



Experimental Philosophy of Technology

Steven R. Kraaijeveld¹

Received: 22 November 2020 / Accepted: 3 March 2021 / Published online: 27 March 2021
© The Author(s) 2021

Abstract

Experimental philosophy is a relatively recent discipline that employs experimental methods to investigate the intuitions, concepts, and assumptions behind traditional philosophical arguments, problems, and theories. While experimental philosophy initially served to interrogate the role that intuitions play in philosophy, it has since branched out to bring empirical methods to bear on problems within a variety of traditional areas of philosophy—including metaphysics, philosophy of language, philosophy of mind, and epistemology. To date, no connection has been made between developments in experimental philosophy and philosophy of technology. In this paper, I develop and defend a research program for an experimental philosophy of technology.

Keywords Experimental philosophy of technology · Experimental philosophy · Philosophy of technology · Ethics of technology · Philosophical method

“Once upon a time there was a world ... which was only meant to experiment with this and that, where one could see by trial and error how everything turned out, what could be made of it.”

—Pär Lagerkvist (1954), *The Experimental World*

1 Introduction

Technology is one of the most distinctive and pervasive features of contemporary life. Its place in society is only likely to become more prominent and its effects more comprehensive. Modern technology has changed our existence—our relationships to ourselves, each other, and the world—to such a degree that some have advocated a new ethics for the technological age (e.g., Jonas, 1984). In this paper, I introduce a

✉ Steven R. Kraaijeveld
steven.kraaijeveld@wur.nl

¹ Wageningen University & Research, Wageningen, The Netherlands

novel methodology for research in philosophy of technology that I call experimental philosophy of technology or “techxphi” for short.

The impetus for this approach stems from two observations. First, there has been virtually no attention for, let alone systematic study of, the role of intuitions in philosophy of technology or people’s intuitions about technology more generally. The recent field of experimental philosophy or “x-phi,” from which experimental philosophy of technology derives its name and methodology, has advanced our knowledge about the role of intuitions across a wide range of philosophical disciplines—from metaphysics and philosophy of language, to philosophy of mind and epistemology (Nichols & Knobe, 2007). To date, however, there has been no attempt to extend this approach to topics in philosophy of technology.

Second, the promising role of experimental methods in philosophy of technology has been too little appreciated. Despite a growing recognition that empirical methods and findings can play an important role in addressing philosophical questions (e.g., Appiah, 2008; Hämäläinen, 2016; Plakias, 2015), and notwithstanding the so-called empirical turn in philosophy of technology (e.g., Achterhuis, 2001; Scharff, 2011), systematic engagement with empirical research has been scarce in philosophy of technology, especially when compared to other disciplines. For instance, while research in experimental ethics has witnessed a proliferation of empirically informed normative work on a wide range of subjects (e.g., on the nature of character, well-being, and the good life), the methodologies and findings from these and related studies have rarely been extended to similar issues in ethics of technology. This is perhaps ironic, because new technologies have afforded researchers an increasingly rich arsenal of experimental research methods. Through advances in virtual and augmented reality, for example, people can be more vividly immersed than ever before in the kinds of moral dilemmas that are characteristic of experimental ethics (Alfano et al., 2018). To experimentally investigate the ways in which people engage with technology itself from a philosophical perspective is only a small, but as yet untaken step away.

My main claim in this paper is that experimental philosophy of technology as a novel methodology can inform and advance current debates within philosophy and ethics of technology, as well as generate new ideas and lines of research in this area. More radically, when techxphi flourishes as a research program, some current assumptions, arguments, and theories in philosophy of technology will turn out to be untenable. My hope is that those who are not persuaded by the more radical implications of experimental philosophy of technology will nevertheless be convinced by its more modest propositions.

I proceed as follows. First, as background, I provide a brief overview of experimental philosophy and the kind of research that has been conducted under its banner. Second, I develop an outline for an experimental philosophy of technology. I address and deflate some potential objections that have been leveled against experimental philosophy generally, which might also be relevant specifically to techxphi. I then distinguish two programs of experimental philosophy of technology—a negative and a positive—and offer examples from classic and more recent research in philosophy and ethics of technology to show how techxphi can make unique and significant contributions. Finally, I provide several general arguments for the distinctive value of experimental philosophy of technology as a methodology for philosophy of

technology, and encourage researchers from all fields with an interest in deep questions about technology to embrace techxphi in a collaborative effort.

2 Experimental Philosophy

Experimental philosophy is an interdisciplinary field of inquiry in which experimental methods are employed to investigate questions that are traditionally within the domain of philosophy (Knobe et al., 2012). Empirical methods and findings are put to use in order to probe and inform philosophical intuitions, assumptions, and concepts. Early experimental philosophy can be distinguished by the explicit goal to empirically investigate philosophical intuitions (see, e.g., Knobe, 2007; Alexander, 2012). As Joshua Alexander and Jonathan Weinberg have put it, experimental philosophy is “unified behind a common methodology and a common aim: the application of methods of experimental psychology to the study of the nature of intuitions” (2007, 56). This inquiry was initially prompted by the perceived widespread use of intuitions in philosophical theorizing.¹ Reliance on intuitions in thought experiments and theory building was considered by early proponents of experimental philosophy to be a pervasive feature of philosophical practice and, importantly, it was at the same time viewed as a practice subject to potentially problematic limitations (e.g., bias and limited generalizability) that could and ought to be tested empirically. To provide an illustration of the kind of work that has been done in this area, especially to familiarize those who may not be acquainted with the field, I will summarize a line of experimental philosophical research on the relation between freedom and moral responsibility.

Traditional philosophical discussions about freedom and moral responsibility have often centered on the assumption that people are natural incompatibilists, meaning that they believe that causal determinism is ultimately incompatible with either free will or moral responsibility (Alexander, 2012). This assumption has been used to place the argumentative burden of proof on philosophical compatibilists rather than incompatibilists. However, one might rightfully ask: Are people, in fact, natural incompatibilists? In an early and innovative set of studies, Nahmias et al. (2004) empirically investigated this assumption. By showing participants a series of vignettes of different scenarios about moral responsibility and free will, they demonstrated that people are actually more likely to be natural *compatibilists* rather than incompatibilists. More concretely, in a study of the so-called Jeremy Caser, Nahmias et al. (2005, 566) presented participants with the following scenario:

Imagine that in the next century we discover all the laws of nature, and we build a supercomputer which can deduce from these laws of nature and from the current state of everything in the world exactly what will be happening in the world at any

¹ The prevalence of intuitions—and, importantly, of reliance on intuitions—in analytic and other areas of philosophy has been a matter of considerable debate. For instance, while Bealer (1996, 1998) has argued that philosophical intuitions belong to philosophy’s “standard operating procedure” (1996, 122), others like Williamson (2004) and Cappelen (2012, 2014) have argued that philosophy does not rely on intuitions as commonly understood. See Deutsch (2015) for an extensive account of the current and complex state of intuitions in (experimental) philosophy. I do not take a stance on this issue here. My scope is limited to the role of intuitions in philosophizing about technology, where, as I will argue, these might play an important yet underappreciated role.

future time. It can look at everything about the way the world is and predict everything about how it will be with 100% accuracy. Suppose that such a supercomputer existed, and it looks at the state of the universe at a certain time on March 25, 2150 AD, 20 years before Jeremy Hall is born. The computer then deduces from this information and the laws of nature that Jeremy will definitely rob Fidelity Bank at 6:00 pm on January 26, 2195. As always, the supercomputer's prediction is correct; Jeremy robs Fidelity Bank at 6:00 pm on January 26, 2195.²

They then asked participants to suspend disbelief and to imagine the scenario to be true exactly as stated. Importantly, they subsequently asked participants whether they thought that, when Jeremy robs the bank, he does so of his own free will. They found that a significant majority of participants (76%) judged Jeremy to have robbed the bank of his own free will. The normative conclusion that follows for Nahmias and his colleagues is that, if the burden of proof should be on anyone, it must be on those who assume people to be natural incompatibilists. After all, this is not what their findings suggest about people's actual beliefs.

Philosophers' intuitions about people's basic beliefs, in this case about free will and moral responsibility, were thus challenged through empirical findings. Contesting intuitions in this way has come to be one of the hallmarks of experimental philosophy. Another trend has been to investigate potential differences or biases in these intuitions. Differences between philosophical intuitions have been studied, and predictable differences have been found, across genders (Buckwalter & Stich, 2014), cultures (Buchtel & Norenzayan, 2008; Haidt & Joseph, 2007), and personality traits (Bartels & Pizarro, 2011).

The main idea behind these findings is that philosophical intuitions are neither stable nor uniform across individuals and groups. There are real differences in the kinds of intuitions that different people have; and philosophical intuitions are at least sometimes sensitive to morally irrelevant features. Moral intuitions were found, for instance, to be subject to order effects (Wiegmann et al., 2012), which casts doubt on the stability and reliability not just of interpersonal but also of intrapersonal intuitions. As an extension of this research, the expertise of professional philosophers, who may be expected to have more robust philosophical intuitions, has also been questioned through experimental means (e.g., Schwitzgebel & Cushman, 2012; see also Weinberg et al., 2010).

These early developments in experimental philosophy, as previously noted, focused on the role of intuitions in philosophy; but the field quickly expanded beyond primarily considering intuitions.³ A recent description of x-phi, for instance, states that its aim is to "systematically collect and analyze empirical data in attempting to answer philosophical questions or solve philosophical problems" (Sytsma & Livengood, 2016, 5). So far, research in experimental philosophy⁴ has covered many

² For the sake of clarity, I omit some of the details of the Jeremy Cases as well as the other studies that are included by Nahmias et al. (2005).

³ Though the field is no longer restricted to intuitions, and non-intuitional work is being conducted, the study of intuitions still plays a dominant role in experiment philosophy (see, e.g., Sytsma & Livengood, 2016).

⁴ It should be noted that there is some debate about what precisely should count as experimental philosophy. For instance, ought it to include only the work of those self-consciously active within the field, or should the work of non-self-identifying researchers also be embraced? Furthermore, if the latter, then what are the criteria that should make it count? See, for an empirical study of philosophers' views on these matters, and for a summary of the debate, Sytsma and Livengood (2016).

different traditional areas of philosophy, including metaphysics (Nichols & Bruno, 2010), philosophy of language (Haukioja, 2015; Machery et al., 2004), philosophy of mind (Knobe & Prinz, 2008; Sytsma, 2014), intentionality (Knobe, 2010), and epistemology (Buckwater, 2010; Feltz & Zarpentine, 2010). Most recently, the case has been made for experimental philosophical bioethics or “bioxphi” (Earp et al., 2020).

I believe that it is time for an experimental philosophy of technology. In the following sections, I provide a generic account of techxphi and defend it against several potential objections. I then outline two different programs for an experimental philosophy of technology, which I develop in connection with some recent lines of research in ethics of technology in order to illustrate what the programs can offer.

3 Experimental Philosophy of Technology

In the most general terms, experiment philosophy of technology applies the methods of experimental philosophy to topics in philosophy of technology. Techxphi involves the combination of empirical methods, like the controlled experiments that are characteristic of psychology, neuroscience, and other social sciences, with philosophical and normative analysis.⁵ A flourishing techxphi would mean a truly interdisciplinary effort between empirically minded philosophers of technology and cognitive and social scientists in order to examine deep questions about technology.⁶

Before distinguishing between the negative and positive programs of experimental philosophy of technology, some objections that have been raised against experimental philosophy in the literature need to be discussed.⁷ This is not the place to fully engage with these criticisms; nevertheless, to the extent that some objections to x-phi might extend to techxphi, I will briefly discuss (and dispel) three of what I take to be the most pressing issues.

First, it has been argued against experimental philosophy that intuitions do not play a role in philosophy (e.g., Cappelen, 2012; Deutsch, 2015; Williamson, 2007). It appears to me, however, to be much less fruitful to argue for or against the use of intuitions in philosophy across the board than to go about it in a more close-grained way. From a pragmatic point of view, it makes better sense to examine specific arguments and theories, and then to examine the presence or absence of intuitions in these arguments and theories—and only then to criticize or defend the role of such intuitions on a case-by-case basis. Given that intuitions in philosophizing

⁵ While I am sympathetic to a wider understanding and utilization of empirical methods in philosophy generally, for instance as it has been eloquently advocated by Hämäläinen (2016), and philosophy of technology more specifically, I limit myself in this paper to a conception of “empirical” and “empirical methodology” as it has been adopted within the field of experimental philosophy, thus retaining continuity. It seems to me that the acceptance and implementation of the methodology of experimental philosophy is already a significant—if not a final—step toward a more empirically informed philosophy of technology.

⁶ I model this statement of intent after Earp et al.’s (2020) call for an experimental bioethics.

⁷ For an overview, see Knobe and Nichols (2017). I draw on their discussion of objections to experimental philosophy in this section.

about technology have yet to be seriously studied, there is a *prima facie* reason to at least make the attempt to ascertain and better understand their function in the field.

Second, it has been argued against experimental philosophy that intuitions should not play a role in philosophy (see, e.g., Knobe & Nichols, 2017). One might grant that philosophers at times rely on intuitions, but nonetheless argue that this is unwarranted. I want to point out here that, if one subscribes to the view that intuitions should not play a role in philosophy, then it naturally follows that one would also try to counteract reliance upon such intuitions. However, in order to do so, one must first identify such intuitions in order to get the argument off the ground. How could one attempt to eliminate reliance on intuitions without first setting out to detect intuitions and the role that they in fact play in philosophical arguments? Yet this project, as should be clear, is not fundamentally opposed to that of experimental philosophy (Knobe & Nichols, 2017)—in fact, it falls squarely within the negative program that I describe in more detail later.

Third, it has been argued against experimental philosophy that it is not properly philosophy (e.g., Sorell, 2018). This is fundamentally a dispute about whether empirical work has a rightful place within the discipline of philosophy. I follow Hartmann et al. (2013) in rejecting the terms of the debate about whether philosophical questions are (or should be) impervious to empirical research, allowing instead that experimental methods can complement—rather than displace—more purely analytical methods. The goal should be to maintain the rigor of philosophical reflection *and* to make good use of empirical data, while at the same time avoiding empirical naivety (Hämäläinen, 2016). As I will show, research in the philosophy of technology stands to gain much from an explicit engagement with empirical research. The question of how to demarcate the field of philosophy, while interesting, is not of concern to the current project, which seeks to generate new knowledge about philosophical issues surrounding technology. Should a purist about philosophical method object that empirical methods have no place in philosophy, then this should not hamper the techphi project. At worst, one might concede that experimental philosophy of technology is a hybrid discipline.

3.1 Two Programs

Experimental philosophy is commonly divided into different programs according to how these are positioned in relation to the traditional role of intuitions in analytic philosophy. The negative and positive programs are taken to be directly concerned with intuitions—the former with undermining them in a “negative” way, and the latter with making progress in philosophy by examining them in a “positive” way (Knobe & Nichols, 2017). A third program is often identified, which is not so much concerned with the role of intuitions in philosophy traditionally, as with the attempt to “make progress on questions that are directly about people’s thoughts and feelings themselves” (Knobe & Nichols, 2017, 5). Sytsma and Livengood (2016) also differentiate between intuitional and non-intuitional programs, according to whether research is or is not, respectively, about intuitions.

With this additional distinction, the positive and negative programs might be called intuitional, and the third non-intuitional.

In the following sections, I will limit myself to specifying two research programs for experimental philosophy of technology: a negative and a positive one. I will go into more detail later, but in the broadest terms, the difference between them is as follows. The negative program uses experimental methods and findings to debunk (or vindicate) intuitions, judgments, and so on in philosophy and ethics of technology. The positive program, on the other hand, uses experimental means more generally to further knowledge and advance debates in philosophy and ethics of technology. The positive program is thus not necessarily tied to the project of debunking or vindication; it is more broadly concerned with making constructive use of experimental data to inform techno-philosophical reflection.

Making distinctions between programs is a metaphilosophical endeavor. For present purposes, I am primarily interested in creating a space for concrete research topics in experimental philosophy of technology. To distinguish between different research programs provides a useful early schema toward this end, even if the research to be conducted in techxphi will likely fall only loosely within different programs and be subject to considerable overlap, as is the case in x-phi more generally (Knobe & Nichols, 2017). I adapt the two programs from x-phi in order to form a cohesive vision of techxphi that also preserves some continuity. The two programs are not, however, directly translated from, nor wholly translatable to, current x-phi programs. Later work might relate the two programs in experimental philosophy of technology back to work in experimental philosophy. As it stands, however, I have chosen to make the respective divisions between negative and positive programs in order to maximize theoretical clarity and, perhaps more importantly, to create as practicable a guide—and as clear a call to action—as possible.

Before moving on, something must be said about intuitions. There is a substantial philosophical literature about the nature, prevalence, and role of intuitions (see, e.g., Pust 2019). I cannot do justice to this rich area here. If the nature and the role of philosophical intuitions can be contested within experimental philosophy, then this can also be done within experimental philosophy of technology. Differently put, the success of an experimental philosophy of technology does not hinge on the adoption of any particular philosophical view of intuitions. For my part, in this paper, where I speak of intuitions, I follow Devitt (2015) in taking a minimally demanding view of intuitions that brings them close to our ordinary ability to recognize intuitions. Intuitions are, in short, what we ordinarily take them to be.

3.2 The Negative Program

The negative program in x-phi centers on using experimental means to demonstrate that certain intuitions in philosophy are unreliable. Within the negative program of techxphi, this aim is extended in the attempt to debunk (or demonstrate as unreliable, unstable, biased, etc.) intuitions, judgments, and assumptions in philosophy of

technology. The “negative” refers to the role of research in this program, which is to critically examine latent intuitions and assumptions that may escape the arguments, concepts, and theories themselves. The operative metaphor here is that of “negative space”—by framing the space around the substantive arguments and theories in philosophy and ethics of technology, which is present but often taken for granted or hidden from analysis, the hope is that a fuller image will emerge. Of course, when one makes explicit the implicit intuitions and assumptions that are at work in theorizing about technology, and when one sets out to experimentally test these intuitions and assumptions, there are two possible outcomes: The intuitions or assumptions in question may be debunked or vindicated.⁸ As such, the negative program of experimental philosophy of technology includes not only debunking but also vindication efforts. Generally stated, the negative program involves experimental investigation of the intuitions and assumptions that feed into techno-philosophical arguments, in order to question or fortify their value.

I will focus on a specific line of research in ethics of technology in order to illustrate what the negative program of experimental philosophy of technology entails and how it can advance current debates. Consider the influential argument by Andreas Matthias that increasingly autonomous machines threaten to create a so-called responsibility gap, which has initiated a line of responses regarding potential responsibility gaps in technology (e.g., Champagne & Tonkens, 2015; Nyholm, 2018a; Tigard, 2020). The original argument by Matthias (2004) may be summarized as follows:

- [1] Traditionally, manufacturers/operators of machines are held morally/legally responsible for its operations.
- [2] Highly autonomous machines create a novel situation where manufacturers/operators are in principle unable to predict the machine’s future behavior.
- [3] One can only be held morally responsible/legally accountable for things one can control.
- [4] Being unable to predict the machine’s future behavior means that manufacturers/operators do not have control over that behavior.

Therefore,

- [5] Manufacturers/operators cannot be morally responsible/liable for its operations.

Therefore,

- [6] A socially undesirable responsibility gap emerges.

In the spirit of the negative program of experimental philosophy of technology, my claim is that there is an intuition at work in this argument—namely, that being

⁸ The empirical findings may also provide inconclusive evidence, of course. However, one hopes that inconclusive findings will lead to additional, recalibrated research, which may then serve in the end to debunk or validate the respective intuition or assumption.

unable to attribute responsibility to someone or something is problematic for individuals/society. Much of the argument seems to turn on this. This becomes clearer when we separate the theoretical from a more practical concern. If the possibility of the actions of highly autonomous machines being impervious to responsibility-attribution were purely theoretical, but never actual—that is, if we somehow knew that we would never actually encounter a situation where we would have difficulties attributing moral responsibility—then much of the problematic nature of a would-be responsibility gap appears to fall away. The possibility of encountering a concrete situation where a responsibility gap *in fact occurs* and *faces us* is what gives the argument much of its force.

Even if one is not entirely convinced by this interpretation, looking at the argument in this way opens up a line of experimental philosophical research that can help us answer at least some pertinent questions about responsibility gaps. One might, for instance, conduct an experiment in which participants are presented with descriptions of moral dilemmas in which machines cause harm and where the complexity of potential responsibility-attribution is systematically varied (i.e., in some cases, a responsible party will be easily identified; in others, with greater difficulty). How do participants deal with this complexity? Do people, in fact, have trouble attributing responsibility to the outcomes produced by highly autonomous machines?⁹ Another way to approach this would be to study the moral judgments of those who are already working with complex, highly autonomous machines. Using an experimental method, engineers, manufacturers, and operators of highly autonomous machines could be presented with moral dilemmas surrounding machines, harm, and responsibility, in order to examine whether they—that is, people actually working closely with technology of this kind—have difficulties attributing responsibility in certain cases. If this research were to show that people (lay and professional) do not, as a matter of fact, encounter difficulties attributing responsibility even where highly complex systems are concerned, then this gives us reason to think that a responsibility gap is not as threatening a social issue as it has often been portrayed. Of course, to the extent that findings would show people to be unable to attribute moral responsibility in such cases, then worries about a responsibility gap turn out to be better grounded.

The role of intuitions is also evident in a recent offshoot of the responsibility gap argument, namely the potential threat of a retribution gap (Danaher, 2016). Here, the argument is the following:

- [1] Human beings are innate retributivists.
- [2] As highly autonomous robots become ubiquitous, they will more frequently cause harm.
- [3] When autonomous robots cause harm, people will seek targets for retributive blame.

⁹ I focus on negative outcomes as this has been the concern of much of the responsibility gap literature. It is an interesting question, however—and one that can and should be taken up by techxphi—whether there are systematic differences in people’s responsibility attributions according to whether the outcomes are negative or positive. Is there a threat of a responsibility gap when highly autonomous robots make good things happen? I leave this as an open question.

[4] It is unlikely that either the robots or their makers will be eligible for retributive blame.

Therefore,

[5] People will seek retribution but fail to find appropriate targets.

Therefore,

[6] Increased robotization will lead to a retribution gap.

I have argued elsewhere (Kraaijeveld, 2020) that this argument essentially involves people's retributive intuitions, and I have applied an evolutionary debunking argument to intuitions in these cases to argue that they are unjustified and thus ought not to be heeded. Although this approach at first glance may appear to fall within the negative program of experimental philosophy of technology, it was by means of an analytic argument that I attempted to undermine the relevant intuitions and, thereby, the retribution gap argument. I did not, in any case, make use of experimental findings. To qualify as experimental philosophy of technology, at least one empirical premise must be combined with at least one normative premise.¹⁰ The argument by Earp et al. (2020) for undermining or vindicating moral judgments in bioethics may be adapted here, in order to offer an approach to debunking/vindicating moral intuitions and judgments in ethics of technology that can be readily and widely applied in the field. The Debunking/Vindication Argument (DVA) for experimental philosophy of technology may be stated as follows:

[1] Moral judgment *M* or moral intuition *I* is mainly influenced by factor/process *F/P*.

[2] *F/P* is an unreliable (reliable) or morally irrelevant (relevant) factor/process.

[3] So, moral judgment *M* or moral intuition *I* is unjustified (vindicated/not defeated).

It must be noted that the normative conclusion [3] is derived from an empirical premise [1] as well as a normative premise [2]. Empirical premise of this kind, based on experimental findings, combined with normative premises of this kind, provided by philosophical reflection, is how the negative program of experimental philosophy of technology can advance knowledge about moral judgment within the ethics of technology.

To return to the retribution gap, there are at least two ways in which one might take a techphi approach to the argument. On one hand, one could conduct an empirical study to examine whether people in fact respond with the kind of retributive intuitions and moral judgments of blame that make the retribution gap potentially problematic. It is worth carefully examining (i.e., empirically and systematically) both the nature and

¹⁰ I do not mean to imply here that all of philosophy of technology is normative or that experimental philosophy of technology can always be formulated in normative terms (neither would be true). I merely want to suggest that, where normative issues are at stake in techno-philosophical arguments, then in order for it to count as *experimental philosophy of technology*, there must be some combination of normative and empirical premises. Thanks to an anonymous reviewer for pressing me to clarify this point.

the scope of these intuitions and moral judgments in cases of harm caused by highly autonomous robots. If it turns out that people are not as prone to retributive intuitions and/or moral judgments of blame in these cases, then this would give us some reason to think that the social, legal, and moral ramifications of a retribution gap may not be as far-reaching as they have previously been considered.

On the other hand, a more sophisticated way of approaching the argument from the perspective of experimental philosophy of technology is to manipulate the relevant retributive intuitions and moral judgments in an experimental study, in order to test whether they are subject to unreliable and/or morally irrelevant factors or processes. Taking for granted that the intuitions and judgments that go into the retribution gap are as described, perhaps they are ultimately not to be heeded because they are empirically demonstrated to be unreliable, unstable, and etc. If it would turn out that retributive intuitions in cases of robot harm without clear candidates for moral blame were influenced by, say, unreliable factor F or morally irrelevant process P , then one could use the DVA as described above to argue that the judgment based on factor F or process P in these cases means that the resulting moral judgment or moral intuition is not justified. If it turns out, for instance, that when a self-driving car crashes, people's moral judgments about blame are significantly influenced by some morally irrelevant feature like the country of origin of the car manufacturer, then this would give us reason to question their validity.

In one sense, the move made in the second approach is similar to the one that I made in applying an evolutionary debunking argument (Kraaijeveld, 2020), in that both try to undermine in some way the status of the relevant intuitions/judgments. The important difference, as should be clear, is that applying the DVA inherently involves an empirical premise. This is what sets it apart as experimental and thus what makes it count as experimental philosophy of technology.

I have been able to cover only a few recent debates in ethics of technology, but the negative program generally, and the DVA in particular, can be applied to a host of other arguments, intuitions, and judgments in philosophy and ethics of technology.

3.3 The Positive Program

The positive program, which in experimental philosophy centers on making progress directly on all sorts of philosophical issues, can for experimental philosophy of technology be viewed largely as a similar effort to experimentally investigate intuitions, assumptions, thoughts, emotions, concepts, and so on that are relevant to topics in philosophy and ethics of technology.¹¹ What are the intuitions that come into

¹¹ I want to reiterate here that I am sympathetic to a wide array of empirical methods in philosophy generally and in philosophy of technology in particular, including those that do not strictly fall within the "experimental" bracket of experimental philosophy (e.g., narrative or hermeneutic approaches). I have chosen this methodology not only because it promises new insights and lines of research for philosophy of technology, but also because it importantly extends the project of experimental philosophy. Nevertheless, I do not think there is any reason why the positive project as I describe it could not accommodate those who seek to gain a richer understanding in philosophy of technology, for instance by investigating the "moral present" (Hämäläinen, 2016). In other words, should empirically minded researchers who are not experimentally inclined (in the sense given to this by experimental philosophy) wish to take on the positive project or identify with it, then they would be welcome indeed.

play when we think about novel and emerging technologies? How are our moral judgments about technology related to the underlying cognitive and psychological processes that give rise to them? Do certain ethical theories about technology rely on the existence of (previously unexamined) empirical matters of fact? The aim of the positive program is to make progress on these and other questions surrounding technology. The nominal “positive” is indicative of the positive spirit of contribution; this program may thus be seen as wider in scope than the negative program, which focuses more narrowly on debunking and vindication attempts.

There are many potentially interesting lines of research that could be taken up under the auspices of the positive program of experimental philosophy of technology. I will specify two major tasks, which may be pursued separated but will ideally be integrated in a meaningful way. First, there is the descriptive task of probing people’s intuitions and judgments about arguments, theories, and dilemmas in philosophy and ethics of technology. This includes folk intuitions and judgments as well as those of experts. There are many areas in ethics and philosophy of technology where intuitions appear, at least on the surface, to play an important role. Whether we are considering the introduction and ramifications of novel technologies or products of technological advances like cultured or in vitro meat (Van der Weele & Tramper, 2014), accident-algorithms for unavoidable collisions of self-driving cars (Nyholm & Smids, 2016),¹² the possibility of robots being good colleagues (Nyholm & Smids, 2019) or being objects for sexual gratification (Danaher et al., 2017) or mutual love (Nyholm & Frank, 2017), or what the role of technology ought to be in a good society (Brey, 2018), intuitions about many of these questions will play an important role. To explicate them in a systematic way, through empirical investigation, will add to our knowledge about these questions and will assist in formulating new ones.

To give a more elaborate example, one area in which experimental philosophy of technology may be especially productive is in discussions surrounding technomoral change. The main claim of the technomoral change approach is that technology co-shapes many if not all aspects of society, including moral norms and values (Swierstra et al., 2009). The notion of technomoral change readily lends itself to the combination of empirical investigation and philosophical analysis that characterizes experimental philosophy of technology. For, in order to know precisely *how* technology co-shapes moral norms and values, one must know at least some of the empirical matters of fact—about the nature, scope, and direction of changes in people’s moral frameworks and about how technology acts as a driving force. Although rare, some good empirical-normative work has already been conducted in this area. Olya Kudina and Peter-Paul Verbeek (2018), for instance, have studied online discussions about the “explorer” version of Google Glass in order to make progress on the ethical variant of the Collingridge dilemma. Additionally, the proposed use of alternative technomoral scenarios to inform public deliberation about New and

¹² When it comes to people’s intuitions about how self-driving cars should respond to particular scenarios and dilemmas, some important work is already being done as part of the Moral Machine project at the MIT Media Lab. For the Moral Machine experiment, see Awad et al. (2018). The online project can be found here: <https://moralmachine.net>. For a discussion, see also the section on “Empirical Ethics” in Nyholm (2018b).

Emerging Science and Technology (NEST) for Technology Assessment (Swierstra et al., 2009) is a promising way of prompting moral intuitions and judgments on a wide range of relevant issues, which appears closely aligned with the goals of a descriptive experimental philosophy of technology program that uses vignettes and moral dilemmas to elicit intuitions. Where experimental philosophy of technology can make its unique contribution, however, is in not taking the relevant intuitions and judgments at face value, but by experimentally manipulating the scenarios so as to learn about the factors that are involved in producing these intuitions and judgments. For instance, are they subject to situational factors, order effects, or framing effects? Which psychological processes underlie different moral judgments in these scenarios? All of the methodological tools of experimental philosophy may be applied here to categorize and scrutinize the relevant intuitions and judgments.

As in experimental ethics more widely, the use of indirect experiments may give us insight into “the nature of some capacity or judgment: for example, whether certain types of moral dilemmas engage particular areas of the brain,” whereas direct experiments can “investigate whether a claim held (or denied) by philosophers is corroborated or falsified,” which “might mean investigating an intuition and whether it is as widely shared as philosophers claim, or it might mean investigating the claim that a certain behavior or trait is widespread or that two factors covary” (Alfano et al., 2018). If the purpose is to show that the relevant intuitions and judgments are unreliable or biased in some way, then this research may be tallied under the negative program. However, there is no reason why assessing the scope and sensitivity of moral judgments about technomoral change, say, or about self-driving cars causing certain kinds of harm, has to involve the attempt to debunk the judgments in question. To the extent that we really just want to know about these judgments, or about whether intuitions are in fact as philosophers of technology have taken them to be, then this may be seen as an important descriptive task in its own right.

Second, there is the *normative* task of using empirical findings to support normative claims about the subject at hand. Some inchoate work has already been done in this area, even if this was not under the name of experimental philosophy of technology.¹³ For example, Kudina (2019) has recently used empirical methods as well as philosophical analysis to investigate the complex interactions between ethics and technology. As previously mentioned, research into technomoral change that uses technomoral scenarios has also been receptive to the value of empirical input, even though one may still press this work to make the role of empirical research more prominent in discussions surrounding the scenarios and the conclusions to which they give rise (cf. Boenink et al., 2010). Research of this kind remains rare. Given that technologies can “rob moral routines of their self-evident invisibility and turn them into topics for discussion, deliberation, modification, and reassertion” (Swierstra & Rip, 2007, 6), experimental philosophy of technology is well-suited to take on the task of elucidating the (changes in) moral routines, norms, and values surrounding technology, in an effort to strengthen normative conclusions about what the role of technologies *should be* given what they *are* or what they *do* within

¹³ No one, as far as I know, has explicitly taken this approach to date.

particular practices. Think, for example, of people's lived experiences in relation to or as affected by technologies. By studying people's actual experiences (e.g., through qualitative research; see Andow, 2016), one can obtain knowledge that transcends the merely observational or anecdotal; and by documenting the experiences (e.g., of injustice or systematic biases), one can move toward conceptualizing and implementing changes for the better.¹⁴

4 Why Techxphi?

While I think that the forgoing discussion of the different programs, respective examples, and suggestions about further and future research already does much to demonstrate the value of an experimental philosophical approach to technology, there are several more general arguments that speak to its potential value.

First, techxphi can provide a unifying methodology for a number of subfields within philosophy of technology. It can do so at the most general level, where two different approaches—humanities and analytic—are often distinguished (Franssen et al., 2018). Given that I have focused much of the discussion (1) on the role of intuitions, and (2) on the role of intuitions within ethics of technology, it may appear that experimental philosophy of technology is best suited to what Carl Mitcham (1994, 39) has called humanities philosophy of technology, which is committed to the “primacy of the humanities over technology” and studies technology from within the social sciences and humanities. As such, work in this area is concerned first and foremost with the moral, social, and political aspects of technology. Within the analytic approach to philosophy of technology, on the other hand, technology itself has primacy and is viewed as a phenomenon of inherent interest, aside from what it might mean to and for human beings. It must be remembered, however, that I have focused on moral intuitions primarily for reasons of space. Experimental philosophy, as a methodology for investigating philosophical problems, has also already contributed to major issues in analytical philosophy, as evidenced by work done in formal semantics (e.g., Liao & Meskin, 2017; Cariani & Rips, 2017; see also Franssen et al., 2018) to offer but one example. There is no reason to think that experimental philosophy of technology cannot likewise further knowledge and sharpen debates in analytic philosophy of technology. In this way, then, experimental philosophy of technology can provide a much-needed methodological bridge between traditional analytic and humanities approaches to philosophy of technology.

Within ethics of technology, techxphi can also provide a unifying methodology across different applied ethics research programs. This is important, because as Philip Brey has pointed out, “very little work is being done to advance the field of technology ethics theoretically or methodologically” (2010, 44). In the same way that the application of a particular moral principle may usefully be extended

¹⁴ I am here still advocating the experimental methods as they have been used in experimental philosophy; but this seems to me a point at which the thin line may meaningfully be crossed from such methods into more descriptive, nonexperimental, qualitative methods.

to diverse and novel technologies, so experimental philosophy of technology can provide a means of gaining knowledge that is not tied to one particular technology (e.g., gene drives or reproductive technologies). As a specific methodology, experimental philosophy of technology constitutes a novel approach to technology ethics. By its espousal of empirical investigation of moral intuitions and its insistence on tying experimental findings to normative philosophical reflection, it offers an approach to ethics of technology alongside others—be it cultural and political approaches, engineering ethics approaches, or approaches that focus on the ethics of specific technologies.

Experimental philosophy of technology also has implications for fields adjacent and related to philosophy of technology as broadly construed, such as Technology Assessment (TA), Science and Technology Studies (STS), and human–robot interaction (HRI), to name but a few of the more obvious candidates. The lack of a systematic attempt to investigate intuitions in these disciplines may be surprising—not least in the field of HRI, where the study of interactions between humans and robots promises to offer a productive ground for testing intuitions as well as a clear role for empirically informed philosophy. Some recent work in this area, however, does already appear to be within the spirit of experimental philosophy of technology as I have outlined it, even though it has not been explicitly conducted or characterized as such.¹⁵ Nijssen et al. (2019), for instance, have experimentally investigated whether anthropomorphic appearance and attributions modulated utilitarian decision-making about robotic agents, and have found that when people attribute affective states to robots, they are less likely to sacrifice them in order to save humans. de Graaf and Malle (2019) have experimentally tested people’s mental state ascriptions to robots, and found that people use the same conceptual toolbox of behavior explanations for human and robot agents, thus indicating inferences of intentionality and mind, even though people applied those tools somewhat differently to explain robot as compared to human behavior. This is the kind of research that, I think, can and should be part of a much more systematic experimental philosophical approach to how human beings perceive and interact with robots and AI.

Additionally, phenomenological and postphenomenological approaches to technology, the latter of which has in any case been marked by an “empirical turn” and aims to “empirically analyze how particular technologies as “the things themselves” *mediate* the relation between humans and their world” (Zwier et al., 2016, 314), can also benefit from engagement with the proposed experimental methodology to further its own ends. In turn, testing some of the concepts and arguments that underlie (post)phenomenology may provide a fertile ground for potential advances in experimental philosophy of technology.

Second, techxphi can offer a means of generating knowledge across different fields without necessarily requiring the specialized knowledge of such fields. The fundamentally interdisciplinary nature of experimental philosophy of technology means that it can provide a dynamic and productive meeting ground for researchers from a wide range of disciplines and backgrounds. One imagines psychologists, neuroscientist, and cognitive scientists with an interest in normative

¹⁵ I would like to thank an anonymous reviewer for bringing this research to my attention.

questions about technology collaborating with empirically minded ethicists and philosophers of technology in order to investigate questions about, say, how and why people arrive at particular moral judgments about new technologies, or about how and why particular technologies influence the way we think about and treat ourselves and others. This point is also relevant to the previously sketched relation between analytic and humanities philosophy of technology. As a methodology, a collaborative techxphi project can probe the arguments of analytically inclined philosophers of technology just as it can test the claims made by those working on technology from within a humanities approach—all without necessarily requiring specialized background in the respective disciplines. As in so many collaborative efforts, of course, dialog is key.

Third, given that intuitions within philosophy of technology and intuitions about technology more generally are relatively unfamiliar territory from the perspective of our evolutionary history as human beings, it is important in and of itself to be conscious of the role that these intuitions might play. To the extent that experimental philosophy of technology uncovers and scrutinizes relevant intuitions, it can provide a body of knowledge relevant both to philosophical work on technology that has already been conducted (e.g., through the negative program as criticism of the roles that intuitions have played), but also—and perhaps especially—with regard to the future (e.g., through the positive program by mapping the functions of intuitions on an ongoing basis). There is a *prima facie* reason to scrutinize our intuitions when it comes to thinking about technology. When one is dealing with evolutionarily unfamiliar territory—as our technology-laden society is today—our intuitions are less likely to be apt to our environment and are thus less likely to be reliable guides for thought and action. In the end, from our current vantage point, one can only speculate about the myriad ways in which technology will develop and influence, let alone shape, human existence. To keep a finger close to the pulse of our intuitions in an ongoing attempt to understand the impact of technology on our thinking—and our thinking on technology—is a contribution that experimental philosophy of technology is well-positioned to make. If experimental philosophy has successfully garnered new insights for a host of philosophical disciplines, as I believe it has, then the burden of proof should be on those who would deny that it has anything to offer philosophy of technology; especially given the unsettled, perpetually changing, and, in many ways, decisive nature of its subject matter. Should this paper result in an increased self-awareness about intuitions among philosophers of technology, then it will already have succeeded in no small measure.

5 Conclusion

The field of experimental philosophy has, through its engagement with experimental methods, become an important means of obtaining knowledge about the intuitions, concepts, and assumptions that lie behind philosophical arguments, problems, and theories across a wide variety of philosophical disciplines. In this paper, I have extended this burgeoning research program to philosophy of technology, providing

both a general outline of how an experimental philosophy of technology might look and a more specific methodology and set of programs that engages with research already being conducted in the field. I have responded to potential objections to an experimental philosophy of technology and I have argued for a number of unique strengths of the approach. Aside from engaging with work that already involves intuitions in techno-philosophical research, a booming experimental philosophy of technology research program can offer a unifying methodology for a diverse set of subfields, a way of generating knowledge across disciplines without necessarily requiring specialized knowledge; and, at the very least, it can make those working in philosophy of technology—and those in society who engage with technology, which is all of us—more mindful of the intuitions about technology that we may, rightly or wrongly, hold.

Acknowledgements I would like to thank the three anonymous reviewers for their very careful reading of the text and for their constructive feedback.

Author Contribution Not applicable.

Data Availability Not applicable.

Code Availability Not applicable.

Declarations

Competing Interests The author declares no competing interests.

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

References

- Achterhuis, H. (Ed.). (2001). *American philosophy of technology: The empirical turn*. R. P. Crease (Transl.). Indiana University Press.
- Alexander, J. (2012). *Experimental philosophy: An introduction*. Polity Press.
- Alexander, J., & Weinberg, J. (2007). Analytic epistemology and experimental philosophy. *Philosophy Compass*, 2(1), 56–80.
- Alfano, M., Loeb, D., & Plakias, A. (2018). Experimental moral philosophy. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. <https://plato-stanford-edu.ezproxy.library.wur.nl/archives/win2018/entries/experimental-moral/>. Accessed 11 Nov 2020.
- Andow, J. (2016). Qualitative tools and experimental philosophy. *Philosophical Psychology*, 29(8), 1128–1141. <https://doi.org/10.1080/09515089.2016.1224826>.
- Appiah, K. A. (2008). *Experiments in ethics*. Harvard University Press.
- Awad, E., Dsouza, S., Kim, R., Schulz, J., Henrich, J., Shariff, A., Bonnefon, J., & Rahwan, I. (2018). The moral machine experiment. *Nature*, 563, 59–64.

- Bartels, D. M., & Pizarro, D. A. (2011). The mismeasure of morals: Antisocial personality traits predict utilitarian responses to moral dilemmas. *Cognition*, *121*(1), 154–161.
- Bealer, G. (1996). On the possibility of philosophical knowledge. *Philosophical Perspectives*, *10*, 1–34.
- Bealer, G. (1998). Intuition and the autonomy of philosophy. In M. DePaul & W. Ramsey (Eds.), *Rethinking intuition*. (pp. 201–239). Rowman & Littlefield.
- Boeninck, M., Swierstra, S., & Stermerding, D. (2010). Anticipating the interaction between technology and morality: A scenario study of experimenting with humans in bionanotechnology. *Studies in Ethics, Law, and Technology*, *4*(2): Article 4. Available at: <http://www.bepress.com/selt/vol4/iss2/art4>. Accessed 11 Nov 2020.
- Brey, P. (2010). Philosophy of technology after the empirical turn. *Techné*, *14*(1), 36–48. <https://doi.org/10.5840/techné20101416>.
- Brey, P. (2018). The strategic role of technology in a good society. *Technology in Society*, *52*, 39–45.
- Buckwater, W. (2010). Knowledge isn't closed on Saturday: A study in ordinary language. *Review of Philosophy and Psychology*, *1*(3), 395–406.
- Buckwater, W., & Stich, S. (2014). Gender and philosophical intuition. In J. Knobe & S. Nichols (Eds.), *Experimental philosophy*. (Vol. 2, pp. 307–346). Oxford University Press.
- Buchtel, E. E., & Norenzayan, A. (2008). Which should you use, intuition or logic? Cultural differences in injunctive norms about reasoning. *Asian Journal of Social Psychology*, *11*(4), 264–273.
- Cappelen, H. (2012). *Philosophy without intuitions*. Oxford University Press.
- Cappelen, H. (2014). X-Phi without intuitions? In A. R. Booth & D. P. Rowbottom (Eds.), *Intuitions*. (pp. 269–286). Oxford University Press.
- Champagne, M., & Tonkens, R. (2015). Bridging the responsibility gap in automated warfare. *Philosophy & Technology*, *28*, 125–137.
- Cariani, F., & Rips, L. J. (2017). Conditionals, context, and the suppression effect. *Cognitive Science*, *41*(3), 540–589. <https://doi.org/10.1111/cogs.12336>.
- Danaher, J. (2016). Robots, law and the retribution gap. *Ethics and Information Technology*, *18*, 299–309.
- Danaher, J., Earp, B. D., & Sandberg, A. (2017). Should we campaign against sex robots? In J. Danaher & N. McArthur (Eds.), *Robot sex: Social and legal implications*. MIT Press.
- de Graaf, M. M. A., & Malle, B. F. (2019). People's explanations of robot behavior subtly reveal mental state inferences. In: *Proceedings of the international conference on human-robot interaction*, HRI'19, Daegu, South Korea.
- Devitt, M. (2015). Relying on intuitions: Where Cappelen and Deutsch go wrong. *Inquiry*, *58*(7–8), 669–699. <https://doi.org/10.1080/0020174X.2015.1084824>.
- Deutsch, M. (2015). *The myth of the intuitive: Experimental philosophy and philosophical method*. The MIT Press.
- Earp, B. D., Demaree-Cottona, J., Dunne, M., Dranseika, V., Everett, J. A. C., Feltz, A., Geller, G., Hanikainen, I. R., Jansen, L. A., Knobe, J., Kolak, J., Latham, S., Lerner, A., May, J., Mercurio, M., Mihailov, E., Rodríguez-Arias, D., Rodríguez Lopez, B., Savulescu, J., ... Tobia, K. (2020). Experimental philosophical bioethics. *AJOB Empirical Bioethics*, *11*(1), 30–33.
- Feltz, A., & Zarpentine, C. (2010). Do you know more when it matters less? *Philosophical Psychology*, *23*(5), 683–706.
- Franssen, M., Lokhorst, G. J., & Van de Poel, I. (2018). Philosophy of technology. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. <https://plato-stanford-edu.ezproxy.library.wur.nl/archives/fall2018/entries/technology/>. Accessed 7 Nov 2020.
- Haidt, J., & Joseph, C. (2007) The moral mind: How five sets of innate intuitions guide the development of many culture-specific virtues, and perhaps even modules. In P. Carruthers, S. Laurence, & S. Stich (Eds.), *The innate mind, Volume 3*. Oxford University Press. 367–392.
- Hämäläinen, N. (2016). *Descriptive ethics: What does moral philosophy know about morality?*. Palgrave Macmillan.
- Hartmann, S., Lisciandra, C., & Machery, E. (2013). Editorial: Formal epistemology meets experimental philosophy. *Synthese*, *190*, 1333–1335. <https://doi.org/10.1007/s11229-013-0269-1>.
- Haukioja, J. (Ed.). (2015). *Advances in experimental philosophy of language*. Bloomsbury Academic.
- Jonas, H. (1984). *The imperative of responsibility: In search of an ethics for the technological age*. H. Jonas & D. Herr (Transl.). The University of Chicago Press.
- Knobe, J. (2007). Experimental philosophy and philosophical significance. *Philosophical Explorations*, *10*(2), 119–121.
- Knobe, J. (2010). Person as scientist, personal as moralist. *Behavioral and Brain Sciences*, *33*(4), 315–329.

- Knobe, J., & Prinz, J. (2008). Intuitions about consciousness: Experimental studies. *Phenomenology and the Cognitive Sciences*, 7(1), 67–83.
- Knobe, J., & Nichols, S. (2017). Experimental philosophy. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. <https://plato-stanford-edu.ezproxy.library.wur.nl/archives/win2017/entries/experimental-philosophy/>. Accessed 7 Nov 2020.
- Knobe, J., Buckwalter, W., Nichols, S., Robbins, P., Sarkissian, H., & Sommers, T. (2012). Experimental philosophy. *Annual Review of Psychology*, 63(1), 81–99. <https://doi.org/10.1146/annurev-psych-120710-100350>.
- Kraaijeveld, S. R. (2020). Debunking (the) retribution (Gap). *Science and Engineering Ethics*, 26(3), 1315–1328. <https://doi.org/10.1007/s11948-019-00148-6>.
- Kudina, O. (2019). *The technological mediation of morality: Value dynamism, and the complex interaction between ethics and technology*. PhD Diss. University of Twente. <https://doi.org/10.3990/1.9789036547444>.
- Kudina, O., & Verbeek, P. (2018). Ethics from within: Google Glass, the Collingridge dilemma, and the mediated value of privacy. *Science, Technology, & Human Values*, 44(2), 291–314. <https://doi.org/10.1177/0162243918793711>.
- Lagerkvist, P. (1954). *The experimental world*. (p. 53). Hill and Wang.
- Liao, S., & Meskin, A. (2017). Aesthetic adjectives: Experimental semantics and context-sensitivity. *Philosophy and Phenomenological Research*, 94(2), 371–398. <https://doi.org/10.1111/phpr.12217>.
- Machery, E., Mallon, R., Nichols, S., & Stich, S. P. (2004). Semantics, cross-cultural style. *Cognition*, 92(3), 1–12.
- Matthias, A. (2004). The responsibility gap: Ascribing responsibility for the actions of learning Automata. *Ethics and Information Technology*, 6, 175–183.
- Mitcham, C. (1994). *Thinking through technology: The path between engineering and philosophy*. Chicago: University of Chicago Press.
- Nahmias, E., Morris, S., Nadelhoffer, T., & Turner, J. (2004). The phenomenology of free will. *Journal of Consciousness Studies*, 11, 162–179.
- Nahmias, E., Morris, S., Nadelhoffer, T., & Turner, J. (2005). Surveying freedom: Folk intuitions about free will and moral responsibility. *Philosophical Psychology*, 18, 561–584.
- Nahmias, E., Morris, S., Nadelhoffer, T., & Turner, J. (2006). Is incompatibilism intuitive? *Philosophy and Phenomenological Research*, 73, 28–53.
- Nichols, S., & Knobe, J. (2007). Moral responsibility and determinism: The cognitive science of folk intuitions. *Noûs*, 41(4), 663–685.
- Nichols, S., & Bruno, M. (2010). Intuitions about personal identity: An empirical study. *Philosophical Psychology*, 23(3), 293–312.
- Nijssen, S. R. R., Müller, B. C. N., van Baaren, R. B., & Paulus, M. (2019). Saving the robot or the human? Robots who feel deserve moral care. *Social Cognition*, 37(1), 41–56.
- Nyholm, S. (2018a). Attributing agency to automated systems: Reflections on human–robot collaborations and responsibility-loci. *Science and Engineering Ethics*, 24, 1201–1219.
- Nyholm, S. (2018b). The ethics of crashes with self-driving cars: A roadmap, I. *Philosophy Compass*, 13(7), e12507.
- Nyholm, S. R., & Frank, L. E. (2017). From sex robots to love robots: Is mutual love with a robot possible? In J. Danaher & N. McArthur (Eds.), *Robot sex: Social and legal implications*. MIT Press.
- Nyholm, S., & Smids, J. (2016). The ethics of accident-algorithms for self-driving cars: An applied trolley problem? *Ethical Theory and Moral Practice*, 19, 1275–1289.
- Nyholm, S., & Smids, J. (2019). Can a robot be a good colleague? *Science and Engineering Ethics*, 26, 2169–2188.
- Plakias, A. (2015). Experimental philosophy. *Oxford Handbooks Online*. <https://doi.org/10.1093/oxfordhb/9780199935314.013.17>.
- Pust, J. (2019). Intuition. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Accessed 9 Nov 2020. <https://plato-stanford-edu.ezproxy.library.wur.nl/archives/sum2019/entries/intuition/>.
- Scharff, R. C. (2011). Empirical technoscience studies in a Comtean world: Too much concreteness?. *Philosophy & Technology*, 25, 153–177.
- Schwitzgebel, E., & Cushman, F. (2012). Expertise in moral reasoning? Order effects on moral judgment in professional philosophers and non-philosophers. *Mind and Language*, 27(2), 135–153.
- Sorell, T. (2018). Experimental philosophy and the history of philosophy. *British Journal for the History of Philosophy*, 26(5), 829–849. <https://doi.org/10.1080/09608788.2017.1320971>.

- Swierstra, T., & Rip, A. (2007). Nano-ethics as NEST-ethics: Patterns of moral argumentation about new and emerging science and technology. *NanoEthics*, 1, 3–20. <https://doi.org/10.1007/s11569-007-0005-8>.
- Swierstra, T., Stemerding, D., & Boenink, M. (2009). Exploring techno-moral change: The case of the obesity pill. In P. Sollie & M. Duwell (Eds.), *Evaluating new technologies*. (pp. 119–138). Springer.
- Sytsma, J. (Ed.). (2014). *Advances in experimental philosophy of mind*. Bloomsbury Academic.
- Sytsma, J., & Livengood, J. (2016). *The theory and practice of experimental philosophy*. Broadview Press.
- Tigard, D. W. (2020). There is no techno-responsibility gap. *Philosophy & Technology*. <https://doi.org/10.1007/s13347-020-00414-7>.
- Van der Weele, C., & Tramper, J. (2014). Cultured meat: Every village its own factory? *Trends in Biotechnology*, 32(6), 294–296.
- Wiegmann, A., Okan, Y., & Nagel, J. (2012). Order effects in moral judgment. *Philosophical Philosophy*, 25(6), 813–836.
- Weinberg, J. M., Gonnerman, C., Buckner, C., & Alexander, J. (2010). Are philosophers expert intuiters? *Philosophical Psychology*, 23(3), 331–355.
- Williamson, T. (2004). Philosophical ‘intuitions’ and scepticism about judgment. *Dialectica*, 58(1), 109–153.
- Williamson, T. (2007). *The philosophy of philosophy*. Blackwell Publishing.
- Zwier, J., Blok, V., & Lemmens, P. (2016). Phenomenology and the empirical turn: A phenomenological analysis of postphenomenology. *Philosophy and Technology*, 29, 313–333. <https://doi.org/10.1007/s13347-016-0221-7>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.