

Regulating Child Sex Robots: Restriction or Experimentation?

By John Danaher, NUI Galway

(pre-publication draft of a paper forthcoming in *Medical Law Review* special issue on ‘Regulating Sexual Boundaries’ 2019)

Abstract: In July 2014, the roboticist Ronald Arkin suggested that child sex robots could be used to treat those with paedophilic predilections in the same way that methadone is used to treat heroin addicts. Taking this onboard, it would seem that there is reason to experiment with the regulation of this technology. But most people seem to disagree with this idea, with legal authorities in both the UK and US taking steps to outlaw such devices. In this paper, I subject these different regulatory attitudes to critical scrutiny. In doing so, I make three main contributions to the debate. First, I present a framework for thinking about the regulatory options that we confront when dealing with child sex robots. Second, I argue that there is a prima facie case for restrictive regulation, but that this is contingent on whether Arkin’s hypothesis has a reasonable prospect of being successfully tested. Third, I argue that Arkin’s hypothesis probably does not have a reasonable prospect of being successfully tested. Consequently, we should proceed with utmost caution when it comes to this technology.

1. Introduction: Thinking the Unthinkable

Sometimes you have to think the unthinkable. On the 13th of June 2018, Dan Donovan, a US Congressman from 11th District of New York, sponsored and successfully passed legislation through the US House of Representatives banning the importation, distribution and sale of child sex robots and child sex dolls in the USA.¹ Donovan’s so-called CREEPER² Act is premised on the belief that the use of such devices ‘not only lead[s] to rape but...make[s] rape easier’ and tends to ‘normalize sex between adults and minors’.³ In

¹ For information see ‘House Passes Donovan Creeper Act Ban’, press release from the Office of Dan Donovan, available at <https://dan.donovan.us/congress/legislation/press-releases/donovan-s-creeper-act-ban>

² For the text of the Act, see <https://www.gpo.gov/fdsys/pkg/BILLS-115hr4655rfs/pdf/BILLS-115hr4655rfs.pdf> (visited on the 23/7/2018)

crafting this legislation, Donovan was not the first to try to restrictively regulate the existence of such devices. In 2017, UK Crown Prosecution Services, worried about the emerging online trade in child sex dolls, decided to clamp down on their importation to the UK. Using a somewhat obscure provision of the 1876 Customs Consolidation Act — a provision that banned the importation of ‘obscene’ or ‘indecent’ objects⁴ — they successfully prosecuted several individuals for purchasing these items online and having them shipped to the UK.⁵ The CPS felt moved to bring these prosecutions because of what they perceived to be a link between the purchase of these devices and other child sex offences.⁶

The idea of child sex robots is unthinkable. But as the actions of Dan Donovan and the CPS clearly demonstrate, some people have started to think about them very seriously indeed. Donovan and the CPS were acting in a much broader context, one in which companies are now manufacturing all manner of sex robots and dolls, and academics are starting to debate and investigate their ethical and social implications.⁷ Consistent with the attitude of Donovan and the CPS, most of the academic discussion of child sex robots, to date, has been negative in its appraisal: their very existence is generally viewed as a bad thing, and not something to be encouraged. Restrictive regulation, up to and including criminalisation of their manufacture, distribution, and use is generally thought to be warranted.⁸ In other words, ‘restrictive regulation’ represents an emerging consensus in the debate about child sex robots.

In this article, I take a look at this emerging consensus and try to determine whether it is, in fact, justified. Could there ever be a case for a more open approach to the regulation of

⁴ Section 42 of the Customs Consolidation Act 1876

⁵ For the relevant press releases from the CPS see: <https://www.cps.gov.uk/south-east/news/prison-sentence-child-sex-doll-importation>; <https://www.cps.gov.uk/south-west/news/plymouth-man-admitted-importing-childlike-sex-dolls> (visited on 23/7/2018)

⁶ Ibid

⁷ Richardson, K., The ‘Asymmetrical’ Relationship: Parallels Between Prostitution and the Development of Sex Robots (2015) 4(3) SIGCAS Computers & Society 290–293; Richardson, K. Sex Robots: The End of Love (London: Polity Press, 2018); Strikwerda, L. Legal and Moral Implications of Child Sex Robots, in Danaher, J. and McArthur, N. (eds) Robot Sex: Social and Ethical Implications (Cambridge, MA: MIT Press 2017); Sparrow, R. ‘Robots, Rape and Representation’ (2017) 9(4) *International Journal of Social Robotics* 465–477; Danaher, J. Robotic Rape and Robotic Child Sexual Abuse: Should they be criminalised? (2017) 11(1) *Criminal Law and Philosophy* 71–95; Danaher, J. and McArthur, N. (eds) Robot Sex: Social and Ethical Implications (Cambridge, MA: MIT Press, 2017); Devlin, K., In Defence of Sex Machines: Why Trying to Ban Sex Robots is wrong, *The Conversation* 17 September 2015, available at <https://theconversation.com/in-defence-of-sex-machines-why-trying-to-ban-sex-robots-is-wrong-47641>; Devlin, K. Turned On: Science, Sex and Robots (London: Bloomsbury Press, 2018); Maras, MH and Shapiro, LR, Child Sex Dolls and Robots: More than Just an Uncanny Valley (2017) *Journal of Internet Law*, December 2017, 4–21.

⁸ Danaher 2017 (n 7); Strikwerda (n 7); and Maras and Shapiro (n 7) all argue for some form of restrictive regulation.

such devices? I will answer this question in four steps. First, I will develop a framework for thinking about the different regulatory approaches we could take toward child sex robots and child sex dolls (I will also clarify the terminology used to describe such devices). This framework will be grounded in legal moralism, a position that sometimes gets a bad rap but is, in my view, wholly defensible. Second, I will use this framework to evaluate the case for restrictive regulation of child sex robots. I will argue that there are some good reasons for thinking that restrictive regulation is warranted, but that the strength of these reasons is contingent on the absence of a prima facie case for thinking that child sex robots could be used for therapeutic purposes. The presence of such a prima facie case might warrant a more open stance. Third, I consider whether such a prima facie case can be made. I argue that this is pretty difficult, for both theoretical reasons (i.e. what we know about the mechanisms of human psychology) and epistemic reasons (i.e. what we know about the limitations of science). Fourth, even though I think we should be sceptical about the therapeutic case for child sex robots, I accept that there is significant uncertainty on this issue and so this might warrant an ‘experimentalist’ approach to the regulation of such devices. I close the article by considering what this experimentalist approach might look like.

In developing this argument I am, to some extent, arguing against myself. In previous work, I have argued that we should adopt an experimentalist approach to the regulation of new technologies when there is uncertainty as to their social effects. This means that we should not prematurely restrict or ban new technologies in the absence of good data. Part of the goal of this article is to interrogate that stance. Could there ever be a reason not to experiment with a new technology? Is the immorality or potential social cost of a new technology ever too high to warrant this? I think child sex robots probably come as close as anything to warranting a stance of extreme precaution. One of the goals of this article is to explain why this might be the case. This inquiry should, consequently, have value for other debates about new technologies.

2. A Framework for Thinking about the Regulation of Child Sex Robots

In order to evaluate the case for restrictive regulation of child sex robots we need to have some clear sense of what we are talking about. First, let’s consider the object of any potential regulation, i.e. the robots themselves. A robot is simply an embodied artificial agent. A ‘sex robot’ is an embodied artificial agent that is designed for (or used for) the purposes of sexual

stimulation and gratification. A ‘child sex robot’ is a sex robot that is designed to look and act like a child.

In addition to serving the ends of sexual stimulation and gratification, Danaher⁹ argues that there are three other conditions that a sex robot (child or adult) will typically meet:¹⁰

(a) It will have ‘humanoid form, i.e., it [will be] intended to represent (and...[be] taken to represent) a human or humanlike being in its appearance.’

(b) It will have ‘human-like movement/behaviour, i.e., it [will be] intended to represent (and...[be] taken to represent) a human or humanlike being in its behaviours and movements.

(c) It will have some degree of artificial intelligence, i.e., it [will be] capable of interpreting and responding to information in its environment, where such response may be minimal (e.g., simple preprogrammed behavioral responses) or more sophisticated (e.g., human-equivalent intelligence).

The focus on ‘human-likeness’ is, as Danaher notes,¹¹ not strictly necessary, but seems appropriate given that the intention behind the manufacture of sex robots, as well as most of the ethical and legal concerns they raise, stems from the desire to create an artificial substitute for human-to-human sex. Furthermore, one of the nice things about this definitional framework is that it allows us to define a spectrum (or more properly a three-dimensional space) along which devices can be classified. In one corner of the dimensional space there are devices with very minimal degrees of human-likeness and in the other corner there are the, still hypothetical, science fictional devices that are practically equivalent (in outward appearances at any rate) to human beings.

This ‘dimensional’ approach to the concept of a child sex robot allows to consider their relationship to child sex dolls. This article is ostensibly about the regulation of child sex robots and yet I opened with two examples of legal attempts to clamp down on child sex

⁹ Danaher, J ‘Should we be thinking about sex robots?’, in Danaher, J. and McArthur, N. (eds) *Robot Sex: Social and Ethical Implications* (Cambridge, MA: MIT Press, 2017); Danaher, J., *Sex Work, Technological Unemployment and the Basic Income Guarantee* (2014) *Journal of Evolution and Technology* 24(1): 113-130; and Danaher ‘Robotic Rape’ (n 7).

¹⁰ Taken, with some modifications, from Danaher ‘Should we be thinking...’ (n 9), p. 4-5

¹¹ Danaher, ‘Should we be thinking..’ (n 9), p 5

dolls. The conflation of the two concepts is no accident and is also something shared in many discussions of this phenomenon. Dan Donovan's CREEPER Act, for example, talks about 'electronic pedophilic robots' in its title, and yet when you turn to the substance of the act it focuses more on 'dolls' than 'robots', seemingly viewing the later as a sub-set of the former. The dimensional approach allows us to look at it in another way. Child sex dolls and child sex robots are related phenomena according to this approach: the difference between them is more a matter of degree than of type. Dolls are human-like in appearance, but not in movement or intelligence. If you were to score them along the three dimensions, you might say that they score a 'zero' on movement and intelligence, but could score highly on the appearance dimension. The line between a doll and a robot will then become blurry rather than bright. It is relatively easy to imagine someone adding features to a doll (e.g. some minimal touch response or recorded vocals) that make it more robotic in nature. As a result, it can make sense to think of dolls as simply very primitive forms of robots. Consequently, for the remainder of this article, I will assume that questions pertaining to the regulation of child sex robots and child sex dolls can be addressed at the same time. Obviously, robots that are extremely human-like may raise greater concerns than inanimate dolls, for reasons to be discussed below, but they are nevertheless connected phenomena and it is appropriate to analyse them at the same time.

So much for the object of regulation. What about its subject? Whose behaviour or conduct would we be trying to regulate? There are three obvious targets: (i) the end users of such devices, i.e. the people who want to use these devices for sexual stimulation and gratification; (ii) the manufacturers of such devices, i.e. those who create the product that caters to the desire for this form of sexual stimulation and gratification; and (iii) any intermediaries in the market between end users and manufacturers, i.e. the people who distribute or retail these devices. Any of these three could be an appropriate target for restrictive, regulatory interference. The usual debates about ensuring the effective use of regulatory resources, and picking a target that will best achieve the regulatory aim, will apply. In the opening examples, it is clear that Dan Donovan's CREEPER Act tries to target users, manufacturers and distributors, presumably in the interests of comprehensiveness. The CPS, on the other hand, are targeting end users, presumably because they are the only available legitimate target since the manufacturers and distributors of these products are located in other jurisdictions. Although it may be more effective to cut off the supply at the

root, this option is not immediately available. Consequently, co-opting a statute from the 1800s strikes them as being the most effective thing to do.

Getting clear about the objects and subjects of regulation is important, but it is all for naught unless we know what it is we are trying to achieve through our regulatory interventions. There are two aspects to this. First we must clarify the regulatory attitude that we think is appropriate when it comes to the subjects and objects of regulation; second we must clarify the regulatory tools that we think are justified by these attitudes. Do we think users/manufacturers of child sex robots are acting immorally? Is the kind of immorality they are engaged in worthy of legal sanction or punishment? Our response to these questions will determine our regulatory attitude. Once again, we think about these attitudes as lying along a spectrum of possibilities.¹² At one end of the spectrum there is the ‘incentive’ attitude, which holds that the use, manufacture and retail of such devices is to be encouraged. One step up from that is the ‘laissez faire’ attitude, which holds that the use, manufacture and retail of child sex dolls is legally permissible and should not be subject to any specific restrictions and limitations (beyond, perhaps, health and safety and product liability regulations). One step up from that you might have an ‘experimental/uncertain’ attitude toward, which holds that the use, manufacture and retail of these devices is something that can be tolerated, perhaps cautiously, since its unclear exactly what attitude we should take towards the phenomenon. Further along the spectrum you have the ‘disincentive’ attitude, which permits use, manufacture and retail, but subjects it to numerous disincentive regulations (e.g. equivalent to the attitude taken towards smoking in many jurisdictions). Finally, you have the ‘restrictive’ attitude, which tries to effectively ban the use, manufacture and retail of these devices. Figure 1 (below) illustrates the spectrum of regulatory attitudes. For the most part, the shift in attitudes along the spectrum tracks closely with the degree of moral disapproval: someone who favours ‘incentive’ regulation must morally approve of the technology whereas someone who favours ‘disincentive’ or ‘restrictive’ regulation must morally disapprove of the technology. The one complication here is the ‘libertarian’ attitude. A libertarian could

¹² The spectrum articulated in the text builds upon the one articulated in Danaher, J., Earp, B. and Sandberg, A., Should we campaign against sex robots? In Danaher, J. and McArthur, N. (eds) *Robot Sex: Social and Ethical Implications* (Cambridge, MA: MIT Press, 2017). The main difference between what is in the text and what was in the Danaher, Earp and Sandberg article is the inclusion of additional regulatory attitudes, specifically the ‘incentive’, ‘experimental’, ‘disincentive’ and the ‘restrictive’ attitudes. In the Danaher, Earp and Sandberg article only three attitudes are identified: libertarianism/laissez faire, regulation, prohibition.

morally disapprove of a technology, whilst as the time favouring the laissez faire regulation of that technology.¹³

<FIGURE 1 – SPECTRUM>

For each of these attitudes, there is an appropriate and corresponding set of regulatory tools. If you are trying to encourage the use of such devices, you might provide tax breaks or other monetary incentives to people who manufacture and distribute them. If you are being cautious about their use, you might insist on a limited rollout and licensing arrangement. If you are trying to discourage or ban the use of such devices, you will impose penalties on those who use, manufacture or retail them. These penalties could include monetary fines, imprisonment or other criminal sanctions (e.g. entry onto a registry of sexual offenders, combined with employment and movement restrictions).

I have set out this framework in some detail because it helps to make sense of the arguments that will be presented in the remainder of this article. My focus will be on the correct regulatory attitude one should take towards child sex robots. In particular, I will consider whether any case can be made for shifting from a restrictive regulatory attitude to an experimental one.

In considering the case for shifting our regulatory attitude, I will be adopting a theoretical position known as ‘legal moralism’. This position sometimes gets a bad rap. This is largely because ‘moralism’ is conflated with populism or conservatism (i.e. the idea that law should ‘track’ some popular, conservative moral beliefs).¹⁴ This is not how I understand or use the term. For me, legal moralism is the view that the immorality of an action, object or state of affairs provides a reason for restrictively regulating it.¹⁵ In other words, if the best available moral theories state that conduct of type X is immoral, then this provides a good reason (though perhaps not an overwhelming or decisive reason) for using legal tools to prevent or

¹³ I am indebted to an anonymous reviewer for raising this point.

¹⁴ This is probably due to the overweening influence of the Hart-Devlin debate about moralism and criminal law in legal education. In that debate, Devlin defended a moralistic position that did conflate moralism with populist conservatism. See Devlin, P., *The Enforcement of Morals* (1959) *Proceedings of the British Academy* 129-151; and Hart, HLA, *Law, Liberty and Morality* (Stanford: Stanford University Press, 1963) for more.

¹⁵ This is consistent with the view of moralism in Duff, A., *Towards a Modest Legal Moralism* (2014) 8 *Criminal Law and Philosophy* 217-235; and Simester, P. and Von Hirsch, A., *Crimes, Harms and Wrongs* (Oxford: Hart Publishing, 2011). It also follows closely the definition offered in Thaysen, J.D. ‘Defining Legal Moralism’ (2015) 16(2) *SATS* 179-201.

limit conduct of type X. Conversely, if the best available moral theories state that X is morally desirable or permissible, then this provides a good reason to use legal tools to encourage or facilitate X. This basic position strikes me as being a relatively uncontroversial, and highly plausible. All it states is that there *should*¹⁶ be some connection between law and ideal/objective morality. Where this undoubtedly gets controversial is in tracing out the precise contours or limits of that connection. According to the popular ‘harm principle’ the connection between law and morality should be quite limited: only immoral conduct that is harmful to others warrants some intervention from the law. Conduct that causes extreme offence or that is harmful to the self should not. I favour a broader understanding of the connection, one that allows for other kinds of immoral conduct to warrant legal intervention. My general reason for favouring it is that I am not thinking solely about criminal law when it comes to justifying regulatory interference. The harm principle *may* have some merit if we are concerned purely with the justification of criminal prohibitions, but if we think beyond the limits of the criminal law, the relationship between law and moral reasons should be, and indeed already is, more broadly conceived. Think for example of the ‘disincentive regulations’ that are placed upon the consumption of cigarettes and alcohol. These regulatory interventions are widely supported (up to a point), but the primary justification for them is that the consumption of both substances is harmful to the person consuming them, not because of the downstream harm they may cause to others.¹⁷ Similarly, prohibitions on desecrating corpses are not justified on the basis that the desecration is harmful to some identifiable other, but on the grounds that it symbolises disrespect and/or is extremely offensive.¹⁸ A broader conception of immorality is at play in both cases.

The effect of this moralistic approach on the present inquiry is obvious: it means that the justification of our regulatory attitude toward child sex robots depends, significantly, on the morality of their use, manufacture and distribution.

¹⁶ Hence it is not a metaphysical thesis claiming that there is some necessary connection between law and morality.

¹⁷ Obviously there are complications here. Some regulatory interventions into smoking and drinking are clearly justified on the basis of harm to others (e.g. indoor smoking bans; drink-driving laws), but other kinds of disincentive regulation (e.g. taxes, tariffs, standardized packaging and so on) are clearly justified on the basis of harm to the individual smoker. They are paternalistic and moralistic interventions.

¹⁸ This example is the centerpiece of Danny Scoccia’s strong defence of pure legal moralism. See Scoccia, D. ‘In defence of pure legal moralism’ (2013) 7(3) *Criminal Law and Philosophy* 513-530.

3. The Case for Restrictive Regulation

So is the use, manufacture and distribution of child sex robots immoral? The emerging consensus among academic commentators is that they are, and that this provides good reason to favour restrictive regulation.¹⁹ However, this emerging consensus is more nuanced than it might first appear. The moral reasons that are taken to justify restrictive regulation are of a particular type, one that is not always taken to be persuasive in legal settings, and they could be disrupted by countervailing considerations. In other words, there is a reasonably strong *prima facie* case for restrictive regulation, but it could be disrupted by a strong case for experimental regulation.

First, let's acknowledge that there are some 'easy cases' in which restrictive regulation could be uncontroversially justified, e.g. when the human-likeness of a child sex robot is such that it is, to all intents and purposes, indistinguishable from an actual human child. The argument in that case would be that since the robot is epistemically indistinguishable from a human, we must treat it as akin to a human (as something that can be harmed by actions performed to it) and so worthy of protection against sexual interference by an adult. In other words, the same moral reasons that justify the restrictive regulation of human child sexual abuse carry over to the robotic case. This argument depends on both the technical and metaphysical possibility of constructing such an epistemically indistinguishable robot. Some people doubt that this will ever be the case, believing that robots can never have the right internal mental constitution to count as moral patients; others think that this is technically possible, but probably a long way from becoming a reality.²⁰ We don't need to become mired in the details of this debate right now. It suffices for us to note that the kinds of child sex robots that could be manufactured with existing technology, and with the technology that is immediately foreseeable, won't rise to this level of sophistication. So this particular case for restrictive regulation is not currently persuasive.²¹

¹⁹ Danaher 'Robotic Rape' (n 7), Strikwerda (n 7); and Maras and Shapiro (n 7)

²⁰ On this point, see: Gunkel, D., *The other question: Can and should robots have rights?* (2018) *Ethics and Information Technology* 20:87–99; and Gunkel, D. *Robot Rights* (Cambridge, MA: MIT Press, 2018)

²¹ So why am I talking about it? Because I think it is important to consider the full spectrum of relevant possibilities and not to limit our discussion to currently available technologies. Technology is something that will change and develop over time and as it does the persuasiveness of the case for or against restrictive regulation may also change. The crucial point here is that if child sex robots reach a point where they are epistemically indistinguishable from human children then any case for experimental regulation is moot.

This does, however, pave the way for another ‘easy’ case for restrictive regulation. As Danaher notes,²² most people think that it is morally justified to punish attempts to commit criminal offences, e.g. attempted murder, attempted burglary and so on. The idea is that a person who has attempted an offence, but failed, is as blameworthy as a person who has succeeded (possibly more so in some cases). This willingness to punish attempts extends, in some jurisdictions, to so-called ‘impossible attempts’, i.e. cases in which someone attempted to commit an offence that was, unbeknownst to them, impossible in the circumstances. For example, there are some famous cases in England and Wales in which people have been convicted for attempting to deal illegal drugs, even though this was strictly impossible because the substances they were carrying were not, in fact, illegal drugs.²³ By similar reasoning, you could argue that someone who sexually gratified themselves with a child sex robot that they took to be an actual child (no matter how implausible this might be to others) should be guilty of attempted child sexual abuse. Or that someone who imported or trafficked a robot that they took to be an actual child should be guilty of attempted child sex trafficking. In this case, it doesn’t matter if the robot was a moral victim, what matters is what the perpetrator perceived it to be. This would seem to be an easy case for restrictive regulation, but the scenario is probably contrived, not something that is likely to arise very often in practice, and might be difficult to prove.

It’s the harder cases that are of more academic interest and practical concern. Suppose that everyone knows that the robots are, in fact, robots and that they are not capable of being moral victims. Is restrictive regulation warranted in such a case? There are two kinds of arguments one can make. The first is the one used by the CPS in justifying their prosecutions for the importation of child sex dolls, and echoed by Dan Donovan in the text of the CREEPER Act: there is some connection between the use of such devices and other kinds of child sexual offences and so restrictive regulation is warranted because of that connection (usually on the grounds that this general class of offences causes harm to children). This connection could be causal — i.e. those who use child sex robots could, as a result, be more likely to commit offences against real children — or it could be more indirect/inverse — i.e. those who use child sex robots tend to be those who also commit other kinds of child sex offence. In the former case, restrictive regulation would be warranted on preventative grounds: it stops people from committing other offences that harm children; in the latter case,

²² Danaher ‘Robotic Rape...’ (n 7)

²³ For example, *R v Shivpuri* [1986] UKHL 2

the idea would be that restrictively regulating the possession of a child sex robot is warranted by the fact that it is an important and relevant indicator that the person is engaging in other child sex related offences.²⁴ Furthermore, since it might be difficult to prove these other offences, but relatively easy to prove possession of the robot, it's useful to have a set of laws in place that prevent or criminalise possession. Hence, restrictive regulation of possession of the robot is a useful and protective way of using the law. In their justification of their prosecutions, the CPS seem to hedge a little bit on the question of the kind of connection between the use of a child sex doll and other child sex offences, possibly favouring the indirect, indicative connection. As they say in one of their press releases:²⁵

“As in other cases elsewhere in the country, this was not just about the importation of a child sex doll. When the defendant's address was searched, large numbers of indecent images of children were also found. This is why cases involving child sex dolls are particularly alarming, as they can be an indicator of other offending.”

They are probably right to hedge since the existence of a clear, linear causal relationship between the use of a child sex robot and other types of child sex offence is as yet unproven, and could be disrupted by future empirical findings.²⁶ Nevertheless, a precautionary approach might provide some justification for restrictive regulation in the absence of clear empirical evidence. This is something I will return to in a moment. Empirical uncertainty of this kind is, however, one of the things that has encouraged academic commentators to present a different argument for restrictive regulation in the harder cases. This argument does not focus

²⁴ To be clear, I am not endorsing this argument. I am merely suggesting that it represents some of the reasoning that is out there, particularly from the CPS press releases. Nevertheless, an anonymous reviewer raises an important objection. This argument sounds like an attempt to justify the restrictive regulation of one thing because it correlates with something else that we believe warrants restriction. But surely mere correlation is not enough to warrant this? Suppose all the people who commit burglaries have tattoos. Suppose we also know that 50% of the population have tattoos and but that it is only 5% of those with tattoos that commit the burglaries. Should we therefore criminalise getting a tattoo? Surely not. I agree that this is implausible, but I suspect that is only because we have independent reasons for thinking that there is no causal link between having a tattoo and committing a burglary, and because the two things seem to have no other obvious affinity or similarity. The case of possessing child sex dolls and other kinds of child sexual offending is different: there we have some suspicion of a causal connection (even if it is disputed/uncertain) and some obvious affinity between the two activities, namely a sexual interest in children or child-like entities. So even if we did endorse the ‘indication argument’ in the case of child sex dolls, I do not think we would thereby be forced to endorse an implausibly broad principle for restrictive regulation. (Note: the tattoo example is taken, with some modifications, from the anonymous reviewer)

²⁵ For the full text, see <https://www.cps.gov.uk/south-east/news/prison-sentence-child-sex-doll-importation> (visited on 23/7/2018)

²⁶ Cox-George, C. and Bewley, S. I, Sex Robot: the Health Implications of the Sex Robot Industry (2018) 44 BMJ Sexual and Reproductive Health 153-154.

on the consequences or indirect effects of child sex robots. Instead, it focuses on what their use (and, by proxy, manufacture and distribution) says about the character or attitudes of the user and the society that tolerates their existence.

There are several different ways of running this argument. Perhaps the most intuitive is to view the argument as being about the *offensiveness* of child sex robots. As noted in the introduction, there is something deeply repugnant (unthinkable) about the idea. We can tolerate a lot in a modern pluralist society, but tolerating the manufacture, distribution and use of child sex robots will probably strike many as a step too far. Even if no actual children are harmed in the process, the idea that people will be sexually gratifying themselves with such objects will be unsettling to many. Of course, offensiveness arguments of this sort don't have a great track record. Many people were (and some still are) offended by the idea of homosexuality, racial equality and gender equality. This doesn't mean that the law should respect their feelings. Offensiveness is not always a good guide to regulatory practice. That said, as many have noted, it is difficult to completely eliminate offensiveness as a ground for restrictive regulation.²⁷ For example, and as mentioned above, desecrating the corpses of dead people is legally prohibited in most countries, and this prohibition does seem to be linked to its offensiveness.²⁸ Similarly, there are certain forms of pornography that are banned in due to their extreme or gross offensiveness. In England and Wales this includes the general offence of possessing extreme pornography,²⁹ as well as more specific bans on the possession of grossly offensive virtual and cartoon like images of children.³⁰ Similar offences exist in other countries too.³¹ So even if you disagree with offensiveness as a general rationale for restrictive regulation, you might agree that some limited category of extremely offensive materials/conduct does warrant restrictive regulation. It seems plausible that if there is such a category of extremely offensive materials, child sex robots would fall within it.

²⁷ Feinberg, J. *Offense to Others* (Oxford: OUP, 1985); and Scoccia, D., In defense of "pure" legal moralism (2013) 7 *Criminal Law and Philosophy* 513–530

²⁸ Scoccia (n 18)

²⁹ Johnson, P. 'Law Morality and Disgust: The Regulation of 'Extreme Pornography' in England and Wales' (2010) 19 *Social and Legal Studies* 147

³⁰ Section 7(7)a of the Protection of Children Act 1978 (as amended) bans indecent pseudo-photographs of children, which includes virtual images; and Section 62 of the Coroners and Justice Act 2009 makes it an offence to possess pornographic cartoon-like or fantasy-like images of children. For a discussion of the latter, see Ost, S. 'Criminalising fabricated images of child pornography: a matter of harm or morality?' (2010) 30(2) *Legal Studies* 230-256.

³¹ For example the Australian Crimes Legislation Amendment (Telecommunications Offences and Other Measures) Act (No. 2) 2004 bans virtual pornographic representations of children.

I would not, however, be in favour this argument because I am not sure that offensiveness provides a moral reason for prohibition or restriction. But I think there are other, more persuasive, ways to run the argument that do provide such reasons. In particular, I think an argument that focuses on the harm to individual moral character/virtue that is implicit in the use (and by proxy manufacture and distribution) of child sex robots is far more credible. Litska Strikwerda has developed one such argument.³² Using Alisdair McIntyre’s theory of virtue — which holds that virtue is associated with the intrinsic goods of certain activities — she argues that the use of child sex robots is non-virtuous because it normalises and encourages a particularly problematic, ‘incomplete’ form of sex. Strikwerda argues that there is an ideal form of sex — a mutual, embodied, consensual, sexual interaction — and then there are lesser forms of sex. The ideal form exemplifies the intrinsic goods of sex and is, consequently, the one that exemplifies virtue. The lesser forms are incomplete because they lack one or more of the features of complete sex (e.g. they are not mutual or embodied or consensual). They consequently do not exemplify virtue and may even corrode it. Some of these lesser forms are tolerable. Masturbation, for example, is tolerable because it does not encourage or endorse a form of sex that flouts norms of consent or mutuality. Other forms are not tolerable because they normalise, or express endorsement for, forms of sex that disregard mutuality and consent. The consumption of wholly computer-generated child pornography is, for Strikwerda, an example of this: even if no children are harmed by its production, the person using it for sexual gratification is endorsing a form of sexual interaction that flouts norms of mutuality and consent. She then argues that the use child sex robots is non-virtuous and intolerable for broadly similar reasons:³³

“...[since] sexual acts between adults and children are always considered imposed, sex acts performed on child sex robots can be seen as replicas of imposed sex acts. This leads to the conclusion that, despite the differences, sex acts performed on child sex robots are, just like the sex act of watching entirely computer-generated child pornography, likely to erode respect for persons, for they are incomplete and flout our sexual mentality, which is based on equality.”

³² Strikwerda (n 7)

³³ Strikwerda (n 7), p 146

In a similar vein, Danaher³⁴ argues that anyone who uses child sex robots would, through such usage, display a defective or highly problematic moral character. Danaher focuses more on robots designed to replicate rape fantasies in his argument, but he also considers its application to child sex robots as well. Adopting an argument originally developed by the aestheticist Stephanie Patridge,³⁵ Danaher argues that having sex with a robot that is designed to look and act like a child has an ‘incorrigible social meaning’ that indicates a lack of moral virtue. To be more precise, he argues that the act symbolises a total disregard for a sexual morality based on equality and respect for persons. Furthermore, using some preliminary findings from human-robot interactions, he argues that a person who was sexually gratified by such an act would either have to (a) overcome or (b) lack an innate moral resistance to performing such an immoral act with a robot.³⁶ Either way, it would be telling us something important about their moral character. This is something that Strikwerda also emphasises her work, pointing out that since people naturally tend to anthropomorphise robots, a person who has sex with a child sex robot is also likely to be anthropomorphising it, which suggests that their attitude toward the robot is likely to be similar to their attitude toward real children.³⁷

There are some objections to this line of reasoning. Making inferences to moral character based on actions performed against non-real/virtual entities has been the subject of much debate in the field of video game ethics.³⁸ Some people argue that standard moral judgments do not apply in these cases. For example, Ostritsch argues that many actions in video game environments can be deemed ‘amoral’ as opposed to immoral/moral. They are just a bit of fun. In the game world, you are stepping outside the ordinary rules of moral life, if only temporarily, and so are not open to the same kinds of moral censure and approbation.³⁹ Could someone make a similar argument about the use of child sex robots? Could it be that

³⁴ Danaher ‘Robotic Rape’ (n 7) – Strikwerda’s article explicitly builds upon the argument set out in this article.

³⁵ Patridge, S. ‘The Incorrigible Social Meaning of Video Game Imagery’ (2010) 13(4) *Ethics and Information Technology* 13(4): 303–312; and Patridge, S., ‘Pornography, ethics and video games’ (2013) 15(1) *Ethics and Information Technology* 25–34

³⁶ Danaher ‘Robotic Rape’ (n 7), pp 87-88

³⁷ Strikwerda (n 7), p 144

³⁸ Tillson, J. Is it distinctively wrong to simulate wrongdoing? (2018) *Ethics and Information Technology* DOI: 10.1007/s10676-018-9463-7; Luck, M., The gamer’s dilemma: an analysis of the arguments for the moral distinction between virtual murder and virtual paedophilia (2009) 11(1) *Ethics and Information Technology* 31–36; Luck, M., Has Ali dissolved the Gamer’s Dilemma? (2018) *Ethics and Information Technology* DOI: 10.1007/s10676-018-9455-7; Ostritsch, S., The amoralist challenge to gaming and the gamer’s moral obligation (2017) 19(2) *Ethics and Information Technology* 117-128.

³⁹ See Ostritsch (n 38), p 118. Analogous arguments can be made about acting in movies or plays. Someone who plays an evil character in a movie (e.g. Daniel Day Lewis in *There Will be Blood*) does not deserve the moral censure that would attach to that character if he were real. That said, there are some complexities

someone who uses them is stepping outside the ordinary rules of moral life and so should not be subject to any moral evaluation?

No. There are two reasons for this. First, the use of child sex robots is not a ‘game’ and so does not exist in some realm of suspended moral animation. The person who gratifies themselves with a child sex robot is not doing so in a fantasy world. Second, as Ostritsch acknowledges, there are cases in which the amoralist approach to video games breaks down. One such case is where the actions within the video game spill over into the real world; another such case is where the game as a whole seems to endorse a morally problematic worldview. For example, a computer game like RapeLay (in which the player is rewarded for raping women) is not just a bit of fun; it is endorsing a worldview in which rape is acceptable and socially rewarded. A person who plays such a game — and gets a kick out of it — must, to some extent, share in that morally problematic worldview. The same reasoning would apply to someone who sexually gratifies themselves with a child sex robot. It would also apply to anyone involved in encouraging or facilitating this usage, although probably to a lesser extent. In other words, they would also display a lack of moral virtue, but it may not be as serious as that displayed by the user.

Another objection to this line of reasoning is that it does not apply to certain cohorts of people, specifically paedophiles. I discuss the specifics of paedophilia in more detail in the next section, but to preempt one of the details from this discussion: the general consensus among those who study paedophiles is that paedophilia is not something that is actively chosen. In other words, people with paedophilia do not (necessarily)⁴⁰ want to be sexually attracted to children; they just find themselves having a strong sexual preference for children. Oftentimes, they are fully aware of the harm it would cause if they acted on this preference. Given the unchosen nature of their sexual orientation, you could argue that a paedophile who seeks out sexual gratification with a child sex robot actually evinces good moral character. Moen and Sterri explicitly make this point in a recent paper, suggesting that paedophiles ‘show respect, care, and concern when they choose to abstain from seeking sexual contact with real children [and] instead [use] non-sentient surrogates’.⁴¹

⁴⁰ Some paedophiles may end up endorsing their sexual preference.

⁴¹ Moen, O.M. and Sterri, A.B. ‘Pedophilia and Computer-Generated Child Pornography.’ In Boonin, D. (ed.), *The Palgrave Handbook of Philosophy and Public Policy* (Cham, Switzerland: Palgrave MacMillan, 2018), at p 375.

This argument has some plausibility, but I would say three things in response. First, even if this argument is correct, it only applies to a small subset of people. The actual prevalence of paedophilia is unknown, but is estimated to be around 5% (and certainly less than 10%) of the general population.⁴² Paedophiles also only represent a subset of the total number of people who engage in child sexual offending, or who have fantasies of sexual contact with children. This is important because the argument about bad moral character would still apply to the majority people who might be inclined to make use of child sex robots: in those cases we would still be entitled to make inferences to bad moral character. Second, whether or not paedophiles actually evince good moral character by making use of ‘non-sentient sexual surrogates’ is, I would submit, largely dependent on the consequences of their doing so.⁴³ If the use of such devices emboldens or encourages them to seek out real children, then it won’t evince good moral character even if no child is directly harmed by the use of the robot; if it clearly discourages contact with real children, then maybe it would indicate some improvement in moral character. I consider this empirical/consequential question in more detail in the next section. Finally, assuming we don’t know whether the use of robots encourages or discourages contact with real children, the mere fact that having sex with a child sex robot is *less bad* than sexually abusing an actual child does not imply that it is good or that it evinces good moral character. One type of conduct can be better than another, without either being particularly good or indicative of good character. In fact the opposite could be true. For example, beating someone up is better than killing them, but both are indicative of bad moral character.⁴⁴

When you combine the problem of bad moral character with empirical uncertainty about the consequences of using child sex robots, I think you get a reasonably good case for restrictive regulation. The fact that the use, manufacture and distribution of such robots evidences a lack of moral virtue (indeed, more than that, actual moral viciousness) provides a moral reason to oppose these practices; and the empirical uncertainty would provide an additional reason to be cautious about them. The restrictive regulation endorsed by both considerations may not be a punitive criminalisation, of course. It is important to emphasise

⁴² For these figures, see Seto, M. Pedophilia (2009) 5 Annual Review of Clinical Psychology 391-407

⁴³ This is a proposition that is argued for at greater length in Danaher, J. ‘The Symbolic-Consequences Argument in the Sex Robot Debate’ in Danaher and McArthur (eds) *Robot Sex: Social and Ethical Implications* (Cambridge, MA: MIT Press, 2017).

⁴⁴ This point has been made in relation to eating meat and eating meat-like surrogates, with some philosophers arguing that eating meat-like surrogates is still bad, even if it is better than eating actual meat. It has also been made in relation to wearing fake fur and fake animal products. See Fischer, B. and Ozturk, B. ‘Facsimiles of Flesh’ (2017) 34(4) *Journal of Applied Philosophy* 489-497 for both of these arguments.

this point. We do not have to send people to jail to restrictively regulate child sex robots. A general ban combined with heavy fines, or inclusion of users on a sex offenders register, or prohibition from running a company (for those involved in manufacture) and so on, could suffice. These are details that will need to be worked out by examining the cost effectiveness of each proposal.

But before we consider those details, we should at least consider whether any case can be made for less restrictive regulation. Bear in mind that the case for restrictive regulation is based on a combination of moral approbation and moral precaution. This could be rebutted if it could be shown that the use, manufacture and distribution of child sex robots had beneficial consequences. This is something that is generally true: at a certain threshold, beneficial consequences tend to trump other moral reasons for action. A resolute Kantian might insist that lying is always wrong and should always be avoided, but most people accept that if it really can be proved to serve a greater good, lying is tolerable. This doesn't mean that the wrongness of lying goes away, but it is overridden by other considerations. The question is whether there is any greater good that could be served by more permissive and experimental regulation of child sex robots. At a robotics conference in 2014, Ronald Arkin made one suggestion in this regard. He suggested that 'robots designed to look and act like children could be used to treat pedophiles the way methadone is used to treat drug addicts'.⁴⁵ Does that do the trick?

4. The Problem with Experimental Regulation

Let me first make a concession. I think that the Arkin hypothesis, if proven, would indeed provide a good reason to favour a more permissive and experimental approach to the regulation of child sex robots. It might make us uncomfortable, and the idea of people making and using such devices might be deeply offensive, but if it could be shown to reduce the harm to actual children, and also improve the lives of people with paedophilic predilections,

⁴⁵ Strikwerda (n 7), p 133

we might have to get over that discomfort and offence. To be clear, this doesn't mean that we should encourage a free for all in the market for child sex robots. At best, I suspect the Arkin hypothesis would give us reason to allow licenced facilities and select therapists to use such devices as part of a larger set of therapeutic interventions. But that's still something of a shift from the restrictive position endorsed in the previous section.

The problem, of course, is that the Arkin hypothesis is unproven. In a recent survey of the empirical literature, it was found that there is simply no evidence one way or the other regarding the beneficial or detrimental consequences of sex robots.⁴⁶ No evidence is no evidence. It's all just theoretical conjecture and supposition at this point in time. This means that when we consider the Arkin hypothesis we cannot think about it in terms of making the case for more permissive regulation in and of itself. Instead, we have to think about it in terms of making the case for some kind experimental regulation, that would open the door to the use and distribution of such devices as part of a test to see whether it is true or false. In other words, the question we need to ask is: should we allow some scientists, clinicians, and therapists to use these devices in order to provide us with the evidence we are currently lacking?

You might think 'of course we should'. Surely it would be worth knowing the answer? If it turns out to be true, we'd be foolish to have passed up the opportunity of improving the welfare of children and people with paedophilic predilections. If it turns out to be false, well nothing ventured, nothing gained. We can return to restrictive regulation if needs be.

I think this is a mistake. Opening the door to experimental regulation of this sort has hidden costs, and we should be extremely cautious about doing so in the case of child sex robots. If we allow scientists and clinicians to test the Arkin hypothesis, it will take time, money and effort to do so. Perhaps it will be 'private' time and money, or perhaps (more likely) it will be 'public' time and money. The exact source doesn't matter so much; what matters is that it will involve some expenditure of time and money. There are, at the same time, lots of other good causes that we could invest the same time and money in. To use the economic jargon, there is a significant *opportunity cost* involved in testing the Arkin hypothesis. If we are going to pay that opportunity cost, then there should be some reasonable

⁴⁶ Cox-George and Bewley (n 26)

prospect of success with regards to the experiment. To be more precise: (i) there should be some reasonable prospect of there being a beneficial effect (i.e. some initially plausible mechanism) and (ii) there should be some reasonable prospect of being able to tell whether there is a beneficial effect or not (i.e. some epistemic access to the truth). If we have neither, then the experiment may not be worth running.

An analogy might help. Suppose someone came to you claiming that they have just designed a perpetual motion machine and are looking for some money so that they can build and test the device. They have some conditions attached to this request. They won't let you inspect the designs to determine their plausibility and you have to trust them to run the experiment themselves and to be honest about their results. Would you invest the money? Clearly you wouldn't. There is no reasonable prospect of success. A perpetual motion machine is, according to the best available scientific knowledge, impossible and even if you wanted to give it a shot, running the experiment in the manner envisaged wouldn't give you the knowledge you require. There is no plausible mechanism and no epistemic access. You need both to have a plausible *prima facie* case for investing your time and money in the project.

Now, I don't believe the situation is quite so bad when it comes to the Arkin hypothesis. The perpetual motion machine is an extreme case. Nevertheless, I believe there are at least four reasons for doubting whether testing the Arkin hypothesis has a reasonable prospect of success and thus whether a *prima facie* case can be made for experimental regulation. The first two reasons question whether there is a plausible mechanism underlying the hypothesis; the third and fourth, which I believe are more decisive, cast doubt on whether we will have epistemic access to the information we need to determine whether the test was a success or not.

The first reason for doubt is that the Arkin hypothesis contradicts existing best practice with regards to the treatment of those with paedophilic predilections. There are many unknowns when it comes our understanding of paedophilia and how best to treat it. This epistemic uncertainty is something I return to in a moment. Nevertheless, the general consensus at present is that paedophilia (which is defined as the sexual interest in

prepubescent children) is not something that can be cured.⁴⁷ Those with paedophilic predilections tend to keep them over the course of their adult lives. This means that paedophilic desires can, at best, be managed and controlled. The current therapeutic preference is to do this by encouraging people not to act on their sexual desires, whether that action takes the form of advances towards actual children or the consumption of child pornography. Instead, clinicians and therapists try to control desires through cognitive behavioural therapy, or suppress them through the use of pharmacological interventions such as steroidal antiandrogens or other forms of ‘chemical castration’.⁴⁸ Consider, for example, the German Dunkelfeld (“Dark Field”) Project which has been running since 2005.⁴⁹ This is one of the world’s largest preventative treatment programmes for people with paedophilic predilections. Their therapeutic interventions focus on managing desires through CBT and drugs, not finding an alternative outlet for them. The project’s organisers have found that this can be an effective treatment strategy, compared to (non-randomised) control groups (Beier 2018).⁵⁰ To use robots as an outlet for paedophilic desires would run contrary to this best practice.⁵¹

The second reason for doubt is that it contradicts what we know about how the brain learns and habituates itself to certain kinds of behaviour. A brief aside is in order here. According to one of the founding theories of psychoanalysis, which has deeper roots and is a

⁴⁷ On this point, see: Seto (n 42); Tenbergen, G., Wittfoth, M., Frieling, H. et al The Neurobiology and Psychology of Pedophilia: Recent advances and challenges (2015) *Frontiers in Neuroscience* available at <https://doi.org/10.3389/fnhum.2015.00344>; Hall, R.C., A Profile of Pedophilia: Definition, Characteristics of Offenders, Recidivism, Treatment Outcomes, and Forensic Issues (2007) 82(4) *Mayo Clinic Proceedings* 457–71; and Harvard Medical School, Pessimism about Pedophilia. *Harvard Mental Health Letter* July 2010, available at https://www.health.harvard.edu/newsletter_article/pessimism-about-pedophilia. This is true only with currently available technology. The prospect of advanced forms of biotechnology that enable us to rewire or reconstitute someone’s biology might change this but this is not on the cards right now and, even if it were, it would completely change the conversation around child sex robots. It’s only in the absence of such technologies that the Arkin hypothesis is worth considering.

⁴⁸ Amelung, T., Kuhle, L.F., Konrad, A. et al (2012). Androgen deprivation therapy of self-identifying, help-seeking pedophiles in the Dunkelfeld (2012) 35(3) *International Journal of Law and Psychiatry* 176-184; Turner, D. and Briken, P., Treatment of Paraphilic Disorders in Sexual Offenders or Men With a Risk of Sexual Offending With Luteinizing Hormone-Releasing Hormone Agonists: An Updated Systematic Review (2018) *Journal of Sexual Medicine* 15(1): 77-93.

⁴⁹ Beier, K.M., Grundmann, D., Kuhle, L.F. et al, The German Dunkelfeld project: a pilot study to prevent child sexual abuse and the use of child abusive images (2015) *Journal of Sexual Medicine* 12(2):529-42; Beier, K.M., Preventing Child Sexual Abuse—The Prevention Project Dunkelfeld (2018) *Journal of Sexual Medicine* available at <https://doi.org/10.1016/j.jsxm.2018.03.008>

⁵⁰ Beier (n 49)

⁵¹ An anonymous reviewer wonders whether this current best practice is influenced by the fact that consuming child pornography (which involves the depiction of real children) makes the user complicit in harm to actual children. There is no obvious indication that this is what is influencing best practice from the literature, though it is plausible. Even still, we are only likely to drop this current best practice if we have some good reason to think that the expression of the sexual desires (in some non-harmful context) is beneficial. But, of course, whether we should experimentally test that hypothesis to acquire such evidence is the issue at hand.

widespread popular belief, humans can get rid of unwanted or negative emotions by releasing/venting them.⁵² This is sometimes known as the ‘cathartic’ theory of emotions. The idea is that unwanted emotions build up over time and if they are not released in a safe or harmless context, they are going to explode out in a harmful one. The emotion of anger is often understood in these terms. How many times have you or people that you know been recommended to ‘get your aggression out’ by playing an aggressive sport or game? Presumably, something akin to the cathartic theory lies behind the Arkin hypothesis: if someone gets their unwanted sexual desires out in a safe/harmless environment, they are less likely to act on them in harmful contexts. The problem is that experimental tests of the cathartic theory generally fail. The vast majority of studies of anger suggest that people who are encouraged to vent their anger become more, not less, likely to act out in anger in other contexts.⁵³ In one famous example, test subjects who were encouraged to vent anger towards someone by imagining them while hitting a punching bag were more likely to punish the same person when given the opportunity.⁵⁴ This is probably because, as Gentiles points out,⁵⁵ the cathartic theory contradicts what we know about the way in which the brain learns and adapts. Repeated exposures to the same kinds of stimuli or repetitions of behaviour tend to reinforce themselves, not ebb away. As Gentiles puts it, the current theory is that the brain becomes what the brain does. If this is right, there is good reason to think that the Arkin hypothesis is based on a flawed mechanism.

But there is some opposition to this line of reasoning. In their discussion of the ethics of computer-generated child pornography, Moen and Sterri⁵⁶ argue that there is good evidence to suggest that the availability of child pornography has a tendency to reduce contact sexual abuse against actual children. In support of this they highlight an internet survey done by Riegel suggesting that paedophiles themselves claim that it is beneficial in reducing sexual urges,⁵⁷ a study by Howitt suggesting that there is no clear causal link between pornography

⁵² Gentiles, D., ‘Catharsis and Media Violence: A Conceptual Analysis’ (2013) *Societies* 3: 491-510

⁵³ Gentiles (n 52); Bushman BJ, Does Venting Anger Feed or Extinguish the Flame? Catharsis, Rumination, Distraction, Anger, and Aggressive Responding (2002) 28(6) *Personality and Social Psychology Bulletin* 724-731; Bushman BJ, Baumeister RJ, Stack AD, Catharsis, Aggression, and Persuasive Influence: Self-Fulfilling or Self-Defeating Prophecies? (1999) 76(3) *Journal of Personality and Social Psychology* 367-376; and Geen RG, Quanty MB (1977). The catharsis of aggression: An evaluation of a hypothesis (1977) 10 *Advances in Experimental Psychology*, 1-37.

⁵⁴ Bushman ‘Does venting...’ (n 53)

⁵⁵ Gentiles (n 52)

⁵⁶ See Moen and Sterri (n 41), p 375-6

⁵⁷ Riegel, D. ‘Effects on Boy-Attracted Pedosexual Males of Viewing Boy Erotica,’ (2004) 33(4) *Archives of Sexual Behavior* 321–323.

use and sex crime,⁵⁸ data on recidivism rates among child sex offenders suggesting that exposure to child pornography is not by itself a risk factor for contact abuse,⁵⁹ and a range of population level statistics suggesting that the increased availability of child pornography does not correlate with an increase in child sex offending.⁶⁰ Given this evidence, Moen and Sterri state that ‘we have most reason to conclude that access to [surrogate sexual activity] will result in less adult-child sex’.⁶¹ This might be taken to suggest that there is enough evidence to at least warrant an experimental test of the Arkin hypothesis.

There are, however, two problems with Moen and Sterri’s argument. The first is that none of the evidence they cite contradicts the claim I make above about the absence of a plausible underlying psychological mechanism. For reasons I will get into below, there are no direct experimental tests of whether the use of child pornography emboldens or inhibits the desire for contact abuse. All the evidence is indirect, either based on individual surveys or statistical inferences from crime data. This is why, in a recent comprehensive survey of the empirical literature on this topic – a survey that is noteworthy for covering the evidence cited by Moen and Sterri – Henshaw, Ogloff and Clough conclude that “causative positions are yet to be directly examined or established within the existing empirical literature base’.⁶² The second problem is that even if we accept the evidence presented by Moen and Sterri it is, at best, a partial representation of the evidence that is out there. Against the claims that they make, we would also need to consider: (i) the fact that there are some studies suggesting a causal link between use of child pornography and contact abuse;⁶³ (ii) the fact that there is evidence suggesting a significant correlation between possession of child pornography and contact sexual abuse, with approximately 30% of those convicted of possessing child pornography also found to engage in contact abuse,⁶⁴ and with the real figure possibly being much higher

⁵⁸ Howitt, D. ‘Pornography and the paedophile: Is it criminogenic’ (1995) 68(1) *British Journal of Medical Psychology* 17.

⁵⁹ Endrass, J. et al., ‘The Consumption of Internet Child Pornography and Violent and Sex Offending,’ (2009) 43(9) *BMC Psychiatry* 1.

⁶⁰ For example: Diamond, M. Eva Jozifkova, E. and Petr Weiss, P. ‘Pornography and Sex Crimes in the Czech Republic’ (1989) 40(5) *Archives of Sexual Behavior* 1037–43; Milton Diamond, ‘Pornography, public acceptance and sex related crime: A review’ (2009) 32 *International Journal of Law and Psychiatry* 304–314.

⁶¹ Moen and Sterri (n 41) p 376

⁶² Henshaw, M., Ogloff, J. and Clough, J. ‘Looking Beyond the Screen: A Critical Review of the Literature on the Online Child Pornography Offender’ (2017) 29(5) *Sexual Abuse* 416-445, p. 431.

⁶³ Bourke, M., & Hernandez, A. (2009). ‘The “Butner Study” redux: A report of the incidence of hands-on child victimization by child pornography offenders’ (2009) 24 *Journal of Family Violence*, 183-191. doi:10.1007/s10896-008-9219-y.

⁶⁴ Eke, A. W., Seto, M. C., & Williams, J. ‘Examining the criminal history and future offending of child pornography offenders: An extended prospective follow-up study’ (2011) 35 *Law and Human Behavior* 466-478.

due to the under-reporting of contact abuse;⁶⁵ (iii) the fact that there is evidence suggesting that those who view more extreme forms of pornography and masturbate to online pornography are more likely to commit contact abuse⁶⁶ (which is significant since the use of child sex robots could be viewed as a kind of extreme interactive and masturbatory pornography); and (iv) the fact that although there is population level data suggesting that there has been a reduction (or at least no obvious increase) in child sexual abuse since the widespread availability of child pornography, this must be interpreted in light of the fact that there are many potential causal factors at play, that there has been a general reduction in crime reported in many countries over the same period,⁶⁷ and that there is in any event significant under-reporting of child sexual abuse.⁶⁸ From this more inclusive survey of the evidence, I conclude that there is still no plausible mechanism supporting the Arkin hypothesis and significant empirical uncertainty as to the likely effects of making child sex robots available. Whether experimental regulation should be introduced in light of such uncertainty is, of course, exactly what I am trying to determine.

The third reason for doubting the wisdom of running the test is that the current institutional framework of scientific inquiry has significant and well-documented problems,⁶⁹ and although there are various attempts to rectify the problems, the problems are still sufficient to make us doubt whether running a test of the Arkin hypothesis, in the current environment, is likely to be successful. This claim needs to be finessed. I am not claiming that we should be generally sceptical of scientific inquiry; I am claiming that we should be sceptical of the capacity for scientific institutions to adequately test the Arkin hypothesis *given the ways in which scientific institutions currently operate*. I doubt whether any such test would give us the clear policy guidance we are looking for when it comes to determining the appropriate regulatory stance.

⁶⁵ Bourke, M., Fragonelli, L., Detar, P. J., Sullivan, M. A., Meyle, E., & O’Riordan, M. (2015). The use of tactical polygraph with sex offenders (2015) 21(3) *Journal of Sexual Aggression* 354-367; and Seto, M. C., Hanson, R. K., & Babchishin, K. M. ‘Contact sexual offending by men with online sexual offenses’ (2011) 23 *Sexual Abuse: A Journal of Research and Treatment* 124-145.

⁶⁶ Houtepan, J., Sijtsema, J.J. and Bogaerts, S. ‘From Child Pornography Offending to Child Sexual Abuse: A Review of Child Pornography Offender Characteristics and Risks for Cross-Over’ (2014) 19 *Aggression and Violent Behaviour* 466-473

⁶⁷ As documented in Zimring, F. *The Great American Crime Decline* (Oxford: OUP, 2008) Pinker, S. *The Better Angels of Our Nature* (London: Penguin 2011); Gash, T. *Criminal* (London, Penguin 2016).

⁶⁸ Martin, E.K. and Silverstone, P.H. ‘How Much Child Sexual Abuse is "Below the Surface," and Can We Help Adults Identify it Early?’ (2013) *Frontiers in Psychology* 15(4): 58 doi: 10.3389/fpsy.2013.00058

⁶⁹ For excellent overviews of these problems, see Hughes, B. *Psychology in Crisis* (London: Palgrave MacMillan, 2018), particularly chapter 1; and Fidler, F. and Wilcox, J. ‘Reproducibility of Scientific Results’ (2018) in Zalta (ed) *Stanford Encyclopedia of Philosophy*, available at <https://plato.stanford.edu/entries/scientific-reproducibility/> (accessed 13/12/2018)

Presumably, if we are to test the Arkin hypothesis this means allowing scientists and therapists to test child sex robots out on patient populations and figure out whether there is a beneficial effect compared to control groups. In theory this seems fine, but there are a number of widely-reported structural flaws in the current institutional framework of science that should give us some pause. In the past couple of decades it has become clear that many (if not most)⁷⁰ reported findings in both biomedical science and psychology (the two disciplines most likely to be involved in testing the Arkin hypothesis) are either misleading or false. There are several documented reasons for this, including: the absence of a replication culture in both disciplines,⁷¹ which is troubling in light of the fact that when studies are replicated the replications tend to either fail or suggest a much smaller effect size than the original studies;⁷² the prevalence, particularly in psychology, of small-scale, low-powered studies with a high risk of false positives;⁷³ clear evidence for publication bias and the ‘file drawer effect’ (i.e. the tendency not to publish or report negative findings);⁷⁴ the widespread use of questionable statistical practices that artificially inflate the number of positive findings;⁷⁵ and ‘Rampant Methodological Flexibility’ across studies that makes it difficult to cross compare results and that also opens the door to confirmation bias.⁷⁶ Sometimes the problems are attributable to outright fraud, but there is a general consensus that the problems are primarily to do with the incentive structures within academic institutions and academic publishing.⁷⁷ This suggests that there is a danger that people investigating the Arkin hypothesis will fail to report negative findings, will cherry-pick or massage the data from their studies to make them more

⁷⁰ Ioannidis, J. ‘Why most published research findings are false’ (2005) 2(8) *PLoS Medicine* e124

⁷¹ Makel, MC., Plucker, JA and Hegarty, B ‘Replications in Psychology Research: How Often Do They Really Occur?’ (2012) 7(6) *Perspectives on Psychological Science* 537–542

⁷² Open Science Collaboration, ‘Estimating the reproducibility in psychological science’ (2015) 349(6251) *Science* aac4716; Begley, C. Glenn and Lee M. Ellis, 2012, ‘Raise Standards for Preclinical Cancer Research: Drug Development’ (2012) 483(7391) *Nature* 531–533.

⁷³ Hughes (n 69) pp 19-20; Szucs, D. and Ioannidis, J. ‘Empirical Assessment of Published Effect Sizes and Power in the Recent Cognitive Neuroscience and Psychology Literature’ (2017) 15(3) *PLoS Biology*, e2000797. doi:10.1371/journal.pbio.2000797

⁷⁴ Fanelli, D. ‘Do Pressures to Publish Increase Scientists’ Bias? An Empirical Support from US States Data’ (2010) 5(4) *PLoS ONE* e10271. doi:10.1371/journal.pone.0010271; Fanelli, D. ‘Negative Results Are Disappearing from Most Disciplines and Countries’ (2012) 90(3) *Scientometrics* 891–904; and Rosenthal, R., ‘The File Drawer Problem and Tolerance for Null Results’ (1979) 86(3) *Psychological Bulletin*, 638–641.

⁷⁵ Simmons, JP., Nelson, L and Simonsohn, U. ‘False-Positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant’ (2011) 22(11) *Psychological Science*, 22(11): 1359–1366; John, LK., Loewenstein, G and Prelec, D. ‘Measuring the Prevalence of Questionable Research Practices With Incentives for Truth Telling’ (2012) 23(5) *Psychological Science*, 524–532; Fraser, H., Parker, T, Nakagawa, S, Barnett, A. and Fidler, F. ‘Questionable Research Practices in Ecology and Evolution’ (2018) 13(7) *PLoS ONE* e0200303. doi:10.1371/journal.pone.0200303

⁷⁶ Hughes, B. (n 69), p 22

⁷⁷ See Fidler, F. and Wilcox, J. (n 69) and Hughes, B. (n 69) for this explanation.

‘significant’, or will be otherwise explicitly or implicitly disincentivised from appropriately testing the hypothesis or accurately representing the data from their tests because it doesn’t serve the interests of their careers or reputations.⁷⁸ There is also the problem of conflicts of interests and the biasing effect of commercial funding to be considered, both of which have plagued biomedical science in the recent past.⁷⁹ To be clear, things are gradually improving with respect to these structural flaws, with several important initiatives in place to improve the quality and transparency of scientific inquiry,⁸⁰ but we need to be conscious of the fact that these reforms are nascent and we will not be testing the Arkin hypothesis in a world of dispassionate, rationally ideal scientists, working within well-functioning institutional contexts.

In this regard, we should also bear in mind what we have learned from the empirical investigation of analogous hypotheses in the past. As noted above, the empirical literature on the link between child pornography and contact child sexual abuse is quite heterogeneous and its policy implications are unclear. The situation is similar with respect to the study of hardcore pornography and its connection to promiscuity, sexual aggression and misogyny. In a recent review of the literature on hardcore pornography, Danaher notes that:⁸¹

Many studies do indeed find that users of pornography are (slightly, but significantly) more likely to embrace promiscuity, engage in risky sexual behavior, have worse relationships, have disturbing attitudes toward women, and be more likely to engage in acts of sexual aggression. But other studies dispute these claims, suggesting that users of pornography are more likely to have progressive attitudes toward women, that pornography can be associated with positive relationship outcomes, and that correlations between pornography use and sexual aggression fail to disentangle cause and effect (i.e., higher pornography consumption may be an effect of negative attitudes and aggression, not a cause). Claims regarding the addictive effects of pornography are also hotly disputed. And

⁷⁸ Everett, J.A. and Earp, B. ‘A tragedy of the (academic) commons: interpreting the replication crisis in psychology as a social dilemma for early-career researchers’ (2015) *Frontiers in Psychology* <https://doi.org/10.3389/fpsyg.2015.01152> (accessed 13/12/2018)

⁷⁹ For a readable discussion see: Goldacre, B., *Bad Pharma: How Medicine is Broken and How we Can Fix It* (London: Fourth Estate, 2012)

⁸⁰ See Fidler, F. and Wilcox, J. (n 69), section 4 on the efforts now in place.

⁸¹ This text is taken from Danaher, J., *The Symbolic Consequences Argument in the Sex Robot Debate*, in Danaher and McArthur (eds) *Robot Sex: Social and Ethical Implications* (Cambridge, MA: MIT Press, 2017), p. 119. The original text includes references to the empirical literature to support the individual assertions but these have been omitted.

virtually every researcher in this field laments the low quality and biased nature of the available evidence.

Historical analogues of this sort are not encouraging. Given that researchers have been busily studying the effects of child pornography and hardcore pornography on real world sexual behaviour for more than forty years, and there is still a lack of clarity about its actual effects, how confident can we be that testing the Arkin hypothesis will prove to be a success? Not confident at all, I would submit. There is a serious risk that we would end up in a similarly murky situation, with constant back and forth between rival groups of scientists and clinicians, and little clearcut policy guidance emerging from the fog of intellectual war. You may, of course, argue that this argument proves too much. If I am correct in my scepticism about the current practices of scientific inquiry then surely we should abandon all scientific investigations of all new hypotheses? But such blanket scepticism does not follow. Other proposed investigations may have compelling moral reasons in their favour, may pose little risk to the public or to the individuals involved in the investigation, and/or may have some prima facie plausible mechanism associated with them. These factors could warrant testing even in an imperfect institutional context. It is the absence of these factors, combined with the opportunity cost involved, that weigh against the Arkin hypothesis. This becomes even more compelling when we consider the final reason against experimentation.

The fourth and final reason to doubt the wisdom of running the test is that there are specific problems with studying therapeutic interventions for those with paedophilic predilections, which should make us doubt whether we will ever know if the test is successful. In other words, whatever you think about my scepticism about the institutions of science as currently constituted, you should be particularly sceptical when it comes to their capacity to investigate this hypothesis. I mentioned already that there are many empirical uncertainties when it comes to the study of paedophilia. The very name of the German Dunkelfeld Project highlights the problem. Most paedophiles, and most acts of child sexual abuse, exist in the ‘dark field’, i.e. outside the view of the majority of society. It’s usually only if someone is arrested for an offence that they seek (or are required to have) therapy. This creates problems for the empirical study of therapeutic interventions, as is repeatedly noted by scholars who study this phenomenon.⁸² The gold standard in biomedical science is

⁸² For example both Seto (n 42) and Beier et al (n 49) make this point.

the randomised controlled trial. But given the hidden nature of most paedophilia, and given some of the ethical issues involved, it's very difficult to run such a trial on those with paedophilic desires. Researchers have to rely on self-identifying and self-reporting individuals, ones who are willing to admit they have a problem or are discovered to have one, and then have to treat them under non-ideal conditions. This doesn't mean that empirical inquiry is impossible, but it has severe limitations. This is not to mention the fact that there is an important distinction between paedophilia and child sexual abuse. Many people who sexually abuse children do not identify or warrant classification as paedophiles; and child sexual abuse often goes underreported.⁸³ So if the hope is that child sex robots could be studied for their effects on the more general problem of child sexual abuse, that hope needs to be qualified.

Taken in combination, these four reasons make the case against opening the door to an experimental form of regulation. Is there any other reason for pushing it open, even if only slightly?

5. What if we want to experiment anyway?

The argument I presented is not overwhelming or decisive. You could accept everything I have said and still think we should try to test the Arkin hypothesis. After all, it really would be worth knowing if the use of child sex robots could have some beneficial effect. You could simply insist that we run the test with utmost caution. This is not the position that I favour. On balance, I think we would be better off not opening this particular Pandora's box, at least not until there have been significant structural reforms to the institutions of scientific inquiry. But I accept that it is a reasonable and defensible position, i.e. one that a person could, in good conscience, make.

If people are persuaded by this view, then there are a number of measures they should take to ensure that the experiment is successful. In particular, they should try to ensure that the test is not marred by the previously documented structural flaws in the institutions of science. To this end, they should follow the guidelines already being established by the various 'open science' initiatives.⁸⁴ They should avoid industry funding of research (or, at

⁸³ This is a point made in Seto (n 42) and Harvard Medical School (n 47)

⁸⁴ See, for example, the guidelines set out in Fidler, F. and Wilcox, J. (n 69), section 4.

least, funding from industry with some obvious conflict of interest); they should establish a database in which all tests are pre-registered; they should encourage and facilitate the publication of all test results, whether positive or negative; and they should ensure that scientists are properly rewarded for these tests, irrespective of their results. There should also be very careful licensing and tracking of those who are running these tests, and of the devices that they use. All experiments should follow the strictest ethical guidelines,⁸⁵ and, perhaps most crucially, a containment strategy should be put in place before we start the testing, to make sure that we can cancel the experiment if it proves too costly and prevent wider distribution of these devices if they prove harmful.

These are the minimal conditions under which a more experimental mode of regulation could be justified. And even still, these conditions would not be sufficient to surmount the ‘dark field’ problem that would continue to undermine the systematic empirical assessment of the hypothesis.

6. Conclusion

I conclude by returning to where I started. Sometimes you do indeed have to think the unthinkable. There is great value to posing new ideas and subjecting them to critical evaluation. When a potential new technology comes on the scene, it is important that we think about it from all angles and carefully assess the costs and benefits. In this article, I have tried to do this for child sex robots. I have argued that there are good moralistic reasons to favour a restrictive approach to the regulation of such devices. I have acknowledged that these reasons could be overridden if there was some therapeutic benefit to the technology. But, to figure out if there is some therapeutic benefit, we need to run a test (or series of tests). This would mean opening the door to a more experimental mode of regulation. We should only do this if the test has some reasonable prospect of being successful – i.e. it is premised on some plausible mechanism and we would be able to figure out if it was effective or not. I have argued that there are good reasons to doubt whether there is a reasonable prospect of success. The restrictive approach to regulation is, in my opinion, the wisest choice.

⁸⁵ Danaher (n 81) discusses the rules that should be followed in any socio-technical experiment using sex robots.

*Acknowledgements: I would like to thank Sven Nyholm and the students who participated in a ‘Philosophy of Risk’ PhD seminar at TU Eindhoven in November 2018 for their comments and questions on a draft version of this paper. I would also like to thank two anonymous reviewers for their valuable feedback and criticisms of a draft version of this paper.