Ethics and Information Technology (2024) 26:59
https://doi.org/10.1007/s10676-024-09798-y

ORIGINAL PAPER

Framing the Gamer's Dilemma

Michael Hemmingsen[1]

Accepted: 22 August 2024

Abstract
The Gamer's Dilemma is a much-discussed issue in video game ethics which probes our seemingly conflicting intuitions about the moral acceptability of virtual murder compared to virtual child molestation. But how we approach this dilemma depends on how we frame it. With this in mind, I identify three ways the dilemma has been conceptualized: the Descriptive Gamer's Investigation, which focuses on empirically explaining the source of our intuitions; the Gamer's Puzzle, which uses the dilemma to explore and test moral or metaphysical theories; and the Applied Gamer's Dilemma, which reconstructs the moral commitments underlying the intuitions to provide action-guidance. Clearly distinguishing these framings allows us to identify the distinct methodologies and criteria of success for each approach. This tripartite framework helps resolve confusions in the debate, highlights the need for experimental philosophy to test proposed resolutions, and opens space for a serious reconsideration of consequentialism. The paper thus provides conceptual clarity to move the discussion forward productively.

Keywords Gamer's Dilemma · Video games · Virtual ethics · Virtual actions · Trolley problem

Introduction

A long-running debate in the ethics of video games is the so-called "Gamer's Dilemma (GD)" (Luck, 2009). This dilemma draws attention to a set of conflicting intuitions. On the one hand, committing virtual murder (VM) in a video game is not usually seen as morally wrong, since VM doesn't harm anyone: it's "just a game." However, while the same line of reasoning seemingly applies to cases of virtual child molestation (VCM), many people hold the intuition that there is nonetheless something wrong with VCM. But it is unclear what distinction (if any) justifies these diverging judgments.

I argue that how we respond to the GD depends on how we conceptualize the dilemma in the first place. I suggest that there are three basic orientations we can take towards the GD: the Descriptive Gamer's Investigation, which focuses on explaining the source of our intuitions; the Gamer's Puzzle, which aims to use the GD as a way of exploring and testing moral or metaphysical theories; and the Applied Gamer's Dilemma, which emphasizes a response to the GD that in some way captures our own latent intuitions and moral commitments. Being clear about what problems the various formulations of the dilemma speak to allows us to be similarly transparent about what strategies are open to us in response. So long as we fail to conceptualize these three problems as separate, I argue, we will be confused about how we should go about answering the GD and what constitutes a successful response.

The descriptive Gamer's investigation

One way of interpreting the GD is purely in terms of the *explanation* of our intuitions in terms of the social sciences or psychology. This framing has no interest in the *justifiability* of the intuitions in question; it focuses purely on the fact that people *have* these intuitions, and those who frame the GD in this way are interested in accounting for why this is: there is no actual *paradox* involved, it isn't a *dilemma*. Hence, I refer to this framing as the "Descriptive Gamer's Investigation (DGI)." Given this framing, there are two main ways of responding to the DGI: rejection and explanation.

Rejection

Rejecting the DGI involves denying the truth of one or more of the premises. For instance, it may turn out, after empirical investigation, that the armchair philosopher's assumption

✉ Michael Hemmingsen
mhemmingsen@thu.edu.tw

[1] Tunghai University, Taichung, Taiwan

Published online: 01 September 2024

that the intuition that VM is not morally wrong is incorrect. If so, there is no inconsistency to be explained. Importantly, it isn't enough to merely point out that isolated individuals don't agree; a rejection denies the idea that there is something worth discussing in the DGI in terms of a generalized attitude.

Paul Formosa, et al. offer something like a rejection strategy in their empirical investigation of gamer's intuitions (Formosa et al., 2023). They found that both VM and VCM were considered morally unacceptable by respondents. Though this ameliorates the dilemma rather than completely rejecting it—while VM and VCM were both morally unacceptable, nonetheless VCM was seen as *more* so—it rejects the GD in its strongest form.[1]

Explanation

An explanation strategy assumes that people's standard intuitions really are what the GD claims they are. It then seeks "to establish causal regularities of a functional kind: causal claims that relate patterns of non-moral input to patterns of moral output" (Kahane, 2013, p. 426). This might be in terms social science, evolutionary biology, neurobiology, psychology (Luck, 2024, p. 4), and so on (Bunge and Skulmowski, 2014, p. 175). Explanation doesn't address the conflict at the level of justification, it focuses on the fact of it, i.e., that the people who affirm that VM is acceptable and that VCM is not at the same time recognize that these two beliefs seem inconsistent.

Bourne and Bourne offer a psychological explanation (Bourne & Bourne, 2019). In their view, what explains our differing intuitions about VM and VCM is our ability to find the motivations of the action intelligible. Murder in video games is usually based in motivations such as "rage, fear, frustration, greed and revenge," motivations which Bourne and Bourne argue are "more readily regarded as intelligible reasons for action than paedophilic sexual desire" (Bourne & Bourne, 2019, p. 138); that is, for most people it is difficult to find sexual desire for a child intelligible. Hence, we have the intuition that the (intelligible) action of murder is morally acceptable, whereas the (unintelligible) action of child molestation is not. The concept of intelligibility, however, doesn't *justify* these intuitions, rather it merely explains why have them.

According to Jens Kjeldgaard-Christiansen, our differing intuitions can be explained by the fact that child molestation is a "purity violation," which attracts "character judgments and moral disgust," a reaction which he argues is "asymmetrically accentuated in fictional and virtual settings" (Kjeldgaard-Christiansen, 2019, p. 99). Additionally, prototypical cases of child molestation—given that they are physical assaults "by a powerful adult on a powerless child" (Kjeldgaard-Christiansen, 2019, p. 95)—evoke a perpetrator-victim relation that prototypical cases of murder do not.

More recently, John Tillson offers two explanations for our differing intuitions: the Taboo Thesis, which points to the fact that we have been conditioned to react negatively to VCM, beyond what the moral dimension of the act justifies; and The Inoculation Thesis, according to which the prevalence of simulated violence conditions acclimatizes us to see VM as acceptable even though it is not (Tillson 2018).[2]

Evaluation of the DGI

The DGI is a perfectly legitimate way of framing the GD. However, with its empirical focus it is not a *philosophical* framing. While it deals with moral intuitions, a philosophical engagement with moral intuitions treats them as evidence for or against possible answers to normative questions, to "get things right, not to explain why we have a certain set of beliefs" (Kahane, 2013, p. 424). With the DGI, however, moral intuitions are treated as data that we seek to causally explain (Kahane, 2013).

We can illustrate these two perspectives with the classic example of the Trolley Problem (Foot, 1967). The Trolley Problem is originally framed as a normative problem, and as such we are encouraged to attempt to view it through the lens of which moral principle or principles we ought to follow, or which moral principles the dilemma helps to clarify. However, we can also treat it as a descriptive problem, asking "why, as a matter of psychological fact, people tend to approve of trading one life to save several lives in some cases but not others" (Greene, 2016, p. 175).

These two perspectives are not necessarily wholly independent of one another: it is possible that understanding *why* people tend to have the attitudes they do might shed some light on moral principles. As Guy Kahane argues, *if* we accept that our moral intuitions (defeasibly) track certain moral principles, then an investigation into the empirical

[1] A distinction can perhaps be drawn between a strong form of the GD (VM is acceptable and VCM is not) and a weaker form (VM is comparatively more acceptable than VCM). Formosa et al. provide a rejection of the strong GD, but not the weak one.

[2] It should be noted that both Formosa et al., Kjeldgaard-Christiansen and Tillson's accounts don't stop with rejection or explanation alone: in the case of Kjeldgaard-Christiansen and Tillson their approach is ultimately philosophical and offers a "dissolution" strategy. However, if we were to take this aspect of their arguments in isolation – treating them as social scientific accounts rather than philosophical ones – then this element amounts to an explanation. In Formosa et al.'s case, while they don't offer more than an empirical analysis in their 2024 paper, it's clear that they are interested in more than that, and see their work as supporting future philosophical analyses of the GD.

source of those intuitions may also shed light on how those principles ought to be understood (or, alternatively, how our intuitions are being misled and should therefore be defeated in this case) (Kahane, 2013). But despite this possible connection between the DGI and a normative answer to the GD, the fact remains that, in itself, the DGI is fundamentally aiming at answering a descriptive question, whereas the GD is usually taken to be asking an ethical one.

It may seem trite to draw attention to the difference between a descriptive and normative approach to the GD. However, in my experience, the philosopher's default—to focus on the normative element—is actually quite unusual. Certainly, when I have taught the GD to non-major undergraduates, they often jump straight to descriptive explanations. Similarly, when discussing the GD with fellow academics in closely related disciplines, a descriptive approach seems to be standard. For instance, a colleague in media studies' initial reaction to the GD was to account for it by noting that "a) there's a long cultural history of 'aestheticizing' murder and b) that history plays into practices of stylisation that distance us from the concept of the actual act" (Holm, 2024). The field outside philosophy with the greatest likelihood of being interested in the GD—game studies—also tends to be descriptive. As such, it is important to be clear about the difference between what is being sought by a normative answer vs. a descriptive one, and how we would go about answering each.

The Gamer's puzzle

Unlike the DGI, the Gamer's Puzzle (GP) is concerned with the normative. However, the focus on the GP is not on the morality of VM and VCM specifically, but on using the GD as a way of probing moral or metaphysical theories or concepts. For instance, Rebecca Davnall takes the GD as a method for investigating the ontological status of virtual actions (Davnall, 2021, p. 226); and Thomas Montefiore and Paul-Mikhail Catapang Podosky use the GD to engage in the "critical evaluation of concepts [relating to virtual worlds and ethics] to determine whether they ought to be removed, revised, or replaced" (Montefiore & Podosky 2024, p. 6). In other words, the GP uses the GD as a *tool* to answer some further question.

Solution

For the GP, the only option is to try to make the three elements of the GD coherent, either by denying that VM is acceptable or that VCM is not, or by explaining how both of these claims can be made consistent. The purpose of the GP as a tool for shedding light on further philosophical questions relies on drawing conclusions from our attempts to square the seemingly inconsistent set of three premises of the GD. We might compare this in some ways to the problem of evil, where God's omnipotence, omnibenevolence and omniscience are apparently in conflict with the existence of evil (Mackie, 1955). The problem of evil often not interesting in isolation, but in terms of what responses to it can tell us about the nature of God (Hunt, 2001; Hutcheson, 1992; McGrath, 1986); of freedom (Plantinga, 1974); of the good (Adams, 1999); virtue (Hemmingsen, 2020); eschatology (Hick, 1966), and so on. Similarly, the GD as approached through the lens of the GP is not a problem to be solved for its own sake, but as a means of establishing the correct moral theory, or the ontological status of virtual actions, or the appropriate concepts for considering virtual worlds.

One of the most common uses to which the GP has been put is testing normative ethical theories, by finding which deeper moral concept or principle is able to generate both the acceptability of VM and unacceptability of VCM. Hence, denying that there is really a conflict between these two claims can be considered as a "solving" strategy. Importantly, it is not essential for the GP if the concept or principle developed to solve the conflict reflects people's actual moral commitments. Hence, I refer to this framing as the Gamer's "Puzzle" instead of the Gamer's "Dilemma."

Evaluation of the GP

Responses to the GP are distinct in two senses. First, the solution to the conflict between our intuitions regarding VM and VCM do not *necessarily* need to ultimately connect back to our concrete moral commitments in a way that would change or ground our attitudes towards them. In other words, the GP doesn't have as its goal satisfactorily articulating our intuitions in a way that identifies and clarifies our actual moral commitments. In saying this, its nonetheless important that the background assumptions (that VM is acceptable and VCM not) are taken as plausible: they are not *purely* hypothetical.

Second, in focusing on some *further question*, it is one step abstracted from our concrete decision-making. That is, the GP is not an *applied ethics* framing, in which our main goal is to work out what we should do or think about certain actions or states of affairs: it focuses instead on using the GD to explore a higher-level issue. This can be contrasted with a framing concerned with real-world questions, such as whether it is morally okay for us to (continue) playing violent video games. The more abstracted focus of the GP is not directly concerned with these kinds of applied questions (even if it can end up shedding light on them).

In other words, if we treat the GD more as a puzzle of normative ethics/metaphysics/conceptual analysis, where the exploration of our intuitions is intended primarily as a way of evaluating and testing some further theory, we have the GP. Of course, just as we wouldn't say that the Trolley

Problem is a waste of time because it's focused on normative rather than applied ethics, framing the GD as a puzzle can also be worthwhile. Like the Trolley Problem, the GP can allow us to probe the logical boundaries of moral or metaphysical theories and concepts, and help us to explore implicit assumptions or bizarre implications of different theories when applied across cases, drawing attention to novel challenges that might not arise from more standard examples. Additionally, the framing of the GP allows us to explore hypothetical answers to the GD in an unconstrained way, providing a starting point that could be built upon and connected to actual moral commitments in subsequent work.

In separating out the GP we are therefore able to see how we can use the GD as tool for inquiry in higher-level questions of philosophy, rather than as an (applied ethics) attempt to answer questions about how we should concretely act in relation to VM and VCM. These two goals may be related, of course, but they are not the same, and it's important to be clear about our aim and the attendant methodological commitments. So long as we are not clear about whether we are framing the GD as a normative/metaphysical or applied question, we cannot be clear about what we're doing, what success looks like, and the methods appropriate to the particular task we're undertaking. It is therefore important to recognize the possibility of the GP as a distinct way of approaching the GD.

The applied Gamer's Dilemma

The Applied Gamer's Dilemma (AGD) stakes a middle ground between the DGI and the GP. The goal of the AGD is to gain insight into the morality of virtual actions by considering, developing and reconstructing the moral commitments that underlie our intuitions about VM and VCM via a process of reflective equilibrium.

This framing requires that responses explain both "the source of our moral intuitions *and* justif[y] why these intuitions are (not) morally reasonable" (Ulbricht, 2023). Of course, we may not be able to articulate those justifications adequately at the start. However, in the context of the AGD it should be expected that the response we end up with broadly speaks to our existing moral commitments (Luck, 2024). For instance, if it turns out that for most people their intuitions about the dilemma are virtue-ethical, and not at all deontological, then a response that solves the AGD in terms of deontology is not adequate to the AGD. In other words, the AGD is not simply concerned with making the set of three premises consistent, but in doing so in a way that draws on and is grounded in what we actually think, however inchoate that thinking is initially. And unlike the DGI, it cannot simply be an *explanation* of those beliefs: the issue in not simply *that* we have certain beliefs or intuitions, but whether those beliefs are ultimately *justified*.

Taking the GD as an applied problem is to distinguish it from the approach of something like the Trolley Problem, in that the Trolley Problem is framed as a closed problem, whereas the AGD is framed as an open one. By this I mean that the details of the Trolley Problem are clear and fixed: we know the precise details about what options are available—pushing the fat man, switching the tracks, etc.—and the consequences of each choice. By contrast, in the AGD the details remain vague: it is still a matter of dispute, for instance, what the consequences of VM or VCM actually are. The emphasis of each problem is therefore not quite the same: the Trolley Problem *qua* Trolley Problem—like the GP—tends more towards a puzzle; in that it is explicitly aimed at plumbing our intuitions in order to reach conclusions about our moral principles. It is not an *applied* problem to be solved (or, rather, resolved) in that we are not taking as our immediate goal the solution to real-life moral dilemmas: we're not imagining that most people will be facing the problem of whether or not to divert an actual trolley. By contrast, the open nature of the AGD focuses on how we should think about the morality of VM and VCM *directly*. Of course, developing and testing moral principles is a *part* of that, but it is not the end goal. We can therefore contrast the Trolley Problem-like GP with the AGD in the following way: the GP is focused on theory-testing, whereas the AGD is focused on action-guidance.[3]

One potential criticism of the distinction between the AGD and the GP, however, is that *if* we can solve the GP, we could then simply apply that theory directly, bypassing the AGD entirely. In which case, the AGD is superfluous. But while a thought experiment like the GP can provide *a* reason in favor of a particular ethical theory, this is not in itself a *decisive* argument for that theory. There is a long road from having *a* reason in favor of a theory, to then using that theory as an uncontroversial basis for a top-down analysis of the GD. Doing so assumes that "the nature and description of the problem or quandary is not in dispute" (Flynn, 2020). The AGD, drawing on the methodologies of applied ethics, is a more bottom-up, open-ended approach, in which the very definition of the problem is still up for discussion. It involves a process of wide reflective equilibrium, in which our considered judgments about specific cases, the principles derived from our normative theories, and other relevant

[3] Of course, the Trolley problem absolutely *can* be framed with action-guidance in mind: it can and has been applied to questions in distributive justice (Kelman & Kreps 2014); the programming of autonomous vehicles (Geisslinger 2021; Wu 2020); decision-making in war (McMahan 2009), and so on. When framed in this more applied way, the focus shifts from an evaluation of our moral commitments alone, to a wider reflective equilibrium with our applied moral decision-making more generally; in which case our goal is to work out what we should do in real-world cases. When framed in this way, the Trolley Problem becomes closer to the AGD.

background theories are brought into coherence through a process of mutual adjustment (Arras, 2007; Daniels, 1979; Tillson, 2018, p. 207). An ethical theory that coheres with our intuitions in a highly abstract thought experiment may still face challenges when applied to a concrete, real-world case with numerous morally salient details. Importantly, too, by maintaining some autonomy from ethical theory or metaphysics, the AGD allows us to continue discussing and making progress on applied questions even while we hold our views about the correct moral theory or correct ontological account of virtual accounts somewhat in abeyance. This is particularly valuable given the ongoing debates and disagreements within normative ethics and metaphysics.

Dissolution

In the normative version of the Trolley Problem, we can "either explain the moral difference between the two scenarios, or… deny that there is such a difference" (Di Nucci, 2014, p. 82). Both of these options are available to the GD as well: we can accept the validity of our intuitions about VM and VCM and deny that they are inconsistent, or we can accept the inconsistency and undermine "the intuitions that ground [the GD] by denying, in some or all cases, that [VM] is morally permissible or that [VCM] is morally impermissible" (Montefiore & Formosa, 2022, p. 1). This latter approach has been called a *dissolution* strategy (Luck, 2023). Importantly, dissolution doesn't deny *that* we have those particular intuitions (like rejection does): it denies instead that the intuitions we have are *justified*. In dissolution, then, we trace our actual moral reasoning (suitably reconstructed) and discover that our views lack support. For example, we might initially think that the principle of autonomy justifies VM but not VCM, but after careful consideration realize that it disallows both. A dissolution can be strengthened when it can provide "a theory of error of how one came to have those mistaken judgements" in the first place (Tillson, 2018, p. 208). In this respect, it benefits from engaging with the DGI.

A dissolution strategy is one of the more popular ones in the literature. It is usually preceded by a "narrowing" strategy, so most of the examples of dissolution will be postponed until later in the paper. However, Tillson offers an example of a standalone dissolution (Tillson, 2018). It is possible, Tillson suggests, that we have psychological difficulty distinguishing between virtual and real actions: what he calls the Squeamishness Thesis. If we combine the Squeamishness Thesis with his Taboo Thesis (that we are conditioned to react in an unwarrantedly negative way to VCM) then we may need to treat VCM like VM. Alternatively, if we combine the Squeamishness Thesis with the Inoculation Thesis (that we illegitimately treat VM as acceptable due to overexposure) then we may need to treat VM as being as bad as VCM. Either way, the conflict between the intuitions disappears because we have found reason to deny our intuitions about one or the other.

Resolution

Secondly, we might *resolve* the GD by rejecting the inconsistency between our intuitions about VM and VCM (Luck, 2023): we identify a difference that justifies our differing intuitions. Resolving is focused on identifying a principle that "grounds and thus explains" the initial intuitions (Kahane, 2013, p. 426). In this sense, a resolution helps us to identify a moral principle—not directly a moral *theory* (Beauchamp & Childress, 1979)—that "both explains moral propositions with more particular content and is the source of our reasons to believe that they are true" (Kahane, 2013, p. 426).

Several attempts have been made to resolve the AGD. Perhaps the earliest is Stephanie Patridge's (Patridge, 2011). In Patridge's view, games can have an incorrigible social meaning. VM, in certain contexts (such as when the representations involve gendered or race-based violence) can have incorrigible social meanings that involve "egregious, long-term, systematic denials of justice" (Patridge, 2011, p. 310)—they, in other words, contain oppression. In *these* cases, VM is morally wrong. However, VM is not morally wrong *in itself*; a particular act of murder gains an oppressive social meaning as a result of additional details like those above. By contrast, VCM unavoidably contains oppressive social meaning, due to it involving particular classes of person (children and adults) and the issues of justice that arise due to the unequal power relations between those groups. If so, then we can affirm that there is a key difference between VM and VCM: VCM is necessarily oppressive, whereas VM is not, and so our differing intuitions about VM and VCM are not inconsistent at all.

Christopher Bartel also offers a resolution, arguing that VCM is necessarily an instance of child pornography, whereas VM is not (Bartel, 2012, p. 165). The immorality of child pornography and its necessary connection to VCM explains why we justifiably think that VCM is wrong, but not VM.

More recently, Morgan Luck argues for a distinction between the two virtual acts on the grounds of their "graveness" (Luck, 2022, p. 1299). For Luck, there is a distinction between "fair-game wrongdoings"—those wrongdoings which can be treated lightly, since they sufficiently lack graveness—and "off-limits wrongdoings"—wrongdoings for which being treated lightly is impermissible, because doing so fails to recognize their graveness. As shown by our willingness to watch murder mysteries, but not child molestation mysteries (at least, not with the same sense of frivolity), murder is a fair-game wrongdoing (it lacks sufficient graveness), but child molestation is not (it is an excessively

grave act). Since playing a game involves treating the virtual acts involved in that game lightly, and since it is okay to treat murder lightly, but not child molestation, then we are justified in thinking that VCM is wrong, but not VM.

Thomas Coghlan and Damian Cox argue that while murder is also a serious moral wrong—"an act of enormous moral significance and a fitting occasion of the most profound remorse" (Coughlan & Cox, 2023)—the harms of death are only *necessarily* deprivation harms: "The dead are deprived of life and everything good and bad that comes with it" (Coughlan & Cox, 2023). By contrast, child molestation involves the creation of profound suffering, not merely deprivation.[4] This then justifies our differing intuitions about the two kinds of virtual actions: VM does not necessarily involve profound suffering, whereas VCM does.

Mattia Ceccinato argues that it is wrong to fantasize about immoral things. As he puts it, to fantasize is "(i) to actively and extensively imagine A (ii) in the context of your own life and (iii) as the object of some pro-attitude, such as desire or approval" (Cecchinato, 2024, pp. 79–80). It is difficult, Ceccinato thinks, to conceive of reasons for virtual sexual assault that don't involve the act of (immorally) fantasizing about immoral things, whereas it is quite easy to do so for murder: killing in a video game is often "a more *impersonal* kind of imagining" (Cecchinato, 2024, p. 80); it is not necessarily a subject of fantasy. As such, VCM is always morally wrong—it always involves immoral fantasies—whereas VM is not, since it doesn't.

Thomas Montefiore and Paul Formosa—drawing on Neil Levy (2002)—suggest that social attitudes about sexual assault as compared to murder might lead VCM to lead to normalizing sexual assault, whereas VM is less likely to do the same in the case of murder, though they note that empirical evidence is required to support these contentions, and so they are at this stage merely speculative (Montefiore & Formosa, 2022, p. 5).

Finally, Gary Young has argued for an approach he refers to as "constructive ecumenical Expressivism" (CEE) (Young, 2014, 2015, 2016). According to this view, our approval or disapproval of an action is crucially connected to its rightness. Since VM is generally not disapproved of and VCM generally is, then this provides a moral distinction between the two. Interestingly, Young's work can be taken as either a solution or a resolution, depending on the emphasis: where the focus is on providing support for CEE as metaethical view, it is a solution; where the emphasis is on providing an action-guiding answer to the GD, it is a resolution.

[4] Of course, this is in relation to the dead themselves – the effect of that death on others is another matter. We are also assuming a distinction between the process of dying, which can involve suffering, and being dead, which does not.

Narrowing

A third kind of strategy is narrowing. What is in common amongst narrowing approaches is that they see the relevant difference between our intuitions about the morality of VM and VCM as involving a confounding factor: our intuitions are led by something that isn't connected to VM or VCM per se. This strategy, then, is to reformulate the GD to remove that confounding factor, to specify "contextually-equal" (Montefiore & Formosa, 2022, p. 7) situations and compare those instead. As Montefiore and Formosa put it, a narrowing strategy involves adding something like "in *x* cases" to each of the three premises. In other words, while in the case of resolutions, a necessary difference is found that distinguishes between VM and VCM, for a narrowing the difference cuts across these categories.

For instance, Rami Ali argues—in a view expanded on by Karim Nader (2020)—that the GD should be narrowed by distinguishing between sporting, storytelling, and simulation games (Ali, 2022). Each of these three types of game "demand different types of ludic engagement from the player," and in Ali's view "this mode of engagement is relevant to evaluating in-game acts" (Ali, 2022). Each of the three kinds of ludic engagement speak to a separate dilemma: relating to the moral status of game acts (competition), virtual acts (simulation), and representations (storytelling). If we want to adequately respond to the GD, we need to narrow the dilemma, recognizing that these three forms of ludic engagement need to be evaluated independently of one another, and then compare acts of VM in a single form of ludic engagement with acts of VCM in that same type only.

Several of the other narrowing strategies focus similarly on the attitudes of the player. For instance, Sebastian Ostritsch argues that we need to distinguish between virtual acts in which the player enjoys the action in the sense of having *fun*—what he calls "endorsing" the act—from cases in which the player does not endorse the virtual act (Ostritsch, 2017, 2021). Similarly, Samuel Ulbricht contrasts cases in which a player "imaginatively transforms" virtual acts into real crimes, in the sense of adopting a maxim in which the player really tries (inevitably unsuccessfully) to "really carry out murder and molestation while playing" (Ulbricht, 2023). For Christopher Bartel, we need to distinguish between virtual acts that are freely chosen, in Harry Frankfurt's sense, i.e., that we identify with the action that we have performed (Bartel, 2015), and those that are not. And Erick Jose Ramirez distinguishes between virtual actions that create "virtually real experiences" and those that don't (Ramirez, 2020). Finally, Young (2013) suggests that we can distinguish between cases in which the motivation is to benefit an overall strategy—what he calls "M $_{(\text{strategic})}$"—to have fun—(M

$_{(enjoyment)}$) or because I have a desire to engage in that activity in real life (M $_{(substitution)}$).

More recently, Young distinguishes between "idle" and "surrogate" fantasies. Idle fantasies can be enjoyed without an accompanying desire to actually engage in that activity, whereas the enjoyment of surrogate fantasies lies in satisfying the desire to actually engage in that activity (Young, 2020). Since it seems plausible that someone who seeks out VCM is highly likely to do so on the basis of a surrogate fantasy, whereas someone who seeks out VM is more likely to be doing so on the basis of an idle one, then our intuitions are warranted, though just so long as we are comparing idle VM to surrogate VCM, rather than like to like.

Narrowing, however, is not a strategy that is sufficient on its own. Once the GD has been narrowed, we must either pursue a dissolving or resolving strategy (Montefiore & Formosa, 2022). Most of the above cases lend themselves to a subsequent dissolving strategy: endorsing VM is just as bad as endorsing VCM (Ostritsch, 2017, 2021); it is wrong to identify with both VM and VCM (Bartel, 2015); both VM and VCM, if imaginatively transformed into real murder and real child molestation, are wrong (Ulbricht, 2023); it is just as wrong to commit virtually real acts of murder as acts of virtually real child molestation (Ramirez, 2020); and surrogate VM is just as bad as surrogate VCM (Young, 2020).

Rami Ali (2022) and Gary Young's (2013) approaches are particularly interesting cases of (potential) narrowing-then-resolving strategies. For both, the GD is narrowed into three kinds: virtual actions that are (probably) morally acceptable in both cases (sporting and M $_{(strategic)}$); virtual actions that are (probably) morally unacceptable in both cases (simulation and M $_{(substitution)}$); and virtual actions for which there is much more room for debate (storytelling and M $_{(enjoyment)}$). While both Ali (2022, p. 273) and Young argue that their approaches dissolve the dilemma—they believe that there is no "basis for the selective prohibition of one virtual act over the other" (Young, 2013, 18) in any of the three categories—I think it can also be worthwhile to consider their distinctions as more open-ended ways of considering how the GD could be understood, introducing the possibility for further dissolving *or* resolving strategies, in particular in terms of their more ambiguous middle category.

Montefiore and Formosa offer a narrowed resolution along these very lines (Montefiore & Formosa, 2022). Drawing on Ramirez's concept of virtually real experiences and Ali's distinction between storytelling and simulation games, they argue that while there is a "lower-level" of game—a game with low fidelity and a strong storytelling focus—in which both VM and VCM are morally permissible; and a "higher-level" of game—with high fidelity and a great deal of player agency (simulation)—in which both actions are morally impermissible, there is a mid-range of fidelity and player agency in which VCM is wrong (because it is treated as inappropriately pornographic) whereas VM is not.

Denial

Since the AGD—like the DGI—is concerned with reconstructing our actual intuitions about VM and VCM, it has an equivalent to the rejection response: denial, or the "amoralist" view. If rejection is the claim that there is nothing to the GD in the sense that its assertion that its intuitions do not represent a generalized *social* phenomenon, denial consists in asserting that the GD doesn't represent *our own* intuitions: we don't share the intuition that VM is morally acceptable, or VCM is morally wrong, or that there is a conflict between these two intuitions, and as such there is no particular reason to engage with the AGD in the first place.

I have found that denial is a common strategy amongst students: some simply flatly deny that they have the intuition that VCM is morally wrong—like VM, it doesn't harm anyone, so there is no reason to be concerned about it. Typically, given that the class requires them to continue engaging with the GD, they do: but their attitude switches either to the DGI—"even if *I* don't share those intuitions, it's clear that others have them: why is this?"—or the GP—"though I don't share the intuitions, treating VM as morally acceptable and VCM as morally wrong, and trying to square those two things, is an interesting intellectual puzzle."

Of course, the fact that some people deny having the intuitions that ground the GD does not mean we should stop talking about it. "I just don't share your intuitions," would probably not be considered an adequate response in an academic paper on the GD. Denial is therefore not wholly individualistic: it matters whether the denial is widely shared or whether it is idiosyncratic. However, denial differs from rejection in that widespread denial undermines not that there is a social phenomenon worth studying, but rather the extent to which we should lean on intuitions in the process of reflective equilibrium.

Benefits of the tripartite formulations

By dividing the GD into three types of basic problem—the DGI; the GP; and the AGD—we gain a range of benefits.

First, we can identify more clearly what criteria must be met by a successful response, as this depends on what we have set out to achieve. If we consider the matter of justification unimportant—if our main concern is *that* we have certain intuitions—then the DGI, and its rejection and explanation strategies, is sufficient. By contrast, if we are treating the GD as a tool for testing moral or metaphysical theories, then we should be focused on solving the GP. Finally, if we see the GD as a matter of exploring and justifying *our own*

Table 1 Possible responses

Strategy	Problem	Description
Rejection	DGI	Denies that the relevant intuitions are widely shared
Explanation	DGI	Explains the source of our intuitions empirically
Solution	GP	Provides a morally significant distinction without concern for our intuitions
Dissolution	AGD	The intuitions we have are not justified
Resolution	AGD	A reconstruction of our intuitions provides a morally significant distinction
Narrowing	AGD	Clarifies confounding factor, allows for like-to-like comparison
Denial	AGD	Denies that the relevant intuitions constitute strong evidence

intuitions about VM and VCM via reflective equilibrium—where we try to develop action-guidance through degrees of fit between "principle and agents' intuitions about particular cases" (Aguiar & Rodríguez-López 2014, p. 195) —then we need to adopt one of the strategies appropriate to the AGD.

Being clear about what we are trying to achieve is especially important in the context of utilizing a successful resolving strategy. Explanation, solving, and resolving each focus on the seeming conflict between VM being morally permissible and VCM being impermissible, but end up speaking to slightly different problems. A resolving strategy must therefore walk a difficult tightrope. It requires that the moral distinction between the two cases be both justified *and* that that justification represents a reconstructed version of our actual thinking about these virtual actions; that it reflects in some way our real moral commitments after a process of reflective equilibrium. But so long as we are not clearly distinguishing between the AGD, the GP, and the DGI, it's easy to mistake a resolution for a solution or explanation, with what is lost from this—a meaningful connection to actual moral commitments that make a resolution action-guiding on the one hand, and a normative grounding for our intuitions on the other—being overlooked.

By distinguishing between these three aims, then, we can avoid the talking-past-each-other that can so often characterize discussions involving scholars with various approaches, goals, and disciplinary backgrounds. The tripartite framing may also provide a useful heuristic for researchers to clarify *their own* aims and methodological commitments when approaching the dilemma, leading to a more focused and rigorous investigations (see Table 1).

Second, recognizing the AGD as distinct from the GP and the DGI, but related to both, points us towards the necessity for engaging in experimental philosophy (Ulatowski, 2024). By recognizing that a resolution requires a connection to our moral commitments, it becomes clear that we need to investigate what those commitments actually are in the case of virtual actions. The various resolution strategies discussed above all presumably *hope* that their reconstructions reflect people's views on the morality of virtual actions. But until we really investigate this question empirically, we simply don't know. It may turn out that none of these supposed resolutions *are* connected to the moral commitments people actually have, in which case they are solutions rather than resolutions. Identifying the action-guiding element of the AGD could therefore spur fruitful collaborations between philosophers and experimental philosophers or psychologists (such as the recent work of Formosa et al., 2023). This kind of work has already been fruitful in relation to the Trolley Problem (Di Nucci, 2014; Greene, 2014, 2016; Kahane, 2013; Matthew Liao et al. 2012); the same can be true of the GD as well.

Third, identifying the GP as a distinct approach draws attention to how much more loosely specified the GD is compared to the Trolley Problem. A more open-ended framing is fine when our goal is to answer questions in applied ethics, but if we are using the GD to make progress on questions in normative ethics, an open problem is simply not going to work (again, recall that the GD, unlike the Trolley Problem, does not stipulate the consequences of VM and VCM in advance). In order to reach useful (normative ethical) conclusions from the GP, we need to generate a variety of specific, detailed cases, such as those we find in the context of the Trolley Problem. For instance, the Trolley Problem allows us to contrast our intuitions about the footbridge variant—where we push a person off the bridge—and the trapdoor variant—where we press a button to open a trapdoor that drops the person. In exploring these two otherwise identical cases, we can draw useful conclusions about morally relevant distinctions. The GD—as broadly, and loosely, stated—does not really allow this. If we're clear that we are considering the GD through the framing of the GP, we can see that, given our explicit aim of theory testing, it is essential that we stipulate details and multiply cases in the same way that occurs in discussions of the Trolley Problem. So long as we remain vague about whether we're pursuing the GP or the AGD, however, the necessity of this is harder to notice.

Finally, being clear about precisely what is required for a successful resolution to the GD also opens up a possible rehabilitation of the most disrespected of the ethical approaches in the context of the GD: consequentialism. Consequentialism is frequently discussed in the context of the GD, but usually only to quickly dismiss it in order to move

on to more "plausible" views. But the very fact that consequentialism is frequently raised as an option—often first—despite its seeming implausibility, is telling: it suggests that our basic intuitions about the morality of virtual actions may very well run along consequentialist lines (certainly, my students' views on the GD are typically consequentialist by default). If so, then this is a reason to think that any attempts to resolve the GD must take consequentialism seriously; that we should do more to try to develop consequentialist resolutions; and that we should investigate the nature of our consequentialist intuitions about VM and VCM as a matter of priority. Of course, it may well be the case that our intuitions are consequentialist, but that no adequate consequentialist resolution can be found. But if we care about connecting the GD to the source of our own intuitions, if those intuitions really do swing towards consequentialism, we have no choice but to give that approach its due.

Conclusion

The aim of this paper was to provide greater conceptual clarity on the different ways the Gamer's Dilemma can be posed. I argued that the dilemma can have three distinct framings: the Descriptive Gamer's Investigation, the Gamer's Puzzle, and the Applied Gamer's Dilemma. By carefully defining the aims, assumptions, and criteria for success of each framing, I showed how we can more clearly evaluate the relevance and impact of different response strategies, such as rejection, explanation, solution, denial, dissolution, resolution and narrowing.

Some of the benefits of this tripartite division include: first, it helps us locate confusions in existing discussions. This may allow for more targeted exchanges between scholars with different aims and methodologies. Second, it points to the necessity of experimental philosophy collaborations to empirically test proposed resolutions against actual moral intuitions and commitments. Third, it highlights how stipulating specific cases and consequences is required when approaching the GD from the perspective of normative ethics. Finally, it suggests that we should reconsider consequentialism as a plausible approach to the GD, given its prima facie alignment with common intuitions about virtual actions.

Data availability We do not analyse or generate any datasets, because our work proceeds within a theoretical approach.

References

Adams, M. M. (1999). *Horrendous evils and the goodness of god*. Cornell University Press.

Aguiar, F., Gaitán, A., & Rodríguez-López, B. (2014). Robust intuitions, experimental ethics and experimental economics: Bringing reflective equilibrium into the lab. In C. Luetge, H. Rusch, & M. Uhl (Eds.), *Experimental Ethics: Toward an Empirical Moral Philosophy* (pp. 191–208). Palgrave Macmillan. https://doi.org/10.1057/9781137409805_13

Ali, R. (2022). The video Gamer's Dilemmas. *Ethics and Information Technology, 24*(18), 6. https://doi.org/10.1007/s10676-022-09638-x

Arras, J. D. (2007). The way we reason now: Reflective equilibrium in bioethics. In B. Steinbock (Ed.), *Oxford handbook of bioethics* (pp. 46–71). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199562411.003.0003

Bartel, C. (2012). Resolving the Gamer's Dilemma. *Ethics and Information Technology, 14*, 11–16. https://doi.org/10.1007/s10676-011-9280-8

Bartel, C. (2015). Free will and moral responsibility in video games. *Ethics and Information Technology, 17*, 285–293. https://doi.org/10.1007/s10676-015-9383-8

Beauchamp, T. L., & Childress, J. F. (1979). *Principles of biomedical ethics*. Oxford University Press.

Bourne, C., & Bourne, E. C. (2019). Players, characters, and the Gamer's Dilemma. *Journal of Aesthetics and Art Criticism, 77*(2), 133–143. https://doi.org/10.1111/jaac.12634

Bunge, A., & Skulmowski, A. (2014). Descriptive and pragmatic levels of empirical ethics: Utilizing the situated character of moral concepts, judgment, and decision-making. In C. Luetge, H. Rusch, & M. Uhl (Eds.), *Experimental ethics: Toward an empirical moral philosophy* (pp. 175–190). Palgrave Macmillan. https://doi.org/10.1057/9781137409805_12

Cecchinato, M. (2024). Double-standard moralism: Why we can be more permissive within our imagination. *British Journal of Aesthetics, 64*(1), 67–87. https://doi.org/10.1093/aesthj/ayad008

Coghlan, T., & Cox, D. (2023). Between death and suffering: Resolving the Gamer's Dilemma. *Ethics and Information Technology, 25*, 37. https://doi.org/10.1007/s10676-023-09711-z

Daniels, N. (1979). Wide reflective equilibrium and theory acceptance in ethics. *The Journal of Philosophy, 76*(5), 256–282. https://doi.org/10.2307/2025881

Davnall, R. (2021). What does the gamer do? *Ethics and Information Technology, 23*, 225–237. https://doi.org/10.1007/s10676-020-09558-8

Di Nucci, E. (2014). Trolleys and double effect in experimental ethics. In C. Luetge, H. Rusch, & M. Uhl (Eds.), *Experimental ethics: Toward an empirical moral philosophy* (pp. 80–93). Palgrave Macmillan. https://doi.org/10.1057/9781137409805_7

J Flynn. 2020. Theory and bioethics. In: Edward N. Zalta, U Nodelman (Eds.) *The stanford encyclopedia of philosophy* (Winter 2022 Edition), Retrieved from https://plato.stanford.edu/archives/win2022/entries/theory-bioethics/

Foot, P. (1967). The Problem of abortion and the doctrine of the double effect. *Oxford Review, 5*, 5–15.

Formosa, P., Montefiore, T., Ghasemi, O., & McEwan, M. (2023). An empirical investigation of the Gamer's Dilemma: A mixed methods study of whether the dilemma exists. *Behaviour & Information Technology*. https://doi.org/10.1080/0144929X.2023.2178837

Geisslinger, M., Poszler, F., Betz, J., Lütge, C., & Lienkamp, M. (2021). Autonomous driving ethics: From trolley problem to ethics of risk. *Philosophy and Technology, 34*, 1033–1055. https://doi.org/10.1007/s13347-021-00449-4

Greene, J. D. (2014). Beyond point-and-shoot morality: why cognitive (neuro)science matters for ethics. *Ethics, 124*(4), 695–726. https://doi.org/10.1086/675875

Greene, J. D. (2016). Solving the trolley problem. In J. Sytsma & W. Buckwalter (Eds.), *A companion to experimental philosophy* (pp.

175–178). Wiley Blackwell. https://doi.org/10.1002/9781118661666.ch11
Hemmingsen, M. (2020). Soul-making and social progress. *International Journal for Philosophy of Religion, 88*(1), 81–96. https://doi.org/10.1007/s11153-019-09726-x
Hick, J. (1966). *Evil and the god of love*. Harper and Row.
Nicholas Holm. 2024. Private communication. April 2.
Hunt, D. P. (2001). Evil and theistic minimalism. *International Journal for Philosophy of Religion, 49*(3), 133–154.
Hutcheson, P. (1992). Omniscience and the problem of evil. *Sophia, 31*(1–2), 53–58. https://doi.org/10.1007/BF02772352
Kahane, G. (2013). The armchair and the trolley: An argument for experimental ethics. *Philosophical Studies, 162*, 421–445. https://doi.org/10.1007/s11098-011-9775-5
Kelman, M., & Kreps, T. A. (2014). Playing with trolleys: Intuitions about the permissibility of aggregation. *Journal of Empirical Legal Studies, 11*(2), 197–226. https://doi.org/10.1111/jels.12039
Kjeldgaard-Christiansen, J. (2019). Splintering the Gamer's Dilemma: Moral intuitions, motivational assumptions, and action prototypes. *Ethics of Information Technology, 22*, 903–102. https://doi.org/10.1007/s10676-019-09518-x
Levy, N. (2002). Virtual child pornography: The eroticization of inequality. *Ethics and Information Technology, 4*, 319–323. https://doi.org/10.1023/A:1021372601566
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. https://doi.org/10.1007/s10676-008-9168-4
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, 50*, 1287–1308. https://doi.org/10.1007/s11406-021-00455-y
Luck, M. (2023). Has montefiore and formosa resisted the Gamer's Dilemma? *Ethics and Information Technology, 25*, 31. https://doi.org/10.1007/s10676-023-09705-x
Luck, M. (2024). Can we solve the Gamer's Dilemma by resisting it? *Ethics and Information Technology, 26*, 39. https://doi.org/10.1007/s10676-024-09781-7
Mackie, J. L. (1955). Evil and omnipotence. *Mind, 64*, 200–212. https://doi.org/10.1093/mind/lxiv.254.200
Matthew Liao, S., Wiegmann, A., Alexander, J., & Vong, G. (2012). Putting the trolley in order: Experimental philosophy and the loop case. *Philosophical Psychology, 25*(5), 661–671. https://doi.org/10.1080/09515089.2011.627536
McGrath, P. J. (1986). Evil and the existence of a finite god. *Analysis, 46*(1), 63–64. https://doi.org/10.1093/analys/46.1.63
McMahan, J. (2009). *Killing in war*. Oxford University Press.
Montefiore, T., & Formosa, P. (2022). Resisting the Gamer's Dilemma. *Ethics and Information Technology, 24*, 31. https://doi.org/10.1007/s10676-022-09655-w
Nader, K. (2020). Virtual competitions and the Gamer's Dilemma. *Ethics and Information Technology, 22*(239–245), 45. https://doi.org/10.1007/s10676-020-09532-4
Ostritsch, S. (2017). The amoralist challenge to gaming and the gamer's moral obligation. *Ethics and Information Technology*. https://doi.org/10.1007/s10676-017-9420-x
Ostritsch, S. (2021). The immorality of computer games: Defending the endorsement view against young's objections. *Ethics and Information Technology, 23*, 199–205. https://doi.org/10.1007/s10676-020-09549-9
Patridge, S. (2011). The incorrigible social meaning of video game imagery. *Ethics and Information Technology, 13*, 303–312. https://doi.org/10.1007/s10676-010-9250-6
Plantinga, A. (1974). *God, freedom, and evil*. Harper and Row.
Ramirez, E. J. (2020). How to (dis)solve the Gamer's Dilemma. *Ethics and Information Technology, 23*, 141–161. https://doi.org/10.1007/s10677-019-10049-z
Tillson, J. (2018). Is it distinctively wrong to simulate doing wrong? *Ethics and Information Technology, 20*, 205–217. https://doi.org/10.1007/s10676-018-9463-7
Ulatowski, J. (2024). Experimental philosophy. In M. Rossberg (Ed.), *The Cambridge handbook of analytic philosophy*. Cambridge University Press.
Ulbricht, S. (2023). A kantian response to the Gamer's Dilemma. *Ethics and Information Technology, 25*, 39. https://doi.org/10.1007/s10676-023-09710-0
Wu, S. S. (2020). Autonomous vehicles, trolley problems, and the law. *Ethics and Information Technology, 22*, 1–13. https://doi.org/10.1007/s10676-019-09506-1
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*. https://doi.org/10.1007/s10676-012-9306-x
Young, G. (2014). A meta-ethical approach to single-player gamespace: Introducing constructive ecumenical expressivism as a means of explaining why moral consensus is not forthcoming. *Ethics and Information Technology, 16*(2), 91–102. https://doi.org/10.1007/s10676-014-9336-7
Young, G. (2015). Violent video games and morality: A meta-ethical approach. *Ethics and Information Technology, 17*(4), 311–321. https://doi.org/10.1007/s10676-016-9386-0
Young, G. (2016). *Resolving the Gamer's Dilemma*. Palgrave Macmillan.
Young, G. (2020). *Fictional immorality and immoral fiction*. Lexington Books.