

Technological Innovation and Natural Law

Philip Woodward

Valparaiso University, Valparaiso, IN, United States

philip.woodward@valpo.edu

Abstract

I discuss three tiers of technological innovation: mild innovation, or the acceleration by technology of a human activity aimed at a good; moderate innovation, or the obviation by technology of an activity aimed at a good; and radical innovation, or the altering by technology of the human condition so as to change what counts as a good. I argue that it is impossible to morally assess proposed innovations within any of these three tiers unless we rehabilitate a natural-law ethical framework. And I offer some moral starting points within such a framework, in connection with innovations of each of the three types.

Keywords

emerging technology – natural-law theory

1 Introduction

The modern pursuit of social improvement through public policy has been successful on many fronts: it has resulted in the spread of democracy and individual liberty, improved health and prosperity, and so on. The successes of the past might lead us to think that our remaining social ills can, likewise, be engineered out of existence. But this is probably not the case. In their 1973 paper “Dilemmas in a General Theory of Planning,” Horst Rittel and Melvin Webber argue that looming social problems are “wicked,” in contrast to the “tame” problems that engineers are accustomed to working through. There are three main features of looming social problems that make them wicked:

1. *Value-relativity*. What counts as a social problem, and whether the problem has been solved, is relative to a set of societal values, and these values are contested.
2. *Unboundedness*. Social phenomena are embedded in a network of causes and effects that has no apparent boundary. The causes of social problems, and the effects of interventions on social problems, can include any aspect of reality whatever (including the distant past and future).
3. *Unrepeatability*. Every social problem is uniquely complex, and each attempted intervention on a social problem permanently alters society (no “trial-and-error” process is possible).

My concern is not with social problems per se, but with problems of a different (though closely related) sort: *technological implementation problems*. A technological implementation problem is the problem of deciding whether a society ought to incorporate an innovative technology into its life. Like social problems, technological implementation problems exhibit the three signs of wickedness.

Technological implementation problems exhibit *unrepeatability* in the sense that innovations quickly “take root” and spread, owing to economic and social pressure. (For example, telephones and computers have become so central to the functioning of most industries that it is virtually impossible to participate in the Western economy without access to them.) This means that there is no going back once a technology is embraced by society.

Moreover, it is very difficult to predict the full social consequences of implementation. Herein lies the *unboundedness* of technological implementation problems. Unrepeatability and unboundedness exacerbate each other: implementing a new technology comes with the risk that we’ll encounter nasty social consequences but be unable to reverse course.

But matters are even worse than this. Technological implementation problems exhibit *value-relativity*, because we lack a shared framework for assessing the social consequences of innovations. The two dominant ethical frameworks used by applied ethicists today (including technological and engineering ethicists) are deontology and consequentialism. Deontological theories (associated with Immanuel Kant and followers) begin with rules for action: they tell us, as we pursue our ends, which behaviors we must, or must not, engage in. Consequentialist theories (associated with John Stuart Mill and followers) start not with rules for action but with a valuable state of affairs that our actions should aim to maximize—in classic formulations, pleasure or happiness (alternatively, a disvaluable state of affairs our actions should minimize—i.e., pain).

A number of scholars have noted the limitations of these two frameworks (or, at any rate, mainstream versions of them) when it comes to addressing

ethical questions about emerging technologies (see, for instance, Schmidt 2014; Vallor 2016). The trouble, in brief, is that neither framework is well positioned to tell us which ends we should be pursuing as a society. This is because the two frameworks, despite their great differences, tend to converge on the same moral sensibility: *respecting individual differences*. (To simplify: deontology arrives at this sensibility via its emphasis on respecting the autonomous wills of others; consequentialism gets there via its subjectivist conception of happiness.) The importance of this moral sensibility is not under dispute. But it fails to address the question, central to technological implementation problems, of what would make a society as a whole better or worse.¹

In sum: technological implementation problems are wicked because technological innovations are *unrepeatable* (they tend to “stick” once introduced), *unbounded* (they cause society-wide changes in the long run that are hard to predict), and recommendable only in a *value-relative* way (the predominant ethical frameworks appealed to in mainstream technological ethics today do not provide us with shared standards for evaluating them). To adapt Otto Neurath’s metaphor: we are fixing our ship while at sea—and the workers can’t agree on whether the ship is a pleasure yacht or a freighter.

To tame the wickedness of technological implementation problems, we will first of all need a different ethical framework. I propose that *natural-law theory* is the right tool for the job.² This is the framework that was dominant from antiquity up to the early modern period. Recognizing that there are multiple, incompatible ethical theories that bear the label “natural law,” I will work with the historically most influential version, viz., that of Thomas Aquinas. According to Aquinas, natural law is the rational self-governance of human beings “according to their nature.” That is, natural law is humanity’s exercise of practical reason in its endeavor to secure the basic goods toward which its

1 I am assuming, that is, that there is more to the societal good than everyone’s respecting each other’s differences. One reason to think so is the ubiquitous need for coordination in a healthy society. Aspects of life as diverse as the judicial system, public utilities, public health, education, and sport cannot function in the absence of widely shared objectives.

2 Schmidt (2014) and Vallor (2016) both propose *virtue ethics* as an alternative (or supplement) to deontology and consequentialism. Their thought is, roughly, that we must cultivate the virtue of wisdom in order to be able to discern which emerging technologies will be salutary. I suspect this is right, so far as it goes. But to say that we must learn to be discerning with respect to technology is not yet to give us criteria for making judgments in particular cases. Virtue ethicists tend to eschew “decision procedures” for right action. Given that our driving question is precisely a request for such a decision procedure (in particular, a procedure for deciding whether to embrace an innovative technology), virtue ethics does not attempt to provide an alternative answer to that of deontology and consequentialism.

nature is objectively ordered—the goods the attaining of which makes for the objective flourishing of human beings as such.

The phrase *human nature* is used in a number of different ways (perhaps even by Aquinas himself). I use it to denote the set of capacities the joint possession of which makes a thing an instance of the kind *human*, and the joint actualization of which makes for *thriving* as a human. The possession of these capacities is consistent with their not being exercised or even with their being seriously impaired or undeveloped. The “basic goods” of human nature will be fixed by these capacities in some way—either as that at which a capacity is aimed or that which is requisite for a capacity to function—in tandem with contingent facts about the broader human condition (the limits of the earth’s resources, for example). These capacities can come in a range of strengths, and this range makes, in part, for natural diversity within the kind *human*.

What *are* the basic goods toward which human nature is ordered? Following Aquinas, we can group them into four categories:

1. preservation of life, health, and safety;
2. reproduction and the nurture of the young;
3. the acquisition of knowledge, especially about God;
4. life in society with other humans.

I’ll call these four types of goods “biotic,” “animal,” “rational,” and “communal,” respectively.³ This list is very abstract, of course. Presumably, each of the four types of goods subdivides into a host of specific goods; I will invoke more specific goods as they become relevant below.⁴ One important clarification is needed at this point. Aquinas stresses that natural law is ordered toward “the common good.” Following Jacques Maritain (1946), I will take this to mean that basic human goods always simultaneously benefit individual persons *and* the society those persons jointly constitute.

Now, natural-law theory has been subjected to criticisms of its own. I will say nothing in response to these criticisms.⁵ My aim is simply to convince the reader that (one version of) natural-law theory has the resources—as rival ethical frameworks do not—to tame the wickedness of technological implementation problems. I frame my discussion in terms of a distinction

3 Following Aristotle, Aquinas grounds both (3) and (4) in our “rational nature.”

4 Limits of space prevent me from generating an exhaustive list of such goods, let alone defending the inclusion of each item on it. For others’ attempts at constructing and defending such lists, see Weil (1952), Finnis (1980), Murphy (2001), Gomez-Lobo (2002), and Kraut (2009).

5 For an excellent discussion of some of the major objections to natural-law theory and an attempt to address them, see Murphy (2011).

between three “tiers” of technological innovation: *mild*, *moderate*, and *radical*. I explain how implementation problems within each tier exhibit value-relativity, unboundedness, and unrepeatability. Then I discuss how natural law in each case dissolves value-relativity (by providing us with objective standards of thriving) and mitigates unboundedness and unrepeatability (by recommending guidelines for decision-making as well as methods for ensuring that long-term effects are salutary).

2 Mild Innovation

We can think of technological innovations as falling into three different tiers: mild, moderate, and radical. In brief, *mild* innovation accelerates or renders more efficient some human activity that is a means toward an end; *moderate* innovation replaces an activity—i.e., it makes an end accessible without engaging in the activity that was traditionally the means required for obtaining it; and *radical* innovation obviates an end and with it any of the activities previously used to obtain it.⁶ By way of illustration, consider logging, the human activity of harvesting trees for their wood. Mild innovation in the domain of logging has occurred via the development of more powerful saws. Moderate innovation would occur if robotic “self-driving” saws were developed. Radical innovation would occur if we developed thermal exoskeletons that obviated the need for shelter of any kind.

Mild technological innovations accelerate or render more efficient a human activity that is a means toward some end. Practically all of the familiar advances in technology throughout history have been more or less mild (in my technical sense of that term). Developments in weaponry, in medicine, in transport, in agriculture, in construction, in manufacturing, and so forth, have rendered a human activity—warring, healing, traveling, acquiring food, erecting shelter, making household goods—more efficient or effective.

I’ll call problems regarding whether to implement a mildly innovative technology *upgrade cases*. Upgrade cases exhibit the three signs of wickedness. Consider, for example, our ongoing implementation of digital information technologies (especially internet-connected devices and the software that

6 There are, of course, other ways to taxonomize technological innovations. Mine is certainly not the most comprehensive taxonomy: it does not, for instance, cover improvements in the aesthetics or safety of human creations. The virtue of my taxonomy is that it recognizes different *kinds* of ways (and not just different *degrees*) in which technologies can alter human life. Sorting technologies according to this type of difference among them assists us in finding the right starting points for our prudential reasoning about them.

runs on them). This is an upgrade case, in that these technologies accelerate our capacity to communicate and retrieve information. The case has exhibited value-relativity insofar as social groups have disagreed about which values are relevant to assessing these technologies. For example, in the context of grade school education, implementation enthusiasts have appealed to the value of *competitiveness in the global economy* whereas critics have appealed to the value of *the joy of immersive play and learning*.

The case has exhibited unboundedness insofar as implementation has restructured society in unanticipated ways: it has allowed political dissidents to organize; changed the ways that romantic relationships form and develop; made Western entertainment cheaply available nearly everywhere in the world; facilitated new forms of espionage and geopolitical influence (such as election tampering); created whole new economic sectors and put other sectors out of business; and much else of which we are as yet only dimly aware.

Finally, it has exhibited unrepeatability insofar as the new technology is “sticky”: we can’t imagine letting go of it or lessening our dependence on it.

Has the implementation of emerging digital communication been good for society? Should we continue to implement the next generation of such tech, whatever it looks like? It should be clear by now just how unwieldy these questions are. But applying the natural-law framework to mild innovations can help us get our bearings: we simply need to ask whether the technology in question secures genuine human goods without quashing others. Now, it’s clear that digital information technologies have helped secure genuine human goods—both rational goods (such as accessing expert information) and communal goods (such as communicating with friends and family at great distances). But it is becoming increasingly apparent that our internet-connected devices are also making certain rational and communal goods *harder* to secure. Andrew Sullivan (2016) describes his growing realization of what his own smartphone was doing to him as follows:

I tried reading books, but that skill now began to elude me. After a couple of pages, my fingers twitched for a keyboard. I tried meditation, but my mind bucked and bridled as I tried to still it. I got a steady workout routine, and it gave me the only relief I could measure for an hour or so a day. But over time in this pervasive virtual world, the online clamor grew louder and louder. Although I spent hours each day, alone and silent, attached to a laptop, it felt as if I were in a constant cacophonous crowd of words and images, sounds and ideas, emotions and tirades—a wind tunnel of deafening, deadening noise.

The few scientific studies that have been done on the topic of the psychological effect of smartphone use suggest that Sullivan's experience is typical: these devices distract their users continuously (even when turned off), resulting in measurably diminished cognitive and interpersonal capacities (Ward et al. 2017; Dwyer, Kushlev, and Dunn 2017).

So, even though smartphones have accelerated our access to certain rational and communal goods (viz., information availability and long-distance communication), they are also inhibiting our access to others (viz., focused, immersive learning and socializing). If the force of acceleration were quite strong and the inhibitory force quite weak, natural law might encourage smartphone usage. But the reverse seems true: smartphones are probably inhibiting the thriving of their users more than they are facilitating it. If that is right, natural law pronounces a negative judgment on the proliferation of smartphone usage (though culture-specific details will no doubt be relevant).

I have been discussing one mild innovation in particular that is currently being implemented. But natural law allows us to formulate more general strategies for taming upgrade cases in the future. Most obviously, natural law helps tame value-relativity by supplying us with the objective goods that any worthy innovation must help us secure. And these goods are not just objective but *common*. That is, natural law requires us to ask whether an innovation will secure objective human goods *for society as a whole*, when implemented *by* society as whole.

What about unboundedness and unrepeatability? Natural-law theory cannot itself offer insight into the likely downstream effects of innovation, of course. But it does counsel a particular rule of thumb under circumstances of unboundedness and unrepeatability: *caution*. If a society is flourishing—if it has a means of securing its objective, common goods—we are not *sacrificing* great goods by opting out of new, more efficient technology that might make other goods harder to secure in the long run. (Thus, natural law helps mitigate the worry that if we say no to an innovative technology, we could be missing out on great goods.) That is, a proposed innovation that would merely *accelerate* society's access to certain goods but that, for all we know, would in the long run actually thwart society's access to other goods is not worth the risk. Of course, if the society *isn't* flourishing, the calculation will be different. Note that this stance of caution is not expressive of a general technophobia or of a worry that technology is "unnatural." A society is "natural" insofar as its basic goods are secured—with or without the assistance of technology.

Natural law, I have been stressing, appeals to the objective, common goods of "a society" as the standard whereby innovations are adjudicated. What sort

of social group counts as a “society”? There are two types of answers: an ideal answer and a fallback answer.

The ideal answer is *any group of people who substantially share a culture*. This criterion, incomplete and vague as it may be, suggests at least that societies are pretty large (e.g., North Americans and Europeans are probably all part of one society). There is a problem with this answer, however. Today’s liberal societies avoid mandating the use of or abstention from particular technologies (or other cultural artifacts). Better to let individuals choose, we have decided, than to give governments the power to impose culture in a top-down way—a power that has consistently been abused.

The fallback answer is *any group of people who care to coordinate a common life with one another*. Societies so understood might be quite small—as small as individual families. But they might be larger. In more homogeneous societies, they might be municipalities; in more pluralistic societies, they might be geographically disparate subcultures. The trouble with the fallback answer is that these smaller communities might yet be subject to overwhelming social and economic pressures from the broader society to implement (or abstain from) certain technologies that they judge deleterious (or salutary). But these communities will at least have some agency over their technological habits.

In short, natural law tames the wickedness of upgrade cases by pointing to objective measures of the flourishing of a society. On the basis of these objective measures, to the extent that a society is already flourishing, natural law recommends a cautious, guilty-until-proven-innocent stance toward mild innovations. To would-be innovators we should say: “You wish to accelerate our means of securing good X. As we already have access to X, unless you can convince us that your accelerated means of access will not worsen our lives in other ways, we decline your offer.”

3 Moderate Innovation

Moderate technological innovations make an end accessible without engaging in the activity that was traditionally the means required for securing it. If we individuate activities by their *means*, then many mild innovations satisfy this description: the printing press rendered obsolete the activity of hand transcription; the typewriter rendered obsolete the activity of typesetting; perhaps voice-to-text software will soon render obsolete the activity of typing. If we individuate activities by their *ends*, then it turns out that moderate innovations have been pretty rare. (Monastic scribery, printing with movable type, and typewriting all share the end of reproducing written text, after all.)

But they are becoming more common, with the development of robotics and artificial intelligence (AI). These technologies are steadily replacing human workers in fields such as construction, manufacturing, banking and financial management, and health care. And there is rampant speculation about eventual applications in such fields as jurisprudence, scientific inquiry, and even the fine arts.⁷

I'll call problems regarding whether to implement a moderately innovative technology *automation cases*. The value-relativity of automation cases is apparent in the ways that proponents and critics of automation argue with one another. Proponents appeal to emancipation: they argue that automation will free human beings from work that is dull, degrading, and poorly paid (e.g., Kelly 2016). Strikingly, critics seem to agree that automation will have an emancipatory effect, but disagree that such an effect would be salutary: it would disrupt the communal interdependence engendered by traditional means of production. Wendell Berry, for example, has gone so far as to say that "the triumph of the industrial economy is the fall of community" (Berry 1992, 133).

Automation cases exhibit unboundedness in that we are unable to tell at this point how much of human life can be automated, at what cost, and on what timeline. Over the last decade, AI has improved dramatically: it can now match or exceed human performance at cognitive skills such as pattern recognition. Some experts have predicted a transformation of the workforce on the scale of the industrial revolution within a generation (Frey and Osborne 2017; Ford 2015). But most of AI's achievements have occurred in the context of narrowly proscribed tasks, such as playing chess or categorizing elements of an image. These task-specific AIs lack a "common sense" that allows them to correct errors or to know when to give up on their assigned task. Most jobs do require the use of common sense from time to time. In order to match human performance at common-sense reasoning, AIs may need to be programmed as artificial *general* intelligences (in addition to whatever task-specific skills they need). Predictions regarding when we will develop artificial general intelligence range from a few years from now to a few hundred years from now (Brooks 2019). Whether automation is about to have a colossal, snowballing effect on human labor or whether it will only ever have a minor, circumscribed effect is not something we can predict before the technology is developed.

7 HRP-4C, a singing and dancing robot, is celebrated as a pop star in Japan. (A video of her performance is available here: <https://www.youtube.com/watch?v=8LXACLbFRKs>.) Robots have recently composed music (Kalegasi 2017), generated images resembling the work of master painters (Schneider and Rea 2018), and written—admittedly incomprehensible—screenplays (Miyamoto 2016).

And finally, it's clear that automation cases exhibit unrepeatability. The driving force behind automation is the promise of economic efficiency. To the extent that it makes good on that promise, there will be enormous economic pressure to implement it, as rival firms compete with each other. Once automation takes hold, it will be here to stay.

Natural law can help us navigate these difficulties. First of all, natural law resolves the standoff between those who appeal to emancipation and those who appeal to communal interdependence as the relevant criterion. From the perspective of natural law, neither of these values is an intrinsic good. Rather, emancipation is an instrumental good insofar as it sets people at leisure to engage in activities that manifest rational goods. And interdependence is likewise an instrumental good insofar as it facilitates communal goods—familiarity, intimacy, a shared way of life, etc. It is these intrinsic (rational and communal) goods that are the nonnegotiables. In certain cases, the best way to attain these goods will involve embracing moderate innovation—particularly, in cases where newfound human leisure will make rational goods more readily attainable without simultaneously thwarting communal goods. Leisure is, as Joseph Pieper (1963) stresses, “the basis of culture”: a person at leisure is free to pursue art, learning, contemplation, and other deeply valuable activities that may not have much market value. No doubt moderate innovation also has the capacity to disrupt communal ways of life. But there can be other ways, besides collaborative economic projects, to cultivate communal goods. (I expand on this point below.)

At the same time, natural law is inconsistent with some types of moderate innovation. Some human activities are not mere means to intrinsic human goods but actually manifest those goods. Here are three human activities that plausibly *could* be automated into oblivion but shouldn't:

1. *The rational goods of aesthetic creativity, performance, and appreciation.* Art is an intrinsic good in human life; it would thus be at odds with human thriving to automate the creation or performance of art.⁸ Still sillier would be to try to automate the human activity of the *appreciation* of art, for delight in beauty is likewise an intrinsic human good.
2. *The communal goods of intimacy and care.* Love is an intrinsic good in human life; it would thus be at odds with human thriving to automate caring encounters between humans. For example, parenting, elder care, health care, social work, and psychotherapy, as well as education in most of its forms, are all activities that should not be short-circuited by

8 There is nothing wrong with creating a robot that makes paintings (say). The trouble is if doing so is meant to, or turns out to, eliminate from human life the same activity.

automation, because they afford opportunities for love to be given and received in substantive ways.⁹

3. *The rational good of practical deliberation.* The exercise of practical reason is an intrinsic good in human life. One reason to think so is that it is one of the few psychological capacities that humans cannot “automate” in themselves: deliberation about rival plans of action in new situations always occurs consciously in humans, unlike nearly every other psychological process (Morsella 2005; Baumeister, Masicampo, and Vohs 2011). It is thus a pre-eminent source of human distinction and dignity.¹⁰ And this means that there should not be robot lawmakers, judges, ethical consultants, and suchlike.

So, natural law can help tame the value-relativity of automation cases, allowing us to discern which moderate innovations would be good and which would be bad for society. What about unboundedness and unrepeatability? Natural law certainly cannot tell us whether and when AI will emancipate us from our work, but it can tell us how to ensure that our newfound leisure facilitates our thriving rather than—as Wendell Berry fears—detracting from it.

First of all, unemployment brings with it loss of income and ensuing loss of biotic goods. Whether automation will create as many new jobs (in a reasonable time frame) as it will replace is a matter of dispute among economists. If it does not, society will have to find ways to de-couple work and income.

Second, leisure can be put to objectively bad uses. For example, in the 2008 Pixar film *Wall-E*, human beings spend their waking hours on floating lounges, receiving an endless stream of video entertainment and good-tasting food. They are (subjectively) happy, but communally and rationally diminished. When given the opportunity to live the indulgent life, untutored human nature puts up little resistance.

The Christian monastic tradition has long distinguished between leisure as mere idleness—the “enemy of the soul,” as St. Benedict says—and leisure as soul cultivation, “a time and space of freedom in which the deep self can find fuller expression” (Casey 2005). Monastics carve out this “space of freedom” by submitting to a strict “rule of life,” a set of norms and practices that order the community toward the good, especially the good of knowing and loving

9 This is not to say that there are no appropriate roles for automation in care-giving. For detailed discussions, see Vallor (2016, chapter 9), VanderLeest and Schuurman (2015), and Schuurman (2016).

10 See Murphy (2001, 116–118) for a different argument that deliberating well is intrinsically and not merely instrumentally good for humans.

God. The idea of freedom through submission may strike one as paradoxical (or worse), but in fact it captures an Aristotelian truism: humans need to be tutored in virtue before they can flourish. This schooling in virtue can happen in a variety of ways. In the past, work and other necessities of life have incidentally provided it. The more emancipated we are from the strictures of work, the more we will need to impose structure on our lives from without. Thus, leisure and asceticism must go hand in hand.

I observed in the last section that centralized, top-down cultural enforcement has frequently inhibited human thriving. Who, then, shall be our cultural “abbot” who enforces a society-wide rule of life? Paralleling what I said above, the only viable answer is that there shall *not* be such a centralized authority—at least not one that administers a full-blown rule. Rather, groups of people who care to coordinate a common life will freely submit themselves to rules they develop together. At the same time, if, as automation spreads, people use their newfound leisure in increasingly vapid, self-degrading ways, it may be necessary for our society to revive a “paternalistic” conception of central authority. Indeed, that conception has never fully gone away. (Note that elementary schooling is a requirement in many liberal societies. Why? Surely the answer is that such schooling is conducive to the common good.)

In sum: though we cannot know the pace at which moderate innovations will replace human activities, we don’t need to be terribly worried about this—provided (1) we safeguard those activities that are intrinsically good for us and (2) we are prepared to subject our newfound leisure to rules of life that orient us toward our communal and rational thriving.

4 Radical Innovation

Whereas mild innovations accelerate human activities and moderate innovations replace human activities, radical innovations remove or alter the ends toward which those activities have been directed by modifying the essential human capacities that fix what is objectively good for humans. Now, human capacities have no doubt been modified over the course of human history by evolutionary and cultural forces (McKenny 2018, 31). Intentional alterations of essential human capacities are basically unprecedented, however. If recent developments in biotechnology—of implants, prosthetics, cognitive enhancers, and gene-editing—proceed uninhibited, we will soon be in a position to tinker with human traits such as physical features, lifespan, childhood developmental stages and pace, cognitive capacity, social and

emotional temperament, physical strength and vulnerability, and perhaps even religious propensity.¹¹

I'll call problems regarding whether to implement a radically innovative technology *enhancement cases*. Here is how enhancement cases exhibit the three signs of wickedness. Regarding value-relativity: according to critics, radical innovation *violates* human dignity, because it involves taking a pejorative, controlling stance toward our fellows or our offspring. According to proponents, radical innovation *upholds* human dignity, because it manifests the uniquely human capacity to freely shape our world and ourselves in salutary ways. Thus, a debate drags on among bioethicists as to whether the givenness or the open-endedness of the human condition is the more salient value relevant to enhancement cases (McKenny 2018, chapters 2 and 4).

Enhancement cases exhibit unboundedness of a particularly bewildering sort. Anticipating the long-term effects of enhancement requires, among other things, anticipating the choices of beings whose priorities might differ from ours. For example, suppose we undergo enhancement with respect to (a) our moral psychology, such that we instinctively care about long-term concerns as much as short-term ones; and (b) our physical endowments, such that extreme temperatures cease to bother us. Would the human race, so enhanced, finally solve our global environmental crisis? Or would it cease to see it as a crisis? We simply cannot know in which direction altered human nature will steer the world—in the direction of paradise or in the direction of Armageddon.

Finally, enhancement cases exhibit unrepeatability insofar as nonenhanced persons will feel social and economic pressure to join the enhanced crowd. Suppose that some parents give their children cognitive enhancing drugs and others do not. Presumably the enhanced children will learn more quickly, get better grades, and ultimately pursue more lucrative and fulfilling careers. Anti-enhancement parents will feel that they are holding their children back.

At first glance, natural law seems ill suited to tame the wickedness of enhancement cases, because it holds up goods of human nature as the standards with respect to which innovations are assessed. But radical innovations seek to alter those standards: "Human nature becomes a variable rather than a constant" (McKenny 2018, 70). Asking whether to embrace a radically innovative technology means asking whether to remain human or to become a thing of a different nature with a different set of basic goods.

11 This list draws on discussions by Vallor (2016) and Agar (2010). Agar mentions geneticist Dean Hamer's claim to have identified the gene that regulates the aptitude for religious experience.

But there are versions of the natural-law framework that are still relevant. Let us call a race of radically enhanced human beings “posthumans” (Bostrom 2008). It no longer makes sense to use a *human* standard to determine whether such beings are thriving, any more than it makes sense to ask whether a frog is a good fish. But humans and posthumans have something important in common: they are *persons*. I suggest that there is a law of self-governance for all persons, human or not—a *personal* natural law. Humanity and posthumanity each stand to personhood as determinate to determinable. Consequently, the list of basic goods for persons will be less specific and probably shorter than the list of basic goods for humans, but there will still be such a list. Presumably, clothing and shelter won’t be on it. But rational and communal goods such as union with God and friendship *will* be on the list, since (I submit) these are goods for persons as such.¹²

This personal natural law can help us tame the wickedness of enhancement problems. Regarding value-relativity: the trouble with the appeal to givenness is that it provides no positive vision for biotechnology (for example, it cannot tell us what is good and right about modern medicine), and the trouble with the appeal to open-endedness is that it provides no *limit* to biotechnology. But the personal natural law gives us a standard that serves both as target and limit—viz., that which is conducive to the thriving of persons as such. In effect, innovations that count as radical from the perspective of *human nature* can now be treated as mild or moderate from the perspective of *personal nature*, so we at least have somewhere to begin our deliberations.

The personal natural law will certainly rule out some radical enhancements as contrary to the objective good for persons. For example, it would be contrary to our nature as persons to substantially alter our need for friendship or union with God (say, by excising loneliness or spiritual longing from our psychological profile), as these are central to our thriving as persons. Indirectly, our needs for education and for parenting facilitate the communal goods of personhood by drawing us into relationships of vulnerability and care with other persons. These needs should be altered only if other mechanisms for interpersonal intimacy are in place.

Plenty of radical innovations will not be ruled out by these considerations, which brings us back to the matters of unboundedness and unrepeatability. How do we decide whether to implement a proposed enhancement (that is not immediately ruled out by the personal natural law), given that we cannot know its long-term effects and that we cannot “take it back” once it is implemented?

12 This claim is familiar from the so-called “personalist” tradition in Catholic ethical thought. See, e.g., Maritain (1946).

In order to answer this question, we need to distinguish between two ways that the relationship between human nature and the goods of persons might be understood. On one view, the goods of persons are attainable equally well—qualitatively, if not quantitatively—by humans as by posthumans; the advantages posthumans enjoy are unrelated to their thriving as persons—for example, they are merely freed from certain physical and cognitive limitations characteristic of humans. On the second view, the goods of persons are *better* attained (both qualitatively and quantitatively) by posthumans than by humans. The idea here would be that certain dimensions of human nature—our limited lifespan, physical vulnerabilities, cognitive bandwidth, and so on—actually inhibit the thriving of other dimensions of our nature—viz., the rational and communal capacities that make us persons.

Suppose the first view is correct: humans don't *need* to become posthuman in order to thrive as persons. We can now treat enhancement cases as mere upgrade cases, from the perspective of the personal natural law. And the criterion we reached in connection with upgrade cases is directly applicable: a proposed innovation that would merely accelerate access to personal goods (or to a greater quantity of them) but that, for all we know, would in the long run actually thwart our access to other goods, is not worth the risk. We should let sleeping dogs lie.

But maybe the first view isn't correct. Might we *need* to enhance ourselves in order to thrive as persons? After all, there are only so many academic subjects one can master in a *human* lifespan, there are only so many friends one can know intimately given *human* memory capacities, and so on. Strikingly, Aquinas seems to endorse an idea along these lines. The beatific vision (our ultimate end), he writes, "is beyond the natural power of any created intellect" (*Summa theologiae* I, Q.12, A.4). Perhaps our capacities must exceed the essentially human range before we can fully thrive as persons.

This question of whether we must cast off our human nature in order to thrive as persons can be answered, I believe, only by grappling with some of the details of Christian theology. Christians affirm that God has become united to human nature in the incarnation, suffering, and death of Christ and has revealed what fully thriving human nature is like in the resurrection and ascension of Christ. Christians understand this arc of descent to death and ascent to glory to reveal the most important truths about human nature: first, that human nature itself—and not some superhuman or posthuman alternative—is capable of the greatest thriving we could wish for; second, that such thriving is possible *because of* and not *in spite of* our human limitations and vulnerabilities.¹³

¹³ For a much more thorough treatment of these matters, see McKenny (2011; 2018). My discussion here is deeply indebted to McKenny's.

How could this be? The answer is that God's grace is needed in order for human nature's highest capacities to be activated. (This is what Aquinas means by saying that the beatific vision is beyond our natural powers.) And we can receive God's grace only when, accepting our status as limited, vulnerable creatures, we abandon ourselves to God's care. Our limitations and vulnerabilities, that is, create the conditions most suitable to occupying a posture of willed dependence on God. Writes St. Paul, "I am therefore happy to boast of my weaknesses, because then the power of Christ will rest upon me" (1 Cor. 12:9). If this Christian vision of human nature is correct, we do not need to transcend our human nature in order to experience full personal thriving. In fact, if we try to, we might actually hurt our chances at full personal thriving, because we might remove the conditions that teach us to depend on God's grace.

This is not to say that the personal natural law rules out all tinkering with our human capacities. It will probably recommend it in certain cases—for example, it would be in keeping with the personal natural law to splice immunizations into our genetic code. But really radical enhancements are not wise: we can't predict whether their long-term effects will be salutary, and we don't need them in order to thrive, anyway. We might actually need to reject them in order to thrive.

5 Conclusion

Let us review the discussion so far. Emerging technologies are positioned to alter the human condition to an unprecedented degree. Careful discernment is called for, but the dominant frameworks used by applied ethicists today do not equip us to do this. As Shannon Vallor (2016, 241) puts it: "The real problem is *knowing what it is that we ought to wish for*. . . . In a neoliberal age, who among us has the courage and genuinely magnanimous moral leadership to point the way to a positive vision of the human future?"

The natural-law framework can fill the void here. I applied this framework to three tiers of technological innovations and concluded the following:

- Regarding *mild* innovations, natural law lists the objective human goods that proposed technological upgrades must be directed toward if they are to earn admittance into human life; and even regarding those upgrades that *are* so ordered, it recommends a guilty-until-proven-innocent stance.
- Regarding *moderate* innovations, natural law forbids the automation of certain intrinsically good human activities; it permits the automation of

plenty of others, but only to the extent that resulting leisure-time in human life is filled with objectively good activities.

- Regarding *radical* innovation, natural law points to the objective goods *for persons* as a fallback standard (once human nature ceases to be a given); regarding enhancements that appear to honor this standard, it recommends a guilty-until-proven-innocent stance, especially insofar as such enhancements target human limits and vulnerabilities conducive to the reception of divine grace.

At every stage the discussion has focused on what *society* must do in order to secure the *common* good. I have twice alluded to the difficulty of implementing the findings of natural law: liberal societies are simply not set up to police citizens' use of technology. One might take the concern one step further. Perhaps we should not even *develop* a "positive vision of the human future." Pluralism is nonnegotiable. We must celebrate the freedom of each individual to follow her own vision of the future. The alternative, as history has shown, is fascism or war.

But we must be careful not to celebrate ideological pluralism inordinately. Tolerance surely has *some* limits in a civil society. In particular, one ought not to tolerate behavior that undermines the possibility of civil society at all. It is a difficult matter to determine where the line is, but surely some uses of technology are so devastating that they should not be tolerated. The limits of tolerance narrow further if we are interested not just in the *preservation* of society but in its *flourishing*. Societies exist only when some values are shared among citizens, and they flourish only when a lot of values are shared among citizens.

Natural-law theory is a high-risk/high-gain moral framework. Because it requires us to agree on the human good in considerable detail, it risks the pain of interminable moral disagreement. But it holds out the promise of a flourishing society in return. There is, then, a third alternative to unbounded ideological pluralism (besides fascism or war): the hard work of pursuing the truth about the human good together in the public square.

References

- Agar, N. (2010). *Humanity's End: Why We Should Reject Radical Enhancement*. Cambridge: MIT Press.
- Baumeister, R.F., Masicampo, E.J., and Vohs, K.D. (2011). Do Conscious Thoughts Cause Behavior? *Annual Review of Psychology* 62, pp. 331–361.
- Berry, W. (1992). *Sex, Economy, Freedom, and Community: Eight Essays*. New York: Pantheon Books.

- Bostrom, N. (2008). Why I Want to Be a Posthuman When I Grow Up. In: B. Gordijn and R. Chadwick, eds., *Medical Enhancement and Posthumanity*, Dordrecht: Springer, pp. 107–136.
- Brooks, R. (2019). AGI Has Been Delayed. <https://rodneymbrooks.com/agi-has-been-delayed/>.
- Casey, M. (2005). *Strangers to the City: Reflections on the Beliefs and Values of the Rule of St. Benedict*. Orleans: Paraclete Press.
- Dwyer, R.J., Kushlev, K., and Dunn, E.W. (2017). Smartphone Use Undermines Enjoyment of Face-to-Face Social Interactions. *Journal of Experimental Social Psychology* 78, pp. 233–239.
- Finnis, J. (1980). *Natural Law and Natural Rights*. New York: Oxford University Press.
- Ford, M. (2015). *Rise of the Robots: Technology and the Threat of a Jobless Future*. New York: Basic Books.
- Frey, C.B., and Osborne, M.A. (2017). The Future of Employment: How Susceptible Are Jobs to Computerisation? *Technological Forecasting and Social Change* 114, pp. 254–280.
- Gomez-Lobo, A. (2002). *Morality and the Human Goods: An Introduction to Natural Law Ethics*. Washington: Georgetown University Press.
- Kaleagasi, B. (2017). A New AI Can Write Music as Well as a Human Composer. *Futurism*, March 2017. <https://futurism.com/a-new-ai-can-write-music-as-well-as-a-human-composer>.
- Kelly, K. (2016). *The Inevitable: Understanding the 12 Technological Forces That Will Shape Our Future*. New York: Penguin.
- Kraut, R. (2009). *What Is Good and Why: The Ethics of Well-Being*. Cambridge: Harvard University Press.
- Maritain, J., and Fitzgerald, J.J. (1946). The Person and the Common Good. *The Review of Politics* 8 (4), pp. 419–455.
- McKenny, G. (2011). Transcendence, Technological Enhancement, and Christian Theology. In: R. Cole-Turner, ed., *Transhumanism and Transcendence: Christian Hope in an Age of Technological Enhancement*, Washington: Georgetown University Press, pp. 177–192.
- McKenny, G. (2018). *Biotechnology, Human Nature, and Christian Ethics*. Cambridge: Cambridge University Press.
- Miyamoto, K. (2016). A.I. Writes a Screenplay, Screenwriters Breathe a Sigh of Relief. *Screencraft*, June 2016. <https://screencraft.org/2016/06/10/a-i-writes-a-screenplay-screenwriters-breathe-a-sigh-of-relief/>.
- Morsella, E. (2005). The Function of Phenomenal States: Supramodular Interaction Theory. *Psychological Review* 112 (4), pp. 1000–1021.
- Murphy, M.C. (2001). *Natural Law and Practical Rationality*. Cambridge: Cambridge University Press.

- Murphy, M.C. (2011). *God and Moral Law: On the Theistic Explanation of Morality*. New York: Oxford University Press.
- Pieper, J. (1963). *Leisure: The Basis of Culture*. Translated by A. Dru. San Francisco: Ignatius Press.
- Rittel, H.W., and Webber, M.M. (1973). Dilemmas in a General Theory of Planning. *Policy Sciences* 4 (2), pp. 155–169.
- Schmidt, J.A. (2014). Changing the Paradigm for Engineering Ethics. *Science and Engineering Ethics* 20 (4), pp. 985–1010.
- Schneider, T., and Rea, N. (2018). Has Artificial Intelligence Given Us the Next Great Art Movement? *Artnet News*, September 2018. <https://news.artnet.com/art-world/ai-art-comes-to-market-is-it-worth-the-hype-1352011>.
- Schuurman, D. (2016). Responsible Automation: Faith and Work in an Age of Intelligent Machines. In: D.H. Kim, ed., *The Wonder and Fear of Technology: Commissioned Essays on Faith and Technology*, New York: Center for Faith and Work, pp. 42–56.
- Sullivan, A. (2016). I Used to Be a Human Being. *New York Magazine*, September 19. <https://nymag.com/intelligencer/2016/09/andrew-sullivan-my-distraction-sickness-and-yours.html>.
- Vallor, S. (2016). *Technology and the Virtues: A Philosophical Guide to a Future Worth Wanting*. New York: Oxford University Press.
- VanderLeest, S., and Schuurman, D. (2015). A Christian Perspective on Artificial Intelligence: How Should Christians Think about Thinking Machines? In: *Proceedings of the 2015 Christian Engineering Conference*, Seattle: Seattle Pacific University, pp. 91–107.
- Ward, A.F., Duke, K., Gneezy, A., and Bos, M.W. (2017). Brain Drain: The Mere Presence of One's Own Smartphone Reduces Available Cognitive Capacity. *Journal of the Association for Consumer Research* 2 (2), pp. 140–154.
- Weil, S. (1952). *The Need for Roots: Prelude to a Declaration of Duties towards Mankind*. Translated by A. Wills. Abingdon: Routledge.