new concept to respond to novel circumstances and unique challenges.

### The conceptual exportation question

There are two things we want to draw out. When it comes to conceptual ethics and conceptual engineering, we need to look at both the *needs* that concepts serve, and what make concepts *pointful*. And discerning this requires looking at the character and circumstances of concept-users. After all, “a given concept is only worth using if certain presuppositions are fulfilled: concept-users must have certain inner needs, pursue certain concerns, and their circumstances have to be propitious to meeting these inner needs using this concept” (Queloz, 2022, p. 14). Once we have discerned this, we can ask certain questions: are our concepts serving these needs well? Would we be better off with new conceptual resources to meet these needs? This approach, while not the only way to consider the character of virtual phenomena, is novel and one that we argue is helpful to explore further.

We are interested in virtual phenomena, and so we’ll now restrict our discussion to this domain. As we’ve iterated through the paper so far, our aim is to offer an additional perspective on normative theorising about virtual phenomena that is strictly about concepts. We want to adapt the Nietzschean approach to conceptual ethics, and its implications for arriving at answers in conceptual engineering, to make it fit for dealing with questions about how to think productively about the normative contours of virtual worlds.

To this end, we offer the *conceptual exportation question* (CEQ) as a guiding methodological query:

CEQ: Should a concept that originated and developed within a non-virtual context be exported into a foreign virtual domain?

The significance of this question, when looked at through Nietzschean conceptual ethics, is that it prompts us to consider what the needs were, and presently are, that a particular concept emerged in response to, and therefore the needs that have given this concept a specific point. Moreover, the question allows us to reflect on whether *those same needs are present in virtual contexts* and whether *the concept, if exported, would have the same point*.

Again, this is all very abstract, so let’s explore concrete examples. Filling in CEQ, we might ask: should *murder* be exported into a foreign virtual domain? Taking lessons from Nietzschean conceptual ethics, we can ask the question: are the needs that produced and sustained the concept *murder* in non-virtual contexts, and thus gave life to its point, the *same needs that are present in virtual situations that appear relevantly similar*? That is, could the concept *murder* have the same point in both non-virtual and virtual reality? We might



ask the same for a host of other normative concepts: sexual assault, gaslighting, theft, lying, etc. We are not willing to settle on answers to these questions here, for our goal is just to introduce a new methodological approach to questions about the normativity of virtual phenomena. Nevertheless, we will say a little bit about how we might go about answering such questions.

In the first place, we need to reflect on the *expressive character* of a concept (Queloz, forthcoming). What is the concept expressing in terms of values, beliefs, commitments, desires, etc.? And are these values, beliefs, commitments, desires, etc. *the kinds of things that would be relevantly related to situations that, if the concept were exported, the concept would then apply?* In other words, are the things we would want to call murder (or sexual assault, gaslighting, theft, lying, etc.) in particular virtual worlds appropriately related to the expressive nature of the concept as it is used in non-virtual reality? This needn't be strictly related to the meaning of words, or the literal content of a concept, and instead the psychological associations which govern practices of applying the concept or term (i.e., lexical effects). Second, we need to reflect on the effects of a concept (Queloz, forthcoming; see also Cappelen, 2018; Marques, 2020; Podosky, 2022). That is, we must look at the consequences of living by the concept, and how it shapes the reality of concept-users. So, we might ponder whether the relevant effects of living by the concept *murder* in non-virtual reality would be similar to the effects of living by that concept in virtual reality, if it were to be exported.

If we take for granted that the concept *murder* expresses a desire for retribution or justice, and by consequence, typically triggering processes of judicial review, then we might say that the concept has been developed to express such sentiments and produce certain outcomes. This is to say that under the right conditions, the concept “would tend to have effects conducive to the satisfaction of those concerns.” (Queloz forthcoming, 18) Perhaps *murder's* expression for retribution is apt in some multi-user game environments (see Johansson, 2009), but justice, particularly legal justice, appears to be irrelevant in other virtual environments—e.g., in single-player simulation games like *GTA IV* (Rockstar, 2008).<sup>2</sup>

What this example is intended to show is that we need to bring into focus concept-users and their circumstances to fully appreciate the need for a concept like *murder*. Then we need to ask what the character of concept-users are like in virtual worlds and whether the circumstance in which they

are situated produces the same needs. After this, we can ask whether the concept *murder*, given its point, can serve the same point of meeting the relevant needs of concept-users, in virtual worlds. This approach, while more complicated than simply holding virtual acts to be the virtual instantiations of their real-world counterparts, highlights the need to consider the relevantly different concept-users in, and features of, virtual reality (across its varying contexts) that would qualify or disqualify a concept for exportation.

This still leaves a curiosity. Suppose that we find that a concept that has its origins and development in non-virtual contexts can indeed serve the same point in virtual worlds. Does this therefore offer an answer to CEQ? This, we take it, is an important site for philosophical investigation.

Some might be inclined to say, ‘yes’. If the needs of a given situation in virtual and non-virtual reality are the same, and the relevant concept can meet these needs, and maintain its point, then one might be satisfied with a positive answer to CEQ: the relevant concept should be exported. Call someone who would give this type of answer *point-focused*. And conversely, if the needs are not the same, even if a relevant concept could, in principle, meet these needs, the point-focused individual would offer a negative answer to CEQ: the relevant concept should not be exported. This is because the point of the concept would differ in virtue of meeting different needs. For example, if the point of the concept *murder* cannot be preserved in virtual environments once exported, then no matter if it happens to serve some other set of needs, *murder* should keep its range of application to non-virtual reality. In sum: point matters, not need.

In contrast to those who are point-focused, there are *need-focused* individuals. Instead of preserving the point of a concept, the need-focused individual is just concerned with whether a concept can meet the needs of a community in a given situation. That is, if a relevant concept serves one set of needs in non-virtual reality, and could, in principle, serve a different set of needs for users in virtual environments, and therefore acquire a new point, then this means that the concept should be exported. This need-focused view is pluralist about points so long as the concept can serve a particular range of needs. For example, if the concept *murder* can serve a different set of needs than those it serves in non-virtual reality, then it should be exported, *even if* it acquires a new point. In sum: need matters, not point.

In the middle are a range of views. Some might think that point-preservation only offers pro-tanto reason to export, and conversely, point-disruption only offers pro-tanto reason not to export. Similarly, one might think that need-satisfaction is pro-tanto reason to export, though this might be weighed against the extent to which a point is disrupted (or ‘pluralised’). And even beyond this, some will want to weigh points and needs against a whole class of other considerations, such as whether the exportation is justice-promoting, whether it

<sup>2</sup> A similar point can be made for legally relevant concepts like *sexual assault*, but perhaps not for non-legal concepts such as *lying*, where retribution or justice is less about carrying out legal processes rather than social ones, such as shaming or soliciting forgiveness.

contributes to the flourishing of those people would live by an exported concept, and whether it's more practically feasible to simply introduce a brand new concept into virtual reality in order to respond to some set of needs, rather than export an already existing non-virtual concept (e.g., it might be more practical to introduce a new wrong-tracking concept that captures cases that appear relevantly similar to murder, rather than exporting the concept *murder*).<sup>3</sup>

We want to stress that we accept that virtual worlds present concept-users with certain needs. The relevant point, however, is whether the needs of concept-users who participate in virtual worlds are radically different to the needs of concept-users in non-virtual reality—and whether this matters at all. We have some tendency to lean in favour of non-exportation in cases where points might be disrupted or pluralised.<sup>4</sup> Given that concepts are responsive to needs, and thus acquire points, it might simply be better to create concepts that serve the needs of concept-users in virtual space, and thus for such concepts to acquire their own unique points that are different and distinguishable without ambiguity from the concepts that serve the needs of concept-users in non-virtual space. Why?

First, there is the problem of *exportation feedback*. If a non-virtual concept is exported into virtual reality, the concept could acquire new psychological associations in its virtual context which might then bleed back into the non-virtual domain. As a potentially unrealistic example, our understanding of the moral significance of *murder* might be compromised due to it developing less serious psychological associations given its application to video games which involve highly realistic depictions of murder. We might

come to feel that murder is less tragic than we ought to. Importantly, this compromises the quality of the concept *murder* in its ability to meet our needs in non-virtual contexts, and because of this, it might change the point of the concept from what it once was.<sup>5</sup> The point here is not to say that this *will* happen. Rather the claim is that it is a potential risk that we must take seriously, and perhaps the best way of dealing with it is a precautionary heuristic: *only* export if there is very good reason to believe that it won't produce harmful or negative exportation feedback.

Second, and expanding on the precautionary heuristic, even if we cannot discern the extent to which exportation might produce harmful or negative feedback, *we can always ask CEQ*. Just because, at one time, we decide not to export a concept, it does not mean that at a future time we will fall on the same judgement. Over time, we might have a better grasp of the normative landscape, the risk of exportation feedback, and a better understanding of our needs, in both virtual and non-virtual space, and this might lead us to the point where we feel much more comfortable with the expanding application of a concept. So, we should not rush into exportation, and accept that it is always on the table. The problem, however, is that it's incredibly difficult to un-ring the bell, as it were. Once we've exported a concept, we might not be able to take it back if it catches on. This is not to say that it's impossible, say through a concerted effort to place a ban on the concept's usage in virtual worlds.

## Objections

This new methodological approach might be unconvincing to some, so we'll take time to work through possible objections. We set up this paper to say that existing approaches to the normativity of virtual wrongdoings cannot always offer practical guidance, perhaps in principle. This might be due to concerns about how we might settle on whether a virtual phenomenon is in fact morally wrong, or how to determine whether a virtual action counts as an action. Such problems appear irresolvable at worst, and resolvable only in the distant future at best. However, it's imperative that we settle on answers to the normativity of virtual phenomena in the present, and we have suggested that the conceptual approach, guided by CEQ, is up to the task.

One might say that the conceptual approach can offer no more practical guidance than either ethics or metaphysics. After all, it appears that there isn't a clear answer to CEQ.

<sup>3</sup> There appears to be a connection between needs/points and the notion of *function*, which is a hot topic in conceptual engineering (Simion & Kelp, 2019; Thomasson, 2020; Haslanger, 2020; Jorem, 2020; Podosky, 2022; Queloz, 2022). Specifically, the notion of an *etiological function*, those functions that exist as a result of selection over generations, looks to be related to the idea of a concept point and partly explains a concept's continued existence. We want to remain silent on the relationship between needs, points, and functions. This is because there is a lot of doubt that concepts even have such functions, or if they do, it is extremely unclear how we should think about them (Cappelen, 2018; Riggs, 2022). Despite this, the fact that concepts have needs and points seems, to us, undeniable. They explain why we have concepts at all. In saying this, however, we want to stress that function-talk is plausibly important to the question of CEQ. If functions provide individuation conditions for concepts, then this will certainly constrain the permissible answers we can give. That is, we cannot export in cases where the individuation conditions of a concept, defined by its function, is violated. But in the absence of knowing what a function is, or whether concepts have functions at all, we do not think that this should be a serious concern.

<sup>4</sup> One reason could be that it violates the individuation conditions of a concept defined by their function. However, we are non-committal here. Our concern is motivated by moral, social, political, and practical reasons.

<sup>5</sup> The same can be said for other important normative concepts, such as *sexual assault*, which could have the regrettable effect of undoing, or diffusing, the profoundly significant work of feminist agitations for the recognition of the seriousness of sexual violence.

As mentioned, one could be point-focused, need-focused, or somewhere in the middle. How would we resolve a dispute between these competing views? We accept that there is a kind of resolvability-problem here. However, we believe that it is of a different kind to the resolvability-problem in ethics and metaphysics. This is because the conceptual approach is largely *pragmatic* in spirit.

If one is point-focused, need-focused, or somewhere in the middle, there is an answer to CEQ that is, in principle, discernible. For instance, if one is point-focused, then what one must do to arrive at a clear answer to CEQ is to first ask what needs gave rise to the concept in non-virtual reality, and therefore endowed the concept with a point, and then ask whether those same needs are present in virtual environments, and whether the concept would be pointful for the same reasons. Or, if one is need-focused, then what one must do to arrive at a clear answer to CEQ is to discover the needs of concept-users in virtual environments, and to ask whether those needs could be served if a non-virtual concept were exported into the virtual domain (even if the concept develops a new point). Or, if someone is in the middle, they need to decide for themselves how they will weigh competing reasons, such as how to weigh concept need against concept point, and how to weigh these against other factors such as justice-promotion, overall flourishing, feasibility, and so on. In sum, it might be difficult to excavate these facts, but such facts are discernible in principle.

Speaking more generally, our goal is to prioritise the possibility of action, and we believe that CEQ allows us to consider a range of different reasons that concern how we want to live our lives, without having to settle on perennial debates in philosophy. CEQ allows us to question why concepts exist, what their purposes are, and whether we would be better off extending our concepts into new domains or instead keeping them in their original domain of use. In saying this, however, our view is not *simply* pragmatic. There are a host of normative reasons that one might appeal to when deciding how to answer CEQ, and whether one should be need-focused, point-focused, or somewhere in the middle. That is, we might reflect on moral, social, or political reasons. We can query, “Should I offer a need-focused response *all-things-considered*?” This, we take it, will depend on the context; the state of the normative landscape. If there is pressing moral, social, or political urgency to answer CEQ, because we need to know how to deal with virtual “murder” or “sexual assault”, then non-pragmatic factors can be introduced into our overall decision-making.<sup>6</sup>

Another objection that one could have is that answering CEQ might not avoid the problem of first having to deal with ‘unanswerable’ questions in ethics and metaphysics.

For example, take metaphysics. To the question of whether *murder* should be exported into virtual contexts, one might say, ‘well, that depends on whether virtual murder is *indeed* murder’, which imposes a metaphysical constraint on conceptual exportation. We might only be permitted to export a concept provided that certain metaphysical questions have been answered. Or, regarding ethics, one might say, ‘one should only export *murder* if the relevant wrong in the virtual world is sufficiently similar to the wrong in non-virtual domains.’ This imposes an ethical constraint on conceptual exportation.

So what we can see is that this seems to take us back to square one: we have to work out tough and seemingly unanswerable questions in ethics and metaphysics to even answer CEQ. However, we don’t think this is right. While many want to use ethics and metaphysics as a constraint on conceptual exportation, this is also a pragmatic decision. That is, there is no pre-theoretical reason to place these constraints on whether to export a concept from one domain to another. Instead, it is simply a decision one makes to guide their normative thinking. We believe that ethical and metaphysical reasons might be important to factor into one’s decision making, but they have to be weighed against other reasons, such as whether a concept can respond to needs, or whether it contributes to flourishing, etc.

Further, even if CEQ is hard to answer, its focus on concepts opens up a new way of thinking about how to understand moral problems in virtual worlds. It might be better to simply introduce a new concept to pick out a novel wrong. This is important to note because it seems that the problem of mapping out the normative contours of virtual reality isn’t so much a matter of trying to say why and how a virtual wrong is similar to some non-virtual wrong, but to simply articulate why a virtual wrong is wrong at all. We can do this without having to spell out the virtual wrong in terms of some non-virtual wrong we are already familiar with. So, instead of exporting the concept *murder*, and risking exportation feedback, it could be best to introduce a new concept, perhaps with a similar label, like *schmurder* or *e-murder* (though, we don’t recommend these words).<sup>7</sup>

## Conclusion

We have argued that while metaphysical and ethical approaches can tell us much about the normativity of virtual wrongdoings, they each face their own set of issues that

<sup>6</sup> We thank a very helpful reviewer for helping to clarify this point.

<sup>7</sup> Though, this might also risk changing the psychological associations with *murder* insofar as people might come to believe that *schmurder* is a kind of murder, maybe with a similar moral significance.

are as of yet unresolved. This has motivated an alternative conceptual approach, which rather than focusing on the ethical or metaphysical aspects of virtual wrongdoings alone, asks the further, prior question of whether a concept that originated and developed within a non-virtual context should be exported into a foreign virtual domain, what we have named the Conceptual Exportation Question. To answer this question, we must not only consider the metaphysical and ethical conditions that constitute, or make ethically salient, a virtual action, as these may be helpful resources, but should first consider whether wrong-tracking concepts used in virtual environments, respond to the same needs and have the same points as they do in non-virtual environments. If they do, then perhaps they ought to be exported regardless of their metaphysical or ethically unclear status, and if they don't, perhaps conceptual exportation should be resisted. Of course, answering the CEQ may not be generalisable given the diversity of virtual phenomena, and may also clash with strongly held intuitions that distinguish virtual from non-virtual phenomena. However, framing the problem in this way adds further nuance to the discussion by raising the as yet unexplored normative conceptual question regarding wrongdoings in virtual worlds.

**Funding** Open Access funding enabled and organized by CAUL and its Member Institutions.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Aarseth, E. J., Perron, B., & Wolf, M. (Eds.). (2014). *The Routledge companion to video game studies*. Routledge. <https://doi.org/10.4324/9780203114261>
- Ali, R. (2015). A new solution to the Gamer's dilemma. *Ethics and Information Technology*, 17(4), 267–274. <https://doi.org/10.1007/s10676-015-9381-x>
- Appiah, A. (1996). Race, culture, identity: Misunderstood connections. *The Tanner Lectures on Human Values*, 17, 51–136.
- Barnes, E. (2020). Gender and gender terms. *Nous*, 54(3), 704–730.
- Bartel, C. (2012). Resolving the Gamer's dilemma. *Ethics and Information Technology*, 14(1), 11–16. <https://doi.org/10.1007/s10676-011-9280-8>
- Brey, P. (2014). The physical and social reality of virtual worlds. In *The Oxford handbook of virtual reality*. Oxford University Press (pp. 42–54). <https://doi.org/10.1093/oxfordhb/9780199826162.013.029>
- Brey, P. (1999). The ethics of representation and action in virtual reality. *Ethics and Information Technology*, 1(1), 5–14.
- Brey, P. (2003). The social ontology of virtual environments. *American Journal of Economics & Sociology*, 62(1), 269–282. <https://doi.org/10.1111/1536-7150.t01-1-00011>
- Burgess, A., Cappelen, H., & Plunkett, D. (2020). In Burgess, H., Cappelen, A., & Plunkett, D. (Eds.), *Conceptual engineering and conceptual ethics*. Oxford University Press.
- Burgess, A., & Plunkett, D. (2013a). Conceptual ethics I. *Philosophy Compass*, 8(12), 1091–1101.
- Burgess, A., & Plunkett, D. (2013b). Conceptual ethics II. *Philosophy Compass*, 8(12), 1102–1110.
- Cappelen, H. (2018). *Fixing language: An essay on conceptual engineering*. Oxford University Press.
- Cappelen, H., & Plunkett, D. (2020). A guided tour of conceptual engineering and conceptual ethics. In A. Burgess, H. Cappelen, & D. Plunkett (Eds.), *Conceptual engineering and conceptual ethics*. Oxford University Press.
- Carroll, N. (1990). *The philosophy of horror: Or paradoxes of the heart*. Routledge.
- Chalmers, D. (2018). The virtual and the real. *Disputatio*, 9, 309–352. <https://doi.org/10.1515/disp-2017-0009>
- Chalmers, D. (2022). *Reality: Virtual worlds and the problems of philosophy*. WW Norton & Company.
- Craft, A. (2007). Sin in cyber-eden: Understanding the metaphysics and morals of virtual worlds. *Ethics & Information Technology*, 9(2), 205.
- Craig, E. (1990). *Knowledge and the state of nature: An essay on conceptual synthesis*. Oxford University Press.
- Davnull, R. (2020). What does the gamer do? *Ethics and Information Technology*. <https://doi.org/10.1007/s10676-020-09558-8>
- Dunn, J. (2012). Virtual worlds and moral evaluation. *Ethics and Information Technology*, 14(4), 255–265. <https://doi.org/10.1007/s10676-012-9298-6>
- Feagin, S. (1983). The pleasures of tragedy. *American Philosophical Quarterly*, 20(1), 95–104.
- Foucault, M. (1961). *Madness and civilization*. Routledge.
- Friend, S. (2008). Imagining fact and fiction. In *New waves in aesthetics* (pp. 150–169).
- Gribble, J. (1983). *Literary education: A reevaluation*. Cambridge University Press.
- Haslanger, S. (2000a). Going on, now in the same way. In A. Burgess, H. Cappelen, & D. Plunkett (Eds.), *Conceptual engineering and conceptual ethics*. Oxford University Press.
- Haslanger, S. (2000b). Gender and race: (What) Are they? (what) do we want them to be? *Nous*, 34(1), 31–55.
- Haslanger, S. (2020). Going On, Not in the Same Way. In A. Burgess, H. Cappelen, & D. Plunkett (Eds.), *Conceptual Engineering and Conceptual Ethics* (p. 0). Oxford University Press. <https://doi.org/10.1093/oso/9780198801856.003.0012>
- Heinrichs, J. (2020). Virtual action. *Ethics and Information Technology*, 23(3), 317–330. <https://doi.org/10.1007/s10676-020-09574-8>
- Hochman, A. (2017). Replacing race: Interactive constructivism about racialized groups. *Ergo*, 4, 61–92.
- Howe, L. (2008). Self and pretence: Playing with identity. *Journal of Social Philosophy*, 39(4), 564–582. <https://doi.org/10.1111/j.1467-9833.2008.00443.x>
- Huizinga, J. (1944). *Homo Ludens: A study of the play-element in culture*. Angelico Press.
- Hurd, H. (1996). The moral magic of consent. *Legal Theory*, 2(2), 121–146. <https://doi.org/10.1017/S1352325200000434>
- Jenkins, K. (2016). Amelioration and inclusion: Gender identity and the concept of woman. *Ethics*, 126(2), 394–421.
- Johansson, M. (2009). Why unreal punishments in response to unreal crimes might actually be a really good thing. *Ethics and*

- Information Technology*, 11(1), 71–79. <https://doi.org/10.1007/s10676-009-9182-1>
- Jorem, S. (2020). The good, the bad, and the insignificant—assessing concept functions for conceptual engineering. *Synthese*, 200(2), 106.
- Juul, J. (2005). *Half-real: Video games between real rules and fictional worlds*. MIT press.
- Juul, J. (2008). *The magic circle and the puzzle piece* (pp. 056–057). University Press.
- Levy, N. (2002). Virtual child pornography: The eroticization of inequality. *Ethics and Information Technology*, 4(4), 319–323. <https://doi.org/10.1023/A:1021372601566>
- Lorenzini, D. (2020). On possibilising genealogy. *Inquiry*. <https://doi.org/10.1080/0020174X.2020.1712227>
- Luck, M. (2009a). The Gamer's dilemma: An analysis of the arguments for the moral distinction between virtual murder and virtual paedophilia. *Ethics and Information Technology*, 11(1), 31–36.
- Luck, M. (2009b). Crashing a virtual funeral: Morality in Mmorpgs. *Journal of Information, Communication and Ethics in Society*, 7(4), 280–285. <https://doi.org/10.1108/14779960911004516>
- Luck, M. (2018). Has Ali dissolved the Gamer's dilemma? *Ethics and Information Technology*, 20(3), 157–162. <https://doi.org/10.1007/s10676-018-9455-7>
- Luck, M. (2022). The grave resolution to the Gamer's dilemma: An argument for a moral distinction between virtual murder and virtual child molestation. *Philosophia*. <https://doi.org/10.1007/s11406-021-00455-y>
- Marques, T. (2020). Amelioration vs. perversion. In *Shifting concepts: The philosophy and psychology of conceptual variability*. Oxford University Press.
- McCormick, M. (2001). Is it wrong to play violent video games? *Ethics and Information Technology*, 3(4), 277–287. <https://doi.org/10.1023/A:1013802119431>
- McDonnell, N., & Wildman, N. (2019). Virtual reality: Digital or fictional? *Disputatio*, 11(55), 371–397. <https://doi.org/10.2478/disp-2019-0004>
- Mildenberger, C. (2017). Virtual killing. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 174(1), 185–203.
- Montefiore, T., & Formosa, P. (2022). Resisting the Gamer's dilemma. *Ethics and Information Technology*, 24(3), 31. <https://doi.org/10.1007/s10676-022-09655-w>
- Nader, K. (2020). Virtual competitions and the Gamer's dilemma. *Ethics and Information Technology*. <https://doi.org/10.1007/s10676-020-09532-4>
- Nietzsche, F. 2006. (1887). *On the genealogy of morality*. Cambridge University Press.
- Novas, C. D. (2020). Carnap meets Foucault: Conceptual engineering and genealogical investigations. *Inquiry*. <https://doi.org/10.1080/0020174X.2020.1860122>
- Öhman, C. (2020). Introducing the pervert's dilemma: A contribution to the critique of Deepfake Pornography. *Ethics and Information Technology*, 22(2). <https://doi.org/10.1007/s10676-019-09522-1>
- Ostritsch, S. (2017). The amoralist challenge to gaming and the Gamer's moral obligation. *Ethics and Information Technology*, 19(2), 117–128. <https://doi.org/10.1007/s10676-017-9420-x>
- Patridge, S. (2011). The incorrigible social meaning of video game imagery. *Ethics and Information Technology*, 13, 303–312.
- Plunkett, D. (2016). Conceptual history, conceptual ethics, and the aims of inquiry: A framework for thinking about the relevance of history/genealogy of concepts to normative inquiry. *Ergo*, 3, 27–64.
- Podosky, P.M.-C. (2022). Can conceptual engineering actually promote social justice? *Synthese*, 200(2), 160.
- Queloz, M. (2022). Genealogy, evaluation, and engineering. *The Monist*, 105, 435–451.
- Queloz, M. (Forthcoming). Nietzsche's conceptual ethics. <https://doi.org/10.1080/0020174X.2022.2164049>
- Ramirez, E. (2020). How to (dis)solve the Gamer's dilemma. *Ethical Theory and Moral Practice*. <https://doi.org/10.1007/s10677-019-10049-z>
- Riggs, J. (2022). Deflating the functional turn in conceptual engineering. *Synthese*, 199, 11555–11586.
- Rockstar Games. (2008). Grand Theft Auto IV (PC version) [Video Game].
- Salen, K., & Zimmerman, E. (2003). Rules of play : game design fundamentals. In K. Salen & E. Zimmerman (Eds.), *Pról de Frank Lantz*. MIT Press.
- Searle, J. R. (1995). The construction of social reality (Nachdr.). Free Press.
- Seddon, R. F. J. (2012). Getting 'virtual' wrongs right. *Ethics and Information Technology*, 15(1). <https://doi.org/10.1007/s10676-012-9304-z>.
- Seddon, R. (2013). Getting 'virtual' wrongs right. *Ethics and Information Technology*, 15(1), 1–11. <https://doi.org/10.1007/s10676-012-9304-z>
- Simon, M., & Kelp, C. (2019). Conceptual innovation, function first. *Nous*, 54(4), 985–1002.
- Stenros, J. (2012). In defence of a magic circle: The social, mental and cultural boundaries of play. *Transactions of the Digital Games Research Association*, 1(2). <https://doi.org/10.26503/todigra.v1i2.10>.
- Thomasson, A. (2020). A pragmatic method for normative conceptual work. In A. Burgess, H. Cappelen, & D. Plunkett (Eds.), *Conceptual engineering and conceptual ethics*. Oxford University Press.
- Tillson, J. (2018). Is it distinctively wrong to simulate doing wrong? *Ethics and Information Technology*, 20(3), 205–217. <https://doi.org/10.1007/s10676-018-9463-7>
- Walton, K. (1990). *Mimesis as make-believe*.
- Walton, K. (2013). Fictionality and imagination reconsidered. In M. Ferraris, A. Voltolini, & C. Babero (Eds.), *Fictionalism to realism: Fictional and other social entities* (pp. 9–26). Cambridge Scholars Publishing.
- Williams, B. (1985). *Ethics and the limits of philosophy*. Harvard University Press.
- Wolfendale, J. (2007). My avatar, my self: Virtual harm and attachment. *Ethics and Information Technology*, 9(2), 111–119. <https://doi.org/10.1007/s10676-006-9125-z>
- Woodcock, S. (2013). Horror films and the argument from reactive attitudes. *Ethical Theory and Moral Practice*, 16(2), 309–324. <https://doi.org/10.1007/s10677-012-9338-7>
- Young, G. (2013). Enacting taboos as a means to an end; but what end? On the morality of motivations for child murder and paedophilia within gamespace. *Ethics and Information Technology*, 15(1), 13–23. <https://doi.org/10.1007/s10676-012-9306-x>
- Young, G. (2017a). Objections to Ostritsch's argument in 'The Amoralist challenge to gaming and the Gamer's moral obligation.' *Ethics and Information Technology*, 19(3), 209–219. <https://doi.org/10.1007/s10676-017-9437-1>
- Young, G. (2017b). Integrating poor taste into the ongoing debate on the morality of violent video games. *The Computer Games Journal*, 6(4), 227–237. <https://doi.org/10.1007/s40869-017-0044-5>
- Young, G. (2019). An expressivist account of the difference between poor taste and immorality. *Ethical Theory and Moral Practice*, 22(2), 465–482. <https://doi.org/10.1007/s10677-019-09998-2>
- Zack, N. (1993). *Race and mixed race*. Temple University Press.