

Robotic Rape and Robotic Child Sexual Abuse: Should they be criminalised?

By John Danaher*

Forthcoming in *Criminal Law and Philosophy* DOI: 10.1007/s11572-014-9362-x

Abstract

Soon there will be sex robots. The creation of such devices raises a host of social, legal and ethical questions. In this article, I focus in on one of them. What if these sex robots are deliberately designed and used to replicate acts of rape and child sexual abuse? Should the creation and use of such robots be criminalised, even if no person is harmed by the acts performed? I offer an argument for thinking that they should be. The argument consists of two premises. The first claims that it can be a proper object of the criminal law to regulate wrongful conduct with no extrinsically harmful effects on others (the moralistic premise). The second claims that the use (and possibly the manufacture) of robots that replicate acts of rape and child sexual abuse would be wrongful, even if such usage had no extrinsically harmful effects on others. I defend both premises of this argument and consider its implications for the criminal law. I do not offer a conclusive argument for criminalisation, nor would I wish to be interpreted as doing so; instead, I offer a tentative argument and a framework for future debate. This framework may also lead one to question the proposed rationales for criminalisation.

Keywords: Criminalisation; Robotics; Rape; Child Sexual Abuse; Legal Moralism; Public Wrongs

* Lecturer in Law, NUI Galway

1. Introduction

Sex robots are already a reality.¹ The current crop provide their users with crude and unsophisticated facsimiles of human sex. Nevertheless, the technology is advancing, and it is probable that more realistic sex robots will be created in the future.² What happens if this technology allows users to engage in acts of *robotic* rape and *robotic* child sexual abuse. Should this type of activity be criminalised? Should the creation of such devices be banned? More narrowly, should they be criminalised even if their use has no extrinsically harmful effects on others?

In this article, I offer an initial, and tentative, argument in favour of criminalisation. The argument consists of two main premises. The first holds that it can be a proper object of the criminal law to prohibit wrongful conduct even if it has no extrinsically harmful effects on others (the *moralistic* premise). The second holds that acts of robotic rape and child sexual abuse (even if they have no extrinsically harmful effects on others) would fall within the class of wrongful conduct that it is a proper object of the criminal law to prohibit (the *wrongness* premise). The result is a *prima facie* case for criminalisation. I then develop this case further by considering the possible extrinsic effects of robotic rape and child sexual abuse, as well as the possible forms that any criminalisation could take. I do not intend, nor would I hope to be interpreted as intending, to provide a conclusive argument in favour of criminalisation. Instead, I hope to provide a framework for future debate about this issue. I acknowledge at the outset that my argument could either support a future case for criminalisation, or cast more doubt on certain rationales for criminalisation.

The remainder of the article consists of six parts. First, I try to clarify the nature of the technology being discussed, its possible uses and the social context in which it may be used. Second, I set out the basic argument and introduce its two main premises. Third, I address the *moralistic* premise, looking specifically at recent defences of legal moralism from Duff and Walls. Fourth, I set out a range of arguments in favour of the *wrongness* premise, relying on claims about the wrongness of virtual acts and on recent findings from moral and cognitive psychology. Fifth, I look at the possible extrinsic effects of such technology, arguing that empirical evidence in relation to extrinsic

¹ The most famous example is TrueCompanion's to male and female sex robots, *Roxy* and *Rocky*. See <http://www.truecompanion.com> (visited 31/7/14)

² Levy, D. *Love and Sex with Robots* (New York: Harper Perennial, 2007) gives an overview of the possible future. I will discuss this in more detail in section 2.

effects is likely to be ambiguous, and, if so, this may actually tip the balance in favour of criminalisation. And then sixth, and finally, I look at the implications of the argument and the possible forms that any criminalisation could take.

2. What are we talking about?

The argument that I present concerns *purely robotic* forms of rape and child sexual abuse. The concept of a “purely robotic” form of either activity is something that is open to misunderstanding. To avoid this, four initial questions must be asked and answered. What is a sex robot? What would count as an act of robotic rape or child sexual abuse? What is meant by referring to those acts as being “purely” robotic? And what sort of social context would such robots be used in?

In answer to the first question, I will propose a definition. For present purposes, a sex robot can be defined as any artifact that is used for sexual stimulation and/or release with the following three properties:³ (i) a humanoid form; (ii) the ability to move; and (iii) some degree of artificial intelligence (*i.e.* some ability to sense, process and respond to signals in its surrounding environment). Sex robots are different from traditional sex toys and other artifacts for sexual stimulation, which will tend to either replicate specific body parts, and so not have a humanoid form, or, if they do have such a form, will tend not to have any degree of artificial intelligence.

Primitive sex robots are in existence today. The most infamous of these is *Roxxy/Rocky*, created by the company TrueCompanion. It is billed as the world’s first sex robot.⁴ *Roxxy/Rocky* takes a humanoid form, has some ability to move, and can interact with its users via a set of pre-programmed personality types.⁵ Videos of *Roxxy/Rocky* are available online and give a sense of how crude the technology currently is.⁶ Nevertheless, there have been significant developments in the field of

³ I have proffered a similar definition in a previous paper [reference omitted]

⁴ For details, see <http://www.truecompanion.com> (visited 31/7/14)

⁵ See <http://www.truecompanion.com/shop/faq> (visited 31/7/14) for details on the personality types. For the female robot, these are described with monikers like “Frigid Farah”, “Wild Wendy” and “Young Yoko”. The latter is particularly interesting in light of the present discussion, she is described as being “oh so young (barely 18) and waiting for you to teach her”. I discuss the example of “Frigid Farah” below as well.

⁶ See <http://www.truecompanion.com/tv/> (visited 31/7/14) for a series of videos. These are reasonably non-explicit.

humanoid robotics and artificial intelligence that give a sense of where the technology could eventually go.⁷

There are two important points to raise about this definition of sex robots. The first is that it is deliberately restricted to “human-like” robots. That is: robots that are designed to provide an artificial facsimile of a real human sex partner. It is possible that manufacturers could create sex robots that are not human-like in appearance or behaviour, *e.g.* animal or non-humanoid alien sex robots. My argument will not focus on those possibilities. I make specific claims about the social meaning of robots with human-like features that may not transfer over to non-humanoid or animal robots. The significance of those forms of robotic sexual encounter is something that others might consider in the future. Furthermore, for my purposes, the degree of human-likeness is determined by the three properties listed above. Thus, the more human-like in appearance, movement, and intelligence the robot is, the more human-like it is. This is important insofar as the argument I present for criminalisation depends on sex robots crossing some minimal threshold of human-likeness. That minimal threshold might be quite low — indeed, it is possible that sex robots like Roxxxxy have already crossed it — but the more human-like the robot is, the more persuasive the argument becomes.

This brings up the second important point about the definition of sex robots. For the remainder of this discussion, the focus will be on human-like sex robots that do not meet the criteria for legal or moral personhood. Thus, I will assume that the sex robots are not capable of being morally harmed by the actions of their users. By making this assumption, I do not rule out the possibility of robots having moral status.⁸ This could one day be possible and if so I think that any such robots should be entitled to all the legal and moral protections afforded to beings with that status. But if that were true, the case for criminalisation would be much more straightforward.⁹ My focus is on the more difficult case in which the robot is not deemed to have moral status and hence cannot be a moral victim.

⁷ For example, human-like appearance, touch and movement have advanced considerably in recent years. For a good example of this, see the HRP-4C dancing robot, which can be viewed at <http://www.youtube.com/watch?v=1C9XBQmlGbo> (accessed 31/7/14). Couple this with advances in AI assistants — such as Apple’s SIRI — and a clearer picture of the future of sex robots emerges.

⁸ Gunkel, D. *The Machine Question: Critical Perspectives on AI, Robots and Ethics* (Cambridge MA: MIT Press, 2012) has a lengthy analysis of this issue.

⁹ Though see Petersen, S. “Designing People to Serve” in Lin, P. Lin, P., Abney, K., and Bekey, G. (eds) *Robot Ethics: The Ethical and Social Implications of Robotics* (Cambridge, MA: MIT Press, 2012). Petersen argues that robots who meet the criteria for personhood could be designed to serve our needs without thereby being morally harmed.

With a clearer sense of what a sex robot is, we can move onto the second question: what would count as an act of robotic rape and child sexual abuse? Robotic child sexual abuse is the easier case. This would arise whenever someone performs sexual activity with a robot that looks and acts like a child.¹⁰ Robotic rape is slightly trickier. If we assume that rape is any non-consensual penetrative sex act with the absence of reasonable belief in consent,¹¹ then two problems would seem to emerge. The first is that any act of robotic sex would seem to be trivially non-consensual: if the robot is not a moral agent, then it is not capable of granting consent.¹² This is true, but I think it can be set to one side as it is possible for robots to be created that deliberately mimic signals of non-consent. Such sex robots may even exist today. One of *Roxxy*'s pre-programmed personalities is called, by her makers, "Frigid Farah". We are told that if "you touched her in a private area, more than likely, she will not be to[o] appreciative of your advance."¹³ Admittedly, this is a pretty incomplete description of how she interacts with her users, but it does suggest a signal of non-consent. Whatever the case may be with *Roxxy*, I take it that anyone who engages in (penetrative) sexual activity with a robot that signals non-consent is engaging in an act of robotic rape.

This raises another problem with classifying acts of robotic rape: what counts as a signal of non-consent? And, perhaps more importantly,¹⁴ what would lead to the absence of a reasonable belief in consent? This is a notoriously contentious issue, both among philosophers who study the nature of sexual consent,¹⁵ and among the general public.¹⁶ I certainly do not wish to downgrade the seriousness of all forms of non-

¹⁰ Interestingly, computer-imaged virtual children have already been used to entrap real paedophiles. "Sweetie" was an ultra-realistic CGI representation of a 10-year old Filipino child that was created by Hans Guyt and his team at Terre des Hommes Netherlands, which is an organisation fighting child exploitation. The representation was used when would-be paedophiles seek visual confirmation of who they are interacting with via on-line chat rooms. It was created and used as a proof of concept. It has subsequently been shut down. For a discussion see Grothaus, M. "Why are we building jailbait sexbots?" *Fast Company Labs*, 14 November 2013. Available at <http://www.fastcolabs.com/3021651/why-are-we-building-jailbait-sexbots> (visited 31/7/14)

¹¹ This is the definition adopted in England and Wales. Section 1, *Sexual Offences Act 2003*.

¹² Alternatively, one might say that it is a category mistake to apply terms like "consent" to an artifact of this sort.

¹³ See, once more, <http://www.truecompanion.com/shop/faq> (visited 31/7/14)

¹⁴ I say more importantly because the argument I'm developing assumes no moral victim and so focuses entirely on the character of the perpetrator.

¹⁵ For example, Wertheimer, A. *Consent to sexual relations* (Cambridge: Cambridge University Press, 2002) discusses the many conditions that may or may not negate consent.

¹⁶ This seems to be confirmed by empirical studies of opinions about reasonable belief see Finch and Munro "Breaking Boundaries? Sexual Consent in the Jury Room" (2006) 26 *Legal Studies* 303-330 (for mock jury experiments); and Carline and Gunby "How an Ordinary Jury Makes Sense of it is a Mystery" (2011) 32 *Liverpool Law Review* 237-250, (for practitioner's opinions).

consensual sex, nor do I wish to introduce morally irrelevant distinctions into the classification of different varieties or types of rape. Nevertheless, I do wish to focus on what we might call paradigmatic signals of non-consent, *e.g.* unambiguous verbal refusals or degrees of forcible resistance. I think this focus is appropriate given that this article is intended to be a first take on the issue, and given that the robots themselves are not assumed to be moral victims. In future, non-paradigmatic signals of non-consent can be addressed.

This brings us to the third question: what is a purely robotic act of rape or child sexual abuse? Very simply, this is an act of rape or child sexual abuse with a robot that has *no extrinsic effects on others*. In other words, an act that does not encourage its perpetrator to engage in similar acts with human beings, or, indeed, discourage him/her from performing such acts with human beings. It is certainly plausible that acts of robotic rape and child sexual abuse could have such extrinsic effects, and in later sections I will speculate about some possible effects, but my initial argument will focus on the purely robotic cases.

That leaves us with the final question, which is really a set of questions concerning the social context in which these robots will be used. First one may wonder who the users will be. In this article, I am tacitly assuming that men will be the likely users. Indeed, one of the arguments I offer in relation to robotic rape is specifically tailored to the social meaning attached to the male-user scenario, so the assumption becomes explicit in that case. As for child sexual abuse, I again assume the majority of users will be male, but the argument I offer could just as easily apply to women in that case. Irrespective of the users, one may also rightly wonder how prevalent the use of such robots will be and whether this will be a significant phenomenon in the near future. It is, of course, difficult to say for sure since we are speculating about a nascent technology. But it's likely that the use of sex robots (in general) will be reasonably significant. Early polls indicate that anywhere from 10 to 17% of the population would be willing to try them out.¹⁷ As for the more specific variants that cater to rape or child

¹⁷ Huffpost/Yougov poll from 2013 found around 10% of American people would be willing to have sex with a robot (see: http://www.huffingtonpost.com/2013/04/10/robot-sex-poll-americans-robotic-lovers-servants-soldiers_n_3037918.html, accessed 21/11/14); A study conducted by Martin Smith of Middlesex University polling 2000 Britons found that 17% would be willing to have sex with a robot (see: *The Guardian* "A Third of Britons fear the rise of the robots" 6 May 2014). Finally, a Pew Report entitled "AI, Robots and the Future of Jobs" also includes the statements of some commentators who believe that sex robots will "become commonplace" (the quote is from Stowe Boyd, head of GigaOM research), see: <http://www.pewinternet.org/files/2014/08/Future-of-AI-Robotics-and-Jobs.pdf> (accessed 21/11/14).

sexual abuse fantasies, they may well prove attractive to a significant percentage of those who are willing to use sex robots, or to those who have certain sexual fantasies. Rape fantasies, for example, are common among men,¹⁸ with some studies suggesting that over 40% of men have them.¹⁹ The thought of acting out those fantasies on a robot that lacks moral status may prove attractive to some given the absence of a moral victim. The same argument would seem to apply to the use of a child robot that lacks moral status. This may prove attractive to those who already fantasize about child abuse, but are unwilling to act out those fantasies on real children. It may also, of course, prove attractive to those who are willing to act out the fantasies on real children but lack access to them or wish to avoid the social repercussions.

This raises some related questions about the distribution and production of such robots. One possibility is that the robots are manufactured by commercial retailers and sold on the market. This is why I include manufacturers within the discussion of my criminalisation argument, below. But there are other possibilities too. For example, the robots could be assembled from parts that are sold for other (more legitimate) purposes. Users could share designs and assembly instructions via online forums. More fancifully, we could imagine a future in which parts could be 3D-printed in the homes of potential users.²⁰ The method of distribution does not matter greatly for my argument as my primary focus is on use and whether that justifies criminalisation.

Finally, the use of robots for sexual stimulation and pleasure raises a number of interesting questions about human sexuality and the social significance of its expression. For example, if the users are predominantly male and the robots are designed to replicate certain stereotypical female attributes, we may wonder about their

¹⁸ On sexual fantasies in general, see Bader, M. *Arousal: The Secret Logic of Sexual Fantasies* (London: Virgin Books, 2003). For studies on the frequency of rape and dominance fantasies among men, see: Schewe, PA, Adam, NM and Ryan, KM “A qualitative analysis of the temptation to use force in sexual relationships” (2009) 24(2) *Violence and Victims* 219-231 (22 out of 83 interviewees reported a temptation to use force during sex) and Zubriggen EL, and Yost MR “Power, desire and pleasure in sexual fantasies” (2004) 41(3) *Journal of Sex Research* 288-300. Female rape fantasies are also widely studied, but there the focus is typically on the woman being raped. It is possible that robots would be constructed to cater to this fantasy, but I ignore that possibility here since it would not seem to raise issues in relation to criminalisation (particular when the social meaning argument advanced in section 5 is factored in).

¹⁹ Crepault, C and Couture, M. “Men’s Erotic Fantasies” (1980) 9(6) *Archives of Sexual Behaviour* 565 (sample size was 92).

²⁰ Immersive virtual reality systems -- such as those provided by Oculus Rift -- are another possibility. I discuss virtual acts below in section 5 and, as I note there, I think there is something distinct about using the physical robot. The use of a virtual reality headset with some form of teledildonics (hardware that simulates sexual friction) would be a borderline case. There are already companies specialising in this. For example, Kiiroo (which focuses on those in long distance relationships) <https://kiiroo.com> (accessed 21/11/14).

contribution to a culture of sexual objectification and oppression (irrespective of whether they are used to act out rape fantasies). I will not be able to explore all these issues in the ensuing discussion -- others may wish to take them up in the future -- but the social meaning of use (particularly in the rape case) is something that features heavily in the argument I will now defend.

3. The Basic Argument

So what is that argument? We can start with a bare, unqualified version. It consists of two premises and a conclusion:

(1) It can be a proper object of the criminal law to regulate conduct that is morally wrong, even if such conduct has no extrinsically harmful effects on others (the *moralistic* premise).

(2) Purely robotic acts of rape and child sexual abuse fall within the class of morally wrong but extrinsically harmless conduct that it can be a proper object of the criminal law to regulate (the *wrongness* premise).

(3) Therefore, it can be a proper object of the criminal law to regulate purely robotic acts of rape and child sexual abuse.

Three comments about this argument at the outset. First, the argument is deliberately vague. It cannot work unless it is shown that purely robotic forms of rape and child sexual abuse fall within the precise subset of that general class that it is a proper object of the criminal law to regulate. In other words, we need to narrow down the scope of premises (1) and (2) before the argument can work. This is something we will do below by considering more precise theories of legal moralism.

Second, the argument uses the term “regulate” rather than, say, “punish” because it does not assume that punishment is always the justifying aim of criminal law.²¹ Instead, the argument assumes that there are certain kinds of wrongdoing that it is a proper object of the criminal law to deal with, even if that doesn’t immediately entail that the wrongdoer ought to be punished (where “punishment” is understood to require

²¹ This is contrary to the theory proposed by Michael Moore in *Placing Blame* (Oxford: OUP, 1996), which suggests that retributive punishment is the goal of criminal law.

harmful treatment).²² This is important both for assessing the overall plausibility of the argument, and its implications.

Third, the argument does not purport to offer a conclusive case in favour of the criminalisation of robotic rape and child sexual abuse. It merely purports to offer a *prima facie* case in favour of criminalisation. That is to say: it argues that robotic rape and child sexual abuse are things we ought to seriously think about criminalising (that fall within the relevant class of conduct), while at the same time acknowledging that the reasons in favour of criminalisation could be overridden by other factors. It will not be possible to address all of these potentially overriding factors in this particular article, but a few of the most salient will be discussed later on.

With this initial statement of the argument out of the way, we can proceed to consider the reasons for endorsing its two main premises.

4. The Moralistic Premise

The first premise of the argument appeals to the concept of legal “moralism”. Unfortunately, this is a vague concept. As popularised through venues like the Hart-Devlin debate,²³ moralism is often associated with an expansive, prudish, and overreaching desire to criminalise private acts of vice. But there are other, less expansive and overreaching, interpretations. At its most basic level, moralism simply demands that immorality (or wrongfulness) is required for criminalisation. This, by itself, is not too controversial. The subsequent implication of this for policies of criminalisation is rather more controversial. Are all immoral acts covered? Does immorality supply a necessary or sufficient condition for criminalisation?²⁴

For the purposes of my argument, I will adopt a partial necessity-based version of moralism. In other words, I will assume that conduct must be wrongful (immoral) before it can be criminalised. In making this assumption, I will not be making the further assumption that wrongfulness is a sufficient condition for criminalisation. Other

²² I have in mind definitions of punishment included in books like Boonin, D. *The Problem of Punishment* (Cambridge: CUP, 2008) and Zimmerman, P. *The Immorality of Punishment* (2012).

²³ See Devlin, P. *The Enforcement of Morals* (London, 1963); and Hart, HLA *Law, Liberty and Morality* (London, 1963)

²⁴ See Simester and Von Hirsch *Crimes, Harms and Wrongs* (Oxford: Hart Publishing, 2011), pp. 22-30 for a discussion of this issue. As they note, one may even think that wrongfulness is a necessary condition for criminalisation, without thinking it provides even a *pro tanto* reason in favour of criminalisation.

factors need to be worked out before we can reach a final determination about whether a particular type of conduct ought to be criminalised. Furthermore, I will not assume that all wrongful conduct is up for grabs in debates about criminalisation. Duff's distinction between *ambitious* and *modest* forms of legal moralism is useful in this regard.²⁵ As he defines those concepts, "ambitious legal moralism" refers to the view that all immoral conduct is apt for criminalisation; "modest legal moralism" refers to the view that only specific sub-types of immoral conduct are apt for criminalisation. Again, for the purposes of this argument, I will adopt a modest version of legal moralism.

In adopting a modest version of legal moralism, a further question is raised: which sub-types of immoral conduct do I think are apt for criminalisation? Given that my focus is on purely robotic acts of rape and child sexual abuse, my argument will only work if the version of moralism implied by premise one is expansive enough to cover immoral conduct that has no extrinsically harmful effects on others. Any such version of moralism would run contrary to some of the leading liberal and libertarian theories of criminal law.²⁶ According to those theories, criminal law should be restricted to regulating immoral conduct that is either harmful to others, poses a significant risk of harm to others, or — at a push — that is offensive to others.²⁷ The problem is that purely robotic acts of rape would involve none of these three things as they would not affect others in any way.²⁸

As it happens, these liberal theories — which one might call "restrictive" versions of moralism — are ones for which I have a lot of sympathy. I think we should be concerned about expanding the scope of criminal law too far,²⁹ especially given that criminalisation is, in many ways, the most serious expression of public disapproval we now have at our disposal. Indeed, it is my sympathy for this view that causes me to be cagey and tentative about the argument presented in this article. Nevertheless, there are

²⁵ Duff, A. "Towards a Modest Legal Moralism" (2014) 8 *Criminal Law and Philosophy* 217-235

²⁶ I speak here of the Millian tradition, first set out in *On Liberty* (London, 1859), and defended and expanded upon by the likes of Hart — *Law, Liberty and Morality* (n 19) — and Feinberg, J. — *The Moral Limits of the Criminal Law* Four Volumes (Oxford: OUP, 1984-88).

²⁷ E.g. Feinberg, J. *Offense to Others* (Oxford: OUP, 1985); see also Simester and Von Hirsch (n 24), Part III: Offence.

²⁸ Just to reiterate, it is not that robotic rape and child sexual abuse could never involve those things — indeed, I think it is highly likely that they could — it is just that my focus is on the purified case.

²⁹ In this sense, I am sympathetic to the view of Husak in *Overcriminalization* (Oxford: OUP, 2009). That said, it may be that the risk of overcriminalization is minimised by the fact that most instances of so-called "morals" legislation cover conduct that is actually perfectly morally acceptable. For example, I do not believe that drug-taking or prostitution is immoral. Thus, for me at least, there is no risk of moralism leading to the criminalisation of those acts.

reasons to doubt restrictive moralism and there are some, recently articulated, and more expansive, versions of moralism that could allow for purely robotic acts of rape and child sexual abuse to be criminalised. I wish to appeal to those versions of moralism in my defence of premise one. I do so largely on an *arguendo* basis. That is to say, I will not explicitly defend these versions of moralism here. Instead, I will simply sketch each theory in some detail, and say something like “if you accepted this version of moralism, premise one would be compelling”. Once this has been done, the focus can switch to the specific arguments about the nature of purely robotic acts of rape and child sexual abuse.

The first, more expansive, version of legal moralism to which I appeal comes from the work Steven Wall.³⁰ He argues that there are four primary domains of moral concern: (i) autonomy; (ii) well-being; (iii) moral character; and (iv) impersonal excellence³¹. The first three relate to the individual, the last to the intrinsic value of the world outside the individual. As Wall point out,³² many traditional theories of criminalisation argue that the criminal law can only govern immoral conduct affecting the first domain, *i.e.* interferences with individual autonomy and, to some extent, the second domain, *i.e.* conduct that negatively impacts on individual well-being. Sometimes the concern for the former domain clashes with the latter, particularly if individual acts of autonomy impact on individual well-being (e.g. acts of self-harm). In fact, this is the quintessential clash animating debates between liberals/libertarians and paternalists.

Wall tries to defend a theory of criminalisation that allows for the criminal law to regulate conduct that affects the third domain as well. His argument for this builds upon a connection between well-being and moral character.³³ It claims that if the criminal law can rightly concern itself with the well-being of the individual, so too should it concern itself with their moral characters. The idea is that moral habits and dispositions are integral to a sense of self-respect,³⁴ and that self-respect in turn is essential to the well-

³⁰ Two sources for this: Wall, S. “Enforcing Morality” (2013) 7 *Criminal Law and Philosophy* 455-471; and Wall, S. “Moral Environmentalism” in Coons, C. and Weber, M. *Paternalism: In Theory and Practice* (Cambridge: Cambridge University Press, 2013)

³¹ By this he means to cover objects and events that have an intrinsic value that is not fully explained by how humans interact and respond to them. For example, the value of the natural (non-man-made) environment.

³² See Wall “Enforcing morality” (n 30), section “Articulating the Presumption”

³³ See “Enforcing morality” (n 30), section “Self-respect and character”

³⁴ Where this is understood as the commitment to and belief in the worth of one’s life and life projects.

lived life. So you cannot hope to have well-being without building a sound moral character.

As stated above, my goal here is not to defend Wall's view. I merely wish to have it granted it for the sake of argument. So my claim is that, if you agree with Wall that the regulation moral character is a legitimate concern of the criminal law, you should embrace premise one, and in due course premise two, of my argument. Nevertheless, it does behoove me to say a little more in defence Wall's view lest it be deemed completely implausible.³⁵

The obvious problem with Wall's view is that it contradicts the liberal commitments to the value of autonomy and moral independence. If the criminal law is going to regulate conduct that impacts upon an individual's moral character, then it seems like the commitment to the individual's right to choose, and to take responsibility for the shape of their moral life, is compromised.

This is an attractive criticism, at least at first glance. But as Wall himself points out, it is not entirely compelling. Autonomy and moral independence are complex values. They can be negatively affected by a whole range of factors, many of them non-legal. With regard to autonomy, criminalisation, depending on its form, need not completely compromise this capacity. In other words, it need not close off options from the individual: it may just make the options harder to take. Furthermore, even if it does, it is possible that by diminishing choice in some ways, we could actually increase it in others.³⁶ This may be particularly true for those in the grip of "compulsive" or "irresistible" desires. These desires may so consume them that they are unable to act in accordance with their preferred and more reflective desires. It is possible that certain types of sex offender (e.g. paedophiles) are in the grip of such desires, which may in turn help the argument pursued below.

³⁵ In addition to the arguments mentioned, see Scoccia, D. "In defense of "pure" legal moralism" (2013) 7 *Criminal Law and Philosophy* 513-530, which argues that liberals cannot consistently deny pure forms of moralism if they wish to uphold a commitment to criminalising things like the desecration of corpses.

³⁶ Wall's himself makes this point in "Moral Environmentalism". The same point is defended in many sources. One, recent and rather nice defence arises in the context of the human enhancement debate. See Earp, B., Sandberg, A., Kahane, G. and Savulescu, J. "When is diminishment a form of enhancement" (2014) 8 *Frontiers in systems neuroscience* 12

With respect to moral independence, responding to an argument from Dworkin,³⁷ Wall highlights that it is very difficult for the law to be morally neutral with respect to the lives of its subjects. The decision not to legislate in a particular domain can have just as much of an impact on moral independence as the decision to legislate in a particular domain.³⁸ To assume that non-criminalisation is always better with respect to moral independence than criminalisation is an error.

Again, the point is not to defend this view. The point is merely that it is possible for one to accept Wall's view without necessarily compromising one's commitment to the core liberal values like autonomy and moral independence. This makes the argument I am defending about purely robotic rape and child sexual abuse slightly more palatable to those of a liberal persuasion.

In any event, Wall's expansive legal moralism is not the only resource upon which I can call. Antony Duff's defence of modest legal moralism also offers up possibilities for the criminalisation of robotic rape and child sexual abuse. Duff's theory centres on the notion of a "public wrong", *i.e.* an instance of immoral conduct that is publicly salient in such a way that it is right and proper for the wrongdoer to be publicly held to account and condemned. The criminal law, for Duff, must be understood as a public institution, with a distinctive social meaning. At a minimum, three conditions must be met before we can properly criminalise any conduct:³⁹ (i) the conduct must be immoral; (ii) it must require or make appropriate a social or collective response; and (iii) the wrongfulness must deserve to be made publicly salient, *i.e.* to receive public attention, censure and condemnation. Duff notes how this theory of criminalisation can extend the scope of criminal law into cases involving no consequential (extrinsic) harm to others, specifically citing extreme pornography and hate speech as examples.⁴⁰

To return to my argument, the point once again is not to defend Duff's particular theory, but to argue that if one accepts it, my case for the criminalisation of purely robotic rape and child sexual abuse becomes plausible. As I will argue below, it is

³⁷ See Wall "Enforcing Morality" (n 30) section "Authenticity and Independence"

³⁸ There are shades here of the libertarian paternalist thesis defended by Sunstein and Thaler in "Libertarian Paternalism" (2003) 93(2) *American Economic Review* 175-179

³⁹ These are taken from Duff, "Towards a Modest Legal Moralism" (n 25) section "Criminalization and wrongdoing"

⁴⁰ *Ibid.* See also: Johnson, P. "Law Morality and Disgust: The Regulation of 'Extreme Pornography' in England and Wales" (2010)19 *Social and Legal Studies* 147 which highlights how expansive moralism seems to have motivated the criminalisation of the possession of extreme pornography.

possible for both types of act to fall within the class of public wrongs and thereby warrant a public response.

To sum up, the moralistic premise is vague in its initial statement. But by appealing to specific theories of moralism, we can add the detail necessary for the rest of the argument to work. In what follows, I will work with the two versions of moralism just outlined. Thus, I will take it that if I can show that purely robotic acts of rape and child sexual abuse are wrongful because they either (a) involve harm to the moral character of those performing them or (b) fall within the class of public wrongs, my argument can succeed.

5. The Wrongness Premise

We now reach the crux of the argument. The wrongness premise claims that purely robotic acts of rape and child sexual abuse do indeed fall within the set of wrongful acts that it is a proper object of the criminal law to regulate. I want to spend the majority of my time defending this premise. I do so initially by arguing that they fall within the scope of Wall's version of legal moralism (i.e. they are acts that are harmful to individual moral character). I will return to Duff's version as this argument develops.

I start with a warm-up argument. There is one obvious case in which purely robotic acts of rape and child sexual abuse would be suggestive of a defective moral character, and in which the criminalisation of such acts would be relatively uncontroversial. The case in question is where the acts would fall within the class of *impossible attempts*. As is well known, the criminal law commonly imposes liability on people for attempting but not completing certain criminal offences.⁴¹ There are many reasons for it to do so, including the moral guilt of the attempter, deterrence and the protection of public safety. In addition to this, some jurisdictions⁴² allow for attempt liability even where completion of the criminal offence is impossible. A classic example is the English case of *R v. Shivpuri*.⁴³ There, the defendant was found with a powdered substance on his person (and in his flat) which he (the defendant) believed to be an

⁴¹ In England and Wales, this is governed by the Criminal Attempts Act 1981, which covers all indictable offences. The official definition of an attempt in that act is (section 1) anything which is "more than merely preparatory" of a complete crime.

⁴² Again, England and Wales is one. See Criminal Attempts Act 1981 s.1 (2)

⁴³ [1986] UKHL 2

illegal drug. It turned out on scientific examination that the substance was not an illegal drug but, rather, some harmless vegetable matter. Nevertheless, the defendant was convicted of an attempted drugs offence. The reason being that he thought he was committing a crime, even if in reality he was not. One can imagine similar scenarios arising in the case of purely robotic acts of rape and child sexual abuse. Provided that the robot is sufficiently realistic, or the defendant sufficiently imperceptive, he (or she)⁴⁴ could have sex with a robot under the belief that it really was a child or a non-consenting adult.

I would submit that in such a case the imposition of a criminal sanction would be justified. This is because in such a case the acts in question would be directly expressive of a desire or intent to do something that we already rightly deem to be immoral and worthy of criminalisation. In other words, in such a case the defendant's acts would themselves be expressive of a deficient moral character. This would fit those acts within the terms set by Wall's version of legal moralism.

But that's an easy case. Where the defendant is so mistaken about the nature of the robot that they think it is really a human being, then of course what they do is harmful to their moral character. What happens when they are not making any such mistake? What if they are fully aware of the purely robotic nature of their acts? Can such acts still be harmful to moral character? That's what the wrongness premise is claiming.

Though the argument is less clear-cut, I think there is still reason to believe that such purely robotic acts would involve harm to moral character. In order to set out those reasons, I will draw upon some work done in relation to immoral acts performed in virtual environments (i.e. in video games), particularly the work of Stephanie Patridge.⁴⁵ As it happens, I think there are important differences between the virtual and robotic scenarios, but I think these speak in favour of the criminalisation of robotic acts, not

⁴⁴ Obviously, in many jurisdictions rape is a gender-specific offence (e.g. Section 1 of the English Sexual Offences Act 2003), I add "or she" here since women could commit sex offences with children, or other types of sexual assault.

⁴⁵ The following three articles are my primary sources from her: Patridge, S. "The Incurable Social Meaning of Video Game Imagery" (2010) 13(4) *Ethics and Information Technology* 303-312; "Pornography, ethics and video games" (2013) 15(1) *Ethics and Information Technology* 25-34; and (with Andrew Jordan) "Against the moralistic fallacy: In Defence of Modest Moralism about Humour" (2010) 15(1) *Ethical Theory and Moral Practice* 84-95. The latter, though about humour, also concerns our ethical responses to fictional representations and so is relevant to this debate.

against. First, I must be clear about the distinction between virtual and robotic acts. As I am defining them here, virtual acts take place in a virtual world and are performed through the medium of a virtual character (or avatar); robotic acts take place in the real world and are not performed through the medium of a character or avatar. One complication is the possibility of immersive virtual reality acts. These would involve acts performed in a virtual world without the medium of a character or avatar. They are a borderline case, though (for reasons stated below) I would treat them as being more akin to robotic acts.

Debates about the ethics of virtual acts have taken place for some years. This is not surprising given that video games routinely encourage and reward their players for undertaking actions that we would deem immoral in the real world: theft, torture, violent assault, murder are all part and parcel of some of the most popular video games. For the most part, we do not question the moral propriety of these acts. But sometimes there is something peculiarly disturbing about the virtual representations that makes us pause for thought. Two examples of this, with particular relevance to the issue of robotic rape and child sexual abuse, would be the games *Custer's Revenge* and *Rapelay*.⁴⁶ *Custer's Revenge* was a 1982 game in which the goal was to direct a crudely-pixelated, naked avatar of General Custer to rape a Native American woman who was tied to a stake. *Rapelay* was a 2006 game in which the goal was to stalk and rape a mother and her two daughters. As such, both games encouraged and rewarded players for engaging in acts of virtual rape, which in the case of *Custer's Revenge* had a strongly racist overtone. Both games, unsurprisingly, attracted controversy and were either banned or had their sale restricted (e.g. *Rapelay* was technically only for sale in Japan, though obviously was available elsewhere via the internet). A similarly moralistic attitude has been defended in relation to acts of virtual child sexual abuse.⁴⁷

If we accept the legitimacy of this attitude, my argument would be supported: If we believe that the moralistic attitude is warranted in these case of virtual rape and child sexual abuse, then it seems a small step to say that it is warranted in the case of robotic

⁴⁶ *Custer's Revenge* was released by Mystique Games in 1982. Mystique were famous for releasing games with sexually explicit content. *Rapelay* was released by the Japanese company Illusion in 2006.

⁴⁷ Levy, N. "Virtual child pornography: The Eroticisation of Sexual Inequality" (2002) 4(4) *Ethics and Information Technology* 319-323. Levy was writing in response to a US Supreme Court decision that protected virtual child pornography on the grounds of free speech. The free speech angle may have a similar legal effect on the debate about robots (in the US anyway) but two points should be borne in mind about this. First, it is not clear that the free speech defence is morally persuasive. Second, the differences between video games and robots may count against it.

rape and child sexual abuse. But should we accept the moralistic attitude in the virtual case?

Some philosophers argue that this moralistic attitude is unsustainable, at least in its present form. The point has been put most forcefully by Morgan Luck in his work on the Gamer's Dilemma.⁴⁸ Luck looks at acts of virtual murder and child sexual abuse and notices that many people accept the former but oppose the latter. His Gamer's Dilemma argues that this asymmetric attitude is indefensible. Either both of these things are impermissible or neither of them is impermissible. We can't have it both ways. Luck defends this view by investigating the common rationales for deeming virtual child sexual abuse impermissible (these include arguments about its extrinsic effects and about the moral character of those who perform it). He finds them to be just as persuasive (or unpersuasive, as the case may be) when applied to virtual murder.

One of his more important observations⁴⁹ (at least from our perspective) comes when he discusses the link between virtual acts and moral character. He accepts that someone engaging in acts of virtual child sexual abuse is intuitively likely to be deriving pleasure from those acts⁵⁰ because of an attraction to actual child sexual abuse, but argues that the same could well be true of those committing virtual acts of murder. In both cases, the virtual acts could be directly expressive of a deficient moral character. At the same time, there could be those whose enjoyment of those acts is not directly expressive of a fondness for their real world equivalents, either because they don't think about real world implications or because they have other reasons for enjoying them (e.g. they like the iconoclastic nature of the representations).⁵¹ Thus, there is an epistemic hurdle when it comes to knowing for sure whether the virtual act is directly expressive of an deficient moral character: it may be for some people, but not for others. Luck doesn't think this hurdle is necessarily a defeater for those who think these acts are impermissible (just as it may not be a hurdle if some people are unaffected by addictive drugs when deciding whether consuming those substances should be impermissible). His point is just that if moral character is harmed in the case of virtual child sexual

⁴⁸ See, variously, Luck, M. "The Gamer's Dilemma (2009) 11(1) *Ethics and Information Technology* 31-36; and Luck, M. and Ellerby, N. "Has Bartel Resolved the Gamer's Dilemma?" (2013) 15(3) *Ethics and Information Technology* 29-33.

⁴⁹ See Luck 2009 (n 48), p. 34

⁵⁰ One exception to this would be the case where the player must perform the act as a means to some other end in the game, i.e. where the purpose of the game is not simply to depict those acts. Luck discusses this in his 2009 article (n 48), p. 34

⁵¹ Patridge 2010 (n 45) discusses this possibility.

abuse it is as likely to be harmed in the case of virtual murder. The question is whether these likelihoods make those act types sufficiently wrongful that they fall should fall within the scope of the criminal law?

I could easily bite the bullet when it comes to Luck's dilemma. I could accept that both types of virtual act are sufficiently wrongful (because of their effects on moral character) and I could use that to bolster the case for criminalising robotic acts of rape and child sexual abuse. The problem then is that the argument would have knock-on implications for the criminalisation of robotic murder. Thus, I would be committed to the view that those who build and use robots for the purposes of recreationally "killing" them, could find their conduct apt for criminalisation too. While this is an interesting possibility, and while I acknowledge that the argument I am developing may have that implication, I do not wish to commit myself to that implication. Instead, I wish to pursue the claim that there is something distinctive about virtual (and, by proxy, robotic) acts of *rape* and *child sexual abuse* that make those act-types uniquely and sufficiently wrongful.

This is where Patridge's work becomes significant. She argues that the person who plays and enjoys the sexual acts depicted in games like *Custer's Revenge* or *Rapelay* is providing us with much better evidence of corrupted moral character than the person who plays and enjoys games involving virtual murder. This is because in playing and enjoying those types of game, a player is highly likely to be expressing one of two morally harmful attitudes: (i) a desire for real rape and child sexual abuse (which we just discussed), or (ii) an insensitivity to important social moral issues. It is this latter possibility that is typically absent in the virtual murder case. In defending the likelihood of this insensitivity she appeals to the notion that certain virtual/fictional⁵² representations have an *incorrigible social meaning*, *i.e.* a social meaning that it is not morally reasonable to deny.

In brief outline, Patridge's argument is this:⁵³

⁵² Patridge intends for her argument to cover all fictional representations, including representations in prose, poetry, jokes and other forms of art. I will only discuss the video game case, hence I will stick to the term "virtual). The difficulty with some of the other forms of fictional representation is that only the original author or artist plays a role in creating the representation. In the case of the video game, the player is actively creating the representation every time they play the game. This makes the argument stronger in the case of the video game. For in the video game scenario, the player must take some responsibility for creating the representation with the morally problematic social meaning.

⁵³ This is my reconstruction of her argument. It is not found in her original work.

(4) If a virtual representation has an incorrigible (and morally problematic) social meaning, then there is a limit on the range of reasonable interpretations of that representation.

(5) Certain virtual representations — e.g. the representations of rape and sexual violence in games like *Custer's Revenge* and *RapeLay* — have incorrigible and morally problematic social meanings.

(6) Therefore, there is a limit on the range of reasonable interpretations of these representations.

(7) If a given game player fails to notice the limited range of reasonable interpretations — i.e. if they interpret outside of that range, for example, by enjoying or laughing at these representations — then they must have an improperly developed moral sensitivity.

(8) Except in certain exceptional cases, we are entitled to draw negative inferences about the moral character of a person with an improperly developed moral sensitivity.

(8) Therefore, if a person interprets virtual representations such as those found in *Custer's Revenge* or *RapeLay* outside of the limited range of reasonable interpretations, we are entitled to draw negative inferences about their moral character.

The subsequent conclusion would be that the negative inferences suggest that the person's moral character is harmed through such virtual acts, as is required by Wall's principle of criminalisation. Let's briefly work through the premises of this argument.

Premise (4) appeals to the notion that video games frequently have some connection to the real world and that we cannot help but draw upon that connection when interpreting and playing the games. So, for example, in playing a game like *Grand Theft Auto*, I bring my knowledge of real world crime gangs and their operation to bear on the experience. I will suspend some degree of disbelief, but not all. If the game

stretches credulity too much, I will enjoy it less. I need to believe in the world created by the game designers to truly enjoy the experience. Indeed, in many contemporary games, ultra-realism is a goal of the game designers. They want the connection between the real world and the game-world to be as close as possible. This degree of connection to the real world can give rise to the phenomenon of incorrigible social meanings. These are meanings that attach to representations within the game. These meanings are about the nature of the social world we actually inhabit; they are negative or morally problematic in nature; and they are such that any morally sensitive player would be unable to deny their existence. For example, a game-world depiction of a simianised black person would clearly have a racist social meaning. Anyone who played a game incorporating such depictions, and failed to notice their racist overtones, would exhibit a disturbing level of moral insensitivity. This is because it is not morally reasonable for them to deny the problematic social meaning of the representation. The degree of insensitivity increases the more realistic and clearcut the social meaning is, and the more inappropriate the player's reaction. So, for example, a player who laughs at the depiction of a simianised black person is displaying greater insensitivity than one who is merely indifferent. This is important when it comes to the robotic rape and child sexual abuse cases.

Premise (5) is just making the specific claim that games like *Custer's Revenge* and *Rapelay* contain representations with morally problematic social meanings. This shouldn't be too contentious a claim. Rape is clearly immoral, but due to a range of historical practices and beliefs, it has been normalised in our culture. Consequently, we live in what is called a "rape culture", one which tolerates, condones and perhaps even encourages acts of rape and sexual assault. Virtual representations of rape, particularly ones in which players are rewarded and encouraged for committing acts of rape, emanate from that culture and serve to bolster it. They are morally problematic (in part) for that reason and it is not morally reasonable for players to deny the social meaning that attaches to those acts. To derive pleasure from those representations, or to laugh them off, is to display a worrisome form of moral insensitivity.

Or is it? This gets us to premises (7) and (8), which can be treated as a pair. To understand them we must ask two questions. First, is it really fair to say that in these cases there is a limit on the range of morally reasonable interpretations? In her discussion, Patridge accepts that in certain cases representations of this sort can have

alternative, morally acceptable, interpretations. For example, she suggests that a black artist could use racist imagery to draw attention to the problem of racism. But it is unlikely that exceptions of this sort will apply to the virtual representations in games like *Custer's Revenge* and *Rapelay*, and probably even more unlikely for them to apply to the case of robotic rape and child sexual abuse. It would seem very difficult for someone who created and/or made sexual use of a rape-robot to claim they are drawing attention to the problem of rape culture or child abuse, or trying to correct for those problems. After all, they are not merely indifferent to the child-like or rape-like nature of the act; they are actively enjoying and seeking out those acts because of their paedophilic or rape-like qualities.

The second question is: are we entitled to draw negative inferences about the moral characters of those who fail to appreciate the morally problematic social meanings of those representations? I would say that we can definitely draw inferences about their lack of moral sensitivity — since that is a purely factual claim — but whether we can pass moral judgment on them is a separate, more complicated issue. It will depend on our views about blameworthiness and innocence. Premise (8) is designed to acknowledge the complexity at stake here. It is possible that people will be blameless for their moral insensitivity and that we shouldn't pass moral judgment on them. But that shouldn't detract from the fact that, in many cases, we will be able to pass judgment and that this is what is important from the perspective of criminalisation. Individual excuses can be dealt with on a case-by-case basis, as they always have been by the law; what matters when it comes to deciding on criminalisation is the wrongness of the general act-type.

It should be noted that in appealing to the concept of incorrigible social meaning, Patridge's argument could apply to other types of virtual representations. For instance, some types of virtual murder could be covered particularly if they include racist elements (lynching or genocide). Nevertheless, representations of sexual crimes, given the social history of sexual oppression and the use of sexual assault to reinforce that oppression, are particularly good examples of representations with morally problematic social meanings.

Understood in this manner, Patridge's argument places the performance of such virtual acts within the scope of Wall's principle of criminalisation. Since the acts

express or implicate deficiencies in moral character, that seems straightforward. But the argument also places the acts within the scope of Duff's principle of criminalisation. After all, through the concept of incorrigible social meaning, Patridge is pinpointing a set of wrongs that are quintessentially public in nature. If an action or representation has a morally problematic incorrigible social meaning, it is because it ignores or denigrates some moral problem that society had or continues to have (e.g. sexual inequality or racial injustice). If someone is insensitive to that social meaning, then that person is failing to relate to the society around them in the appropriate manner. It is plausible then to argue that what they are doing should be a matter of public concern and should fall within the terms of Duff's principle.⁵⁴

But all this is to speak of virtual representations and virtual acts. What about the purely robotic acts? My point would simply be that the arguments about moral character and moral insensitivity apply *a fortiori* to the case of purely robotic acts. In other words, I would claim that those who engage in such acts either (a) directly express a deficient moral character because they desire real-world rape and child sexual abuse or (b) demonstrate a socially problematic form of moral sensitivity.

In the case of the video game there is some physical and mental distance between the game player and the virtual representation of what they are doing.⁵⁵ This distance will always make it more difficult to draw inferences about the player's actual moral character. In the case of purely robotic acts of rape and child sexual abuse, this distance is massively reduced and it is consequently easier to make those inferences. We can be much more confident then that the general act-types involve harm to moral character, or include wrongs done to the public at large, than we can be in the virtual case. It will help if we develop this argument by responding to a more nuanced version of the "distance" criticism.

This more nuanced version comes from the work Gert Gooskens.⁵⁶ He claims that we can never say that a virtual act is "right" or "wrong" because there is something about the phenomenology of virtual acts that prevents the legitimate application of those moral predicates. He relies on Husserl's concept of the "image-world" when making

⁵⁴ Duff himself seems to accept this when he acknowledges that forms of hate speech or extreme pornography would fall under the scope of his principle. See Duff (n 25) p. 232.

⁵⁵ An exception to this would be the case of truly immersive virtual reality, such as that being developed by Oculus Rift. I would view that as being equivalent to the purely robotic case I discuss in the text.

⁵⁶ Gooskens, G. "The Ethical Status of Virtual Actions" (2010) 17(1) *Ethical Perspectives* 59-78

this claim. For Husserl, the image-world refers to any representation, real or fictional. The phenomenology of the image-world refers to our subjective engagements with such representations. According to Husserl, these subjective engagements always work with an “as-if” modifier in place. That is to say, whenever we engage with an image we do so with the subjective awareness that it is not really what it represents itself as being. It is only “as-if” it is that thing. For example, if I am looking at a photograph of my girlfriend, I know that I am not actually looking at her, it is only “as-if” I am looking at her. Gooskens argues that all our engagements with virtual acts and representations involve this kind of “image consciousness”, i.e. the awareness that we are doing things through an image that is not real, but only “as if” it is real.

The key Husserlian move in all this is to argue that the “as-if” modifier doesn’t simply apply to our perception of the image-world; it applies to all our other mental engagements with the image world. Thus, if I intend to do things through an image — as is the case in the video game — then I do not really intend the actions that are represented. It is only “as-if” I intended those things. The presence of this “as-if” modifier is what, for Gooskens, blocks the application of moral predicates like “right” or “wrong”. It suggests that what the player doing isn’t truly morally harmful (to themselves), because with the image-world there is always separation between the real self and the as-if self.⁵⁷

Even in the virtual realm, this argument can be challenged. The reliance on the Husserlian model could, for example, be challenged on both theoretical and empirical grounds. Theoretically, there is no reason why mental responses to the image world, e.g. emotional responses, would be any less real than they would be to the real world equivalents. If I cry in response to a harrowing painting of a starving child, it seems wrong to say that my emotional response is not real but merely “as if” it were real. I am still *really* crying. Furthermore, empirically, there is evidence to suggest that certain types of virtual experience can be disconcertingly real.⁵⁸

⁵⁷ Gooskens is, arguably, inconsistent on this matter since he acknowledges that moral discomfort arises whenever there is a breakdown between the real world and the image world. In acknowledging the possibility of such a breakdown, he may be undermining his own argument.

⁵⁸ See Metzinger, T. *The Ego Tunnel* (New York: Basic Books, 2010) for a description of experiments involving virtual reality domains which can lead people to completely occupy a false first-person perspective.

But I need not pursue those challenges here. Instead, I will simply insist that the phenomenology of purely robotic acts is different from the phenomenology of virtual acts. The purely robotic act does not involve action through an image or other medium. It involves direct physical contact with an artificial being. That means the personal distance between the act and one's true intentions or desires is lessened in this case. It is not as-if you are someone else when performing these acts; you are yourself. Hence we can be more certain of the inferences we draw about the moral character of the person who engages in those acts.

Of course, a proponent of Gooskens's argument could step in at this point and argue that the artificiality of the robot is what matters, not the action through an image. This artificiality still allows you to create moral distance between yourself and the seemingly immoral act. Two points in response to this. First, this still wouldn't undermine the Partridge-style argument about insensitivity to the incorrigible social meaning of what one is doing. Second, there is evidence to suggest that it is much more difficult for people to perform immoral acts with physical representations of real things than it is for them to think about those things abstractly or perform them in a virtual environment. This suggests that moral faculties are much more deeply engaged when performing acts with artificial beings. It is worth considering this evidence, briefly, in order to round out the argument for the wrongness premise.

The evidence must be appreciated in its context. In recent years, a variety of studies have been done on human moral reasoning. One of the key insights from these studies is that human moral reasoning follows a "dual track" process. One track issues quick, intuitive responses to moral decision problems; another is more reflective and calculating.⁵⁹ Philosophers argue over the implications this has for our understanding of morality. The classic model of moral analysis tends to be heavily intuition-driven and so reliant on the "fast" track; but there are some who argue that knowledge of the evolutionary history of that fast track should cause us to doubt those intuitive judgments.⁶⁰ I cannot adjudicate on that debate here, but I for the purposes of my argument I will assume that the classic model is more-or-less correct. Consequently, I will assume that the fast track functions in a particular way in people of sound moral character.

⁵⁹ Greene, G. *Moral Tribes* (London: Penguin, 2013) contains a good overview of the work that has been done in this field.

⁶⁰ See, generally, Kahane, G. "Evolutionary Debunking Arguments" (2011) 45(1) *Nous* 103-125

This is by way of background. The evidence I am interested in comes from a study looking at how people responded when asked to perform immoral acts with artificial objects.⁶¹ For example, smashing a fake baby's head off a table, or striking another person on the knee with a fake hammer. The study found that, even though people knew that the baby was not real, or that the hammer could not do any real damage, they found it very difficult to do so. The suggestion being that their intuitive, fast-track, system of moral judgment resisted the actions. This resistance could only be overridden by the slower, more detached system of reasoning.⁶²

My belief is that this study has something of interest to say about the cases of robotic rape and child sexual abuse. More precisely, I believe it suggests that those who perform and enjoy performing purely robotic acts of rape and child sexual abuse must either (a) have an inherently defective system of intuitive moral judgments or (b) have worked to repress or overcome the intuitive resistance to such acts. In either case, there is reason to think that actions of this type are harmful to their moral character and are hence apt for criminalisation. For these reasons, I would say that arguments about the virtual realm apply *a fortiori* to the robotic realm.

A final comment is in order about social meaning and sexual objectification.⁶³ Robotic sex (of all kinds) might be deemed to have a problematic social meaning insofar as it expresses the purest form of sexual objectification. If the robots are just mere objects (and I am assuming throughout this article that they are not persons), then all forms of robotic sex involve the reduction of sexual partners to objects.⁶⁴ If sexual objectification of this sort has a problematic social meaning, and if insensitivity to this meaning is indicative of a deficient moral character, one might complain that the argument I have presented is too narrow and could be expanded to cover all forms of robotic sex. This could either amount to a *reductio* of my argument (in that it now seems to warrant too much criminalisation) or a reason to expand the scope of criminalisation.

⁶¹ Cushman, F., Gray, K. et al "Simulating murder: the aversion to harmful action" (2012) 12(1) *Emotion* 2

⁶² Greene (n 59) also discusses the study, p. 36 and p. 228

⁶³ I am indebted to an anonymous reviewer for encouraging me to make this point.

⁶⁴ With a certain amount of stretching all seven of Nussbaum's concepts of objectification would all seem to apply to the robotic sex case, i.e. instrumentality, fungibility, violability, inertness ownership, denial of autonomy, and denial subjectivity (though the last two are tricky given that there isn't really a "denial" taken place). Nussbaum, M. "Objectification" (1995) 24(4) *Philosophy and Public Affairs* 249

While there may be some who wish to push for that expanded scope, I think my narrower argument is still to be preferred. There are several reasons for this. First, the wrongness of objectification is itself open to doubt. As many have pointed out, the objectification of persons is not always morally wrong: I may use another human for a particular end without problematically denying their agency or autonomy (for example, they may consent to a degree of objectification).⁶⁵ The narrower argument applies even if you accept this point. Second, and perhaps more importantly, it is not clear that the objectification concern is relevant when it comes to robots anyway. If robots simply are objects, then surely there is nothing morally problematic in using them as such (just as there is nothing morally problematic about using a hammer as an object). I accept, however, that this may not do full justice to the nuances of the robot case. Some have argued that robots blur the object-agent distinction and may, hence, belong to a unique ontological category.⁶⁶ If this is so there may be something “iffy” about using robots in particular ways. What I want to suggest, however, is that this iffiness would come from other aspects of that usage, not just from the mere objectification. In particular, I want to suggest that if they are used in a way that generates a distinctly problematic social meaning, or expresses other (non-objectification related) moral defects, this should be a cause for concern. This is what the argument I outlined above claims. Using an object (or something in between an object and an agent) in a way that explicitly contributes to or condones rape culture, or expresses dubious sexual desires, warrants greater concern than, say, using an object to replicate a loving, mutually advantageous sexual relationship. This is another reason for preferring the narrower argument.

Summary of the Case for the Wrongness Premise

This discussion has taken a number of twists and turns. It will be helpful if, before we proceed, we can summarise the basic case for the wrongness premise. That case makes two main claims. First, that those who engage in acts of purely robotic rape and child sexual abuse either demonstrate an immoral desire for the real world equivalents of those acts, and/or an disturbing moral insensitivity to the social meaning of those acts. These two arguments are derived from work done on the ethics of virtual acts (*i.e.* acts performed within video games), but they apply *a fortiori* to robotic acts. They do so

⁶⁵ See Marino, P “The Ethics of Sexual Objectification: Autonomy and Consent” (2008) 51(4) *Inquiry* 345-364 for a consent-based defence of some objectification.

⁶⁶ Calo, R. “Robotics and the Lessons of Cyberlaw” (2015) 103 *California Law Review* 102, p. 133, citing Peter Kahn.

because there is a much tighter connection between moral character and robotic acts than there is between moral character and virtual acts.

If this basic argument is correct, then it highlights how purely robotic acts of rape and child sexual abuse are wrong in a way that makes them apt for criminalisation under either Wall or Duff's principles of criminalisation. They are apt under the former because they suggest that purely robotic acts of rape and child sexual abuse involve harm to moral character; and they are apt under the latter because they involve insensitivity to the social meaning of those acts and so naturally fall within the domain of public wrongs.

With the wrongness premise defended, my initial *prima facie* case for criminalisation is complete. The question is whether there are any defeaters for that *prima facie* argument. In what remains, I want to consider two families of defeaters. The first will try to undermine the argument by focusing on the possible extrinsic effects of rape-bot or child-bot usage. The second will focus on problems arising from the practical implementation of any criminalisation policy.

6. What about extrinsic effects?

So far, I have focused on the purely robotic case — *i.e.* the case in which there are no extrinsic effects of robot usage — but presumably engaging in sexual activities with a robot will have *some kind* of extrinsic effects.⁶⁷ Could it be that those extrinsic effects defeat the *prima facie* argument? To answer that question, it will help to narrow and classify the possible extrinsic effects.

For the purposes of this article, the most relevant types of extrinsic effect would be those that relate to other kinds of criminal activity, specifically real-world rape and child sexual abuse. Why? Because it would seem like only those kinds of effects would have sufficient weight to cancel the *prima facie* argument. This isn't to deny that engaging in acts of robotic rape and child sexual abuse could have other extrinsic effects — e.g. they could make you happier or sadder, they could improve or damage your relationship with other human beings, or they could cause you to withdraw from

⁶⁷ At an absolute minimum, there would be the opportunity cost effects, *i.e.* spending time having sex with a robot prevents one from doing other things.

the real-world — but these would not be sufficiently good to alter the criminalisation argument.

With this in mind, there are, broadly speaking, three possible extrinsic effects that are of interest. The first would be that engaging in acts of robotic rape and child sexual abuse *significantly increases* the likelihood of someone performing their real-world equivalents. The second would be that engaging in acts of robotic rape and child sexual abuse *significantly decreases* the likelihood of someone performing their real-world equivalents. And the third would be that it has *no significant effect* or *an ambiguous effect* on the likelihood of someone performing their real-world equivalents.

If the first possibility turns out to be true, then we would simply bolster the case for criminalisation. For then, not only would we have intrinsic reasons for deeming the conduct apt for criminalisation, we would also have some extrinsic reasons. These extrinsic reasons would not only fit with Wall and Duff's principles of criminalisation, they would also fit with the more classic liberal ones. For if there is a significant increase in the risk of real-world rapes or real-world acts of child sexual abuse, there are grounds for criminalising in terms of potential harm to others and protection of the public.

If the second possibility turns out to be true, the situation would be rather different. Then we would have an extrinsic reason for dropping the case for criminalisation. The intrinsic reasons for objecting to robotic rape and child sexual abuse would not be negated (*i.e.* the premises and conclusions of the arguments I offered would still stand) but they would be outweighed by the extrinsic benefits. This is clearly and obviously true. Since we are assuming there is no moral victim in the robotic case, real-world rape and child sexual abuse are obviously much worse than their robotic equivalents. Any reduction in the risk of real-world acts would be worth it. I would suggest that the possibility of such a reduction should be actively and carefully researched. It would have important benefits for society, and also for those in the grip of certain forms of paraphilia. People already advocate much more drastic therapies for such people — *e.g.* chemical castration⁶⁸. The permitted use of a sex robot would be less invasive, and arguably preferable.

⁶⁸ See the discussion in Douglas, T. et al "Coercion, Incarceration, and Chemical Castration: An Argument from Autonomy" (2013) 10 (3) *Bioethical Inquiries* 393-405.

If the third possibility turns out to be true, then we are as we were. The *prima facie* argument would still stand, undefeated by extrinsic considerations. In fact, there may even be a weak reason to think that in the case of ambiguous or inconclusive evidence, we actually strengthen the case for criminalisation, particularly when it comes to robotic rape.

Leslie Green has recently defended the view that law can, and perhaps should, be used to improve social morality (*i.e.* the set of social beliefs and dispositions about what kinds of behaviour are permitted, forbidden and obliged).⁶⁹ And he uses social beliefs about rape and consent to sex as an example. He notes:⁷⁰

“...our social morality of sex is multiply defective. And this is not news. Whether or not our societies actually embody a “rape culture” [footnote omitted], their moralities are flawed in ways that are obviously sexist and, because of that, significantly limit the power of women to plan their lives. The asymmetrical character of the flaws is important... Women who have not been, and may never be, victims of rape nonetheless modify their lives for fear of it, or because their parents fear it: they feel obliged to avoid certain places, to be careful going out at certain times, to give thought to how their dress and manner will be received — in general, to remain unwillingly vigilant.”

To correct for this defective social morality, Green suggests that law can be used to encourage men to seek unambiguous, clear signals of consent. The law can do this in a number of ways. Reforming rape law and practice is one obvious route, making it clear to men that making assumptions about the sexual consent of women in the absence of clear signals will not be tolerated. But reforms to other areas of law could also be important:⁷¹

“...thinking that only rape law stands a chance of shaping attitudes to rape is no better than jurisprudential homeopathy. In fact, it may be that the most influential laws in shaping our social morality here will not be the laws of rape at all. For example, since consent functions through performatives, other laws could protect

⁶⁹ Green, L. “Should Law Improve Morality?” (2013) 7 *Criminal Law and Philosophy* 473-494

⁷⁰ *Ibid*, p. 491

⁷¹ *Ibid* p. 492

the 'felicity conditions' for those acts. They could help ensure, for instance, that refusals are taken seriously, and are actually heard as refusals."

My contention is that the criminalisation of robotic rape (specifically) could form part of this package of laws that help to improve social sexual morality. In the first instance by banning the manufacture and/or criminalising the use of rape robots. And, perhaps, in the second instance by ensuring that those sex robots that are made cannot be used for sexual purposes until they have given clear, unambiguous signals of consent.

Some might worry that this argument assumes the very thing we have just assumed not to be the case, namely: that there is some robust connection between robotic acts and real-world acts. But that is the wrong way to look at the argument. The argument is about how we should deal with the possible risks to social sexual morality. The idea being mooted is that we should not risk doing anything that could exacerbate our problems with sexual consent, even if there is no strong evidence suggesting that the use of rape robots increases or decreases the likelihood of real-world rape. In particular, we should avoid permitting and condoning the use of sex robots that do not take signals of non-consent seriously.

All this is to speak in purely abstract and hypothetical terms. It might be useful if we could review some actual empirical evidence on this matter. But we don't have anything concrete yet because sex robots are in their infancy. That means we should proceed with caution, accepting that anything we say about extrinsic effects is open to doubt and revision. With that in mind, I think the work that has been done on the link between pornography and real-world sex crimes has some lessons for this debate. Pornography is sexually explicit media that is used for sexual stimulation and release. In this sense, it is analogous to the use of sex robots. Is there some link between the consumption of pornography and other harmful extrinsic effects that could provide guidance to the debate about robot sex? To better understand the issues, we need to distinguish between varieties of pornography, particularly between misogynist pornography (which is sometimes argued to give rise to sexual violence against women), and child pornography (consumption of which is usually deemed intrinsically wrong, and possibly a contributor to real-world instances of child sexual abuse).

Focusing first on misogynist pornography, one of the notable things about research into its real world effects is that it is something of a mixed bag. There are several studies showing a correlative link between pornography use and increased tolerance for sexual aggression,⁷² and there are vehement defenders of this view.⁷³ There are, however, other studies suggesting that the link is minimal,⁷⁴ or dependent on personality type or prior attitudes toward sexual aggression.⁷⁵ Furthermore, there are studies that find no positive correlation between pornography use and sex crimes, and possibly even a negative correlation.⁷⁶

The research on links between child pornography and real-world acts of child sexual abuse is similarly mixed. For example, the work of Bourke and Hernandez and Ost supports the view that usage of child pornography increases the risk to real-world children.⁷⁷ But others argue that the link does not hold up to scrutiny. For example, several studies show that possession of child pornography is a minimal risk factor for future incidents of child sexual abuse.⁷⁸ And there are those, like Carissa Hessick, who have argued that the law is far too quick to assume a connection between child pornography possession and child sexual abuse.⁷⁹

⁷² For example: Malamuth, NM and Chech, J.V.P “The Effects of Mass Media Exposure on Acceptance of Violence Against Women: a Field Experiment (1981) 15 *Journal of Research in Personality* 436-446; Zillman, D. and Bryant, J. “Pornography, Sexual Callousness and the Trivialization of Rape” (1982) 32 *Journal of Communication* 10-21; Zillman, D. and Bryant, J. “Pornography and Men’s Sexual Callousness Toward Women” in Zillmann and Bryant (eds) *Pornography: Research Advances and Policy Considerations* (Hillsdale, NJ: Erlbaum, 1989); and Allen et al “Exposure to Pornography and Acceptance of Rape Myths (1996) 45 *Journal of Communication* 5-26

⁷³ Dines, G. *Pornland: How Porn has Hijacked our Sexuality* (Beacon Press, 2011)

⁷⁴ Barak, A. and Fisher, W.A. “Effects of Interactive Computer Erotica on Men’s Attitudes and Behavior toward Women: an Experimental Study” (1997) 13 *Computers in Human Behaviour* 353-369; and Barak, A., Fisher, WA, Belfry, S. and Lashambe, D. “Sex, Guys and Cyberspace: Effects of Internet Pornography and Individual Differences on Men’s Attitude Toward Women” (1999) 11 *Journal of Psychology and Human Sexuality* 63-91.

⁷⁵ Malamuth, NM, Addison, T. and Koss, M. “Pornography and sexual aggression: are there reliable effects and can we understand them?” (2000) 11 *Annual Review of Sex Research* 26-91; and Hald, GM, Malamuth, NM and Lange, T “Pornography and Sexist Attitudes Among Heterosexuals” (2013) 63 *Journal of Communication* 638-660

⁷⁶ Ferguson, CJ and Hartley, RD “The Pleasure is momentary...the expense damnable? The influence of pornography on rape and sexual assault” (2009) 14 *Aggression and Violent Behavior* 323-329; and Diamond, M. “Pornography, Public Acceptance and Sex Related Crime: A Review” (2009) 32 *International Journal of Law and Psychiatry* 304-314.

⁷⁷ Bourke, ML and Hernandez, AE “The ‘Butner Study’ Redux: A Report of the Incidence of Hands-on Child Victimization by Child Pornography Offenders” (2009) 24(3) *Journal of Family Violence* 183-191; Ost, S. “Children at risk: legal and society perceptions of the potential threat the possession of child pornography poses to society” (2002) 29 *Journal of Law and Society* 436-460.

⁷⁸ Endrass, J et al “The Consumption of Internet Child Pornography and Violent and Sex Offending” *BMC Psychiatry* 14 July 2009 doi: 10.1186/1471-244X-9-43; and Seto, MC, Hanson, RK and Babchishin KM “Contact Sexual Offending by Men with Online sexual offending” (2011) 23(1) *Sex Abuse* 124-145 (suggesting that there may be a distinct, low-risk, class of online offender)

⁷⁹ Hessick, CB “Disentangling child pornography from child sex abuse” (2011) 88 *Washington University Law Review* 853-902

If this research is anything to go by, then the third possibility (ambiguous extrinsic effects) may be the most probable one. This would leave the prima facie argument undefeated or, if we accept Green's approach, add another weak reason for criminalisation. However, there are alternative perspectives on the debate about the extrinsic effects of pornography.⁸⁰ Karen Boyle, for instance, has argued that the "effects discourse" is unhelpful.⁸¹ Instead of focusing on the effects of pornography (of either type), she argues that we should focus on its production and consumption, and how that production and consumption may harm or contribute to the abuse and oppression of women and children. There are others too, like AW Eaton, who argue that a more diverse set of potential harms, and a probabilistic concept of causation, should be adopted in this debate.⁸² Could those alternative perspectives have something to say about the extrinsic effects of robotic rape and child sexual abuse? Perhaps, but once again I suspect that they would only tend to support the criminalisation argument. Boyle focuses on cases involving moral victims (e.g. real women and children who may be harmed in the production process). Obviously, I have explicitly avoided the assumption that robots could be moral victims. So at first glance her argument seems inapplicable. But in encouraging us to focus on the production and consumption phases, and the harmful properties of those phases, she encourages something very similar to what I have done in previous sections of this article, which was to highlight the intrinsically harmful aspects of producing and using robots for replicating acts of rape and child sexual abuse. The same is true of Eaton's argument. By encouraging us to focus on a broader and more diverse set of harms and a looser causal link, she provides more reasons to be concerned about the production and consumption of misogynist pornography. Such a broader focus may also lend support to the criminalisation of robotic rape and child sexual abuse.

But, of course, all this is to reason by analogy. There are differences between the use of pornography and the use of sex robots. Those differences could strengthen the link between sex robot usage and real-world effects, though in which direction remains to be seen. This is something that requires much closer scrutiny in the future.

⁸⁰ I am indebted to an anonymous reviewer for encouraging these additional reflections.

⁸¹ Boyle, K *Media and Violence* (Sage Publications, 2005), chapter 2; and Boyle, K. "The Pornography Debates: Beyond Cause and Effect" (2000) 23(2) *Women's Studies International Forum* 187-195.

⁸² Eaton, AW "A Sensible Antiporn Feminism" (2007) 117 *Ethics* 674-715

7. What about the costs of implementation?

Another possible set of defeaters to the *prima facie* argument would concern its practicality. Even if we could agree that acts of purely robotic rape and child sexual abuse were morally wrong, and that their wrongness made them apt for criminalisation, we might worry about the effects of implementing a criminal prohibition. Would the moral costs of such a prohibition outweigh the moral benefits? Would it be wrong to use the full weight of the criminal justice system to address this issue? To answer those questions, it is worth thinking about the possible targets of criminalisation; the way in which breaches of the law could be detected; and the form of punishment (if any) that might be brought to bear on those who break the law. When we think about these issues in more detail, we see that criminalisation need not come with excessive practical/moral costs.

Let's start by thinking about the possible targets of any criminalisation policy. The argument in preceding sections focused on the wrongness of the actual acts of robotic rape and child sexual abuse. This suggests that the prime targets for criminalisation would be the users of those sex robots, *i.e.* the people who perform the wrongful acts. But we might worry that targeting users is too difficult and too authoritarian: they might be a diverse, hard-to-identify group of people and we might even have some sympathy for them if they suffer from forms of paraphilia or other sexual disorder. However, this would not necessarily block the case for criminalisation. As with other crimes involving personal "vice", suppliers and distributors could also be targeted on the grounds that they facilitate and encourage the wrongful acts. Indeed, we might exclusively or preferentially target them, as is now done for illicit drugs in many countries.

But even if we could agree on a set of targets, we might worry that identifying breaches of the law would be too costly. To identify users, for example, would seem to require a considerable invasion of privacy. Presumably people would use sex robots in their homes, away from the public eye. To successfully enforce a prohibition on use would therefore require some significant surveillance and monitoring of such private spheres. If we value privacy, this cost may outweigh the benefit of criminalisation. But, of course, we already live in a world in which privacy is massively compromised and round-the-clock mass surveillance is easily implemented. The heavy reliance on the internet and the increasing connectivity of the world around us, have taken care of

that.⁸³ If sex robots are among the devices with internet connectivity — and given that this is essentially the norm nowadays, it seems likely that this would be the case⁸⁴ — then monitoring and surveilling their usage may not be impractical. This may not be a good thing, due to the privacy invasion it entails, but we are arguably going through a reevaluation of privacy anyway. It may be that our awareness and tolerance of increasing surveillance and monitoring means we care less about privacy than we used to. Such a reevaluation would undercut this defeater to the *prima facie* argument. In any event, this defeater assumes that users will be the primary targets of criminalisation. If it is manufacturers and suppliers, the privacy concerns would be considerably reduced. We already tolerate heavy (criminal) regulation of the manufacture and supply of various products.

Finally, we turn to the question of how those who breach the law are to be treated. I would readily admit that throwing users and manufacturers in jail would be heavy-handed, and probably disproportionate to the wrongness in question (this may depend, however, on volume of usage and evidence as to extrinsic effects). But incarceration is not the only possible response to a breach of the criminal law. Fines could be adopted for manufacturers and users. Those caught using rape or child sexual abuse robots could be required to undergo some form of sex therapy or education. Models of this sort are already in place for drug offences in different jurisdictions. The availability of such options would allow the *prima facie* argument to stand. Furthermore, such treatments/punishments are consistent with the underlying principles of criminalisation. After all, addressing the harm to moral character, or requiring public condemnation and accountability for certain wrongs, does not require harsh treatment or harm to be inflicted upon the perpetrators of those wrongs.

Again, the point here is not to endorse any particular model of criminalisation. The point is merely to show that the implementation of a criminalisation policy need not be excessively costly.

⁸³ The obvious example of this comes from the Edward Snowden NSA leaks, and the revelations about data collection by US government agencies. For an overview, see Greenwald, G. *No Place to Hide* (Metropolitan Books, 2014) and Harding, L. *The Snowden Files* (Guardian Faber, 2014)

⁸⁴ Obviously, the risks of surveillance might create a black market in devices that don't have this connectivity. However, this wouldn't necessarily improve things. Users would need to make sure that they leave no electronic signature of their purchase or use of such objects (i.e. no text messaging about the devices, no electronic purchases, no pictures posted to online forums). This is extremely difficult to achieve.

8. Conclusion

Sex robots are coming, whether we like it or not. And there will, no doubt, be a temptation to create sex robots that cater to paedophilic tastes and rape fantasies. The question pursued in this article was whether such devices should be regulated by the criminal law. Although I remain tentative and sceptical, I have tried to present a plausible, *prima facie* argument for thinking that they should be.

The argument was based on two premises. The moralistic premise followed some recent work done on principles of criminalisation and proposed that acts that are morally wrong can be apt for criminalisation. Specifically, it followed work from Steven Wall and Anthony Duff which held that acts that are harmful to moral character, or which amount to public wrongs, can be apt for criminalisation. The second premise of the argument — the wrongness premise — then argued that those who engage in purely robotic acts of rape and child sexual abuse (i.e. acts with no extrinsic effects) could find that their actions fall within the scope of Wall and Duff's principles of criminalisation. They could fall within the scope of the former because there are reasons for thinking that such usage harms moral character; and they could fall within the scope of the latter because enjoyment of such acts evinces a disturbing insensitivity to important social moral issues.

This *prima facie* argument could be defeated by concerns about the costs of implementation and the possible extrinsic effects of such sex robot usage. I have, however, tried to show how difficult it could be to find successful defeaters. Nevertheless, and in conclusion, I want to highlight once more the importance of obtaining good evidence relating to the extrinsic effects of using such sex robots. If robust and significant evidence was found suggesting that such usage decreased the incidence of real-world sex crimes, then we would have good reason not to criminalise it. This is something that should be carefully researched in the future. This is just the beginning of the debate. It is important it does not get too far ahead of reality.

Acknowledgements: The author would like to thank Neil McArthur, Ezio di Nucci and two anonymous reviewers for feedback on a previous draft of this paper.