



Agential Risks: A Comprehensive Introduction

Phil Torres
X-Risks Institute

philosophytorres@gmail.com

Journal of Evolution and Technology - Vol. 26 Issue 2 – August 2016 - pgs 31-47

Abstract

The greatest existential threats to humanity stem from increasingly powerful advanced technologies. Yet the “risk potential” of such tools can only be realized when coupled with a suitable agent who, through error or terror, could use the tool to bring about an existential catastrophe. While the existential risk literature has provided many accounts of how advanced technologies might be misused and abused to cause unprecedented harm, no scholar has yet explored the other half of the agent-tool coupling, namely the *agent*. This paper aims to correct this failure by offering a comprehensive overview of what we could call “agential riskology.” Only by studying the unique properties of different agential risk types can one acquire an accurate picture of the existential danger before us.

1. A new subfield

The field of existential risk studies, or existential riskology, can be traced back to a 1996 book by the philosopher John Leslie.¹ In the early 2000s, the field emerged as a more formal discipline of active scholarship, led primarily by transhumanists (Bostrom 2005). Numerous institutions dedicated to understanding the greatest threats to our collective future have since been founded, such as the Future of Life Institute, the Centre for the Study of Existential Risks (Cambridge), and the Future of Humanity Institute (Oxford). Despite these signs of progress, the field remains in something like a “pre-paradigmatic” stage, whereby a comprehensive research program has yet to be firmly established.

A particularly problematic gap in the scholarship stems from the failure of existential riskologists to take seriously the range of agents who might use advanced technologies to initiate a catastrophe. One finds only occasional references in the literature to “psychopaths,” “hate groups,” “terrorists,” and “malevolent governments,” typically without any further details about the unique properties of these entities. This mistake is tantamount to asserting, “Future technologies – I’ll refrain from saying which ones, how they might be used, the properties that make them dangerous, *and so on* – could annihilate humanity.” Just as it’s crucial to study the properties of advanced technologies, so too is it crucial to study the properties of agents. The present paper aims to rectify this shortcoming: in effect, it establishes a new subfield of *agential riskology*.

The paper is organized as follows: the next section establishes some basic terminology. Sections 3 and 4 examine the phenomena of agential terror and agential error, respectively. The penultimate section then argues that we should expect the threat of ecoterrorism and apocalyptic terrorism to increase nontrivially in the coming decades.

2. Definitions

An “existential risk” is an event that results in either total annihilation *or* a permanent and severe reduction in our quality of life (Bostrom 2002). Let’s refer to the definiens’ first disjunct as an “extinction risk” and the second disjunct as a “stagnation risk.” Extinction risks are terminal for our species, but stagnation risks are survivable, although they entail an irreversible state of significant deprivation, perhaps resulting in the life opportunities of contemporary North Koreans or our ancestors from the Paleolithic. From a transhumanist perspective, both scenarios would prevent us from reaching a posthuman state in which one or more of our “core capacities” are augmented beyond their natural limits (Bostrom 2008). I use the term “existential risk” to reference either scenario, while “extinction risk” and “stagnation risk” refer to specific existential circumstances.

Existential risks are defined by their consequences, not their probability or etiology. With respect to the latter, we can identify three broad categories of existential risk types. First, there are risks posed by nature, such as supervolcanic eruptions, global pandemics, asteroid/comet impacts, supernovae, black hole explosions or mergers, galactic center outbursts, and gamma-ray bursts. These form our *cosmic risk background* and they have no direct, immediate connection to human activity – that is, except insofar as advanced technologies could enable us to neutralize them. For example, we could deflect an incoming asteroid with a spacecraft or devise a vaccine to contain a deadly pathogen that might otherwise cause a global outbreak of infection.

Second, there are anthropogenic risks like climate change and biodiversity loss. These are the accidental byproducts of industrial civilization. As elaborated below, both are slow-motion catastrophes that will almost certainly lower the “conflict thresholds” that ensure peace between state and nonstate actors. They will, in other words, exacerbate existing geopolitical tensions and introduce entirely new struggles. Climate change and biodiversity loss could thus be considered “context risks” whose most significant effects are to modulate the dangers posed by virtually every other existential risk facing humanity – including those from nature.² Other anthropogenic risks include physics disasters (such as the Large Hadron Collider) and accidentally contacting hostile aliens through Active SETI projects.

The third category subsumes risks that arise from the misuse and abuse of advanced “dual-use” technologies. The property of “dual usability” refers to the moral ambiguity of such technologies, which can be used for either good or bad.³ The very same centrifuges that can enrich uranium for nuclear power plants can also enrich uranium for nuclear bombs, and the very same technique (such as CRISPR/Cas9) that might enable scientists to cure diseases could also enable terrorists to synthesize a designer pathogen. According to many existential riskologists, advanced technologies constitute the greatest threat to our collective future. Not only are many of these technologies becoming exponentially more powerful – thereby making it possible to manipulate and rearrange the physical world in unprecedented new ways – but some are becoming increasingly accessible to groups and individuals as well. Consequently, the total number of token agents capable of inflicting harm on society is growing.

While the existential risk literature offers many sophisticated accounts of how such *tools* could be used to cause a catastrophe, almost no one has examined the various *agents* (with one notable exception) who might want to do this and why.⁴ Let’s define a “tool” as any technology that an agent could use to achieve its ends, and an “agent” as any entity, independent of its material substrate, with the capacity to choose its

own actions in the world. This lacuna is problematic because the “risk potential” of advanced technologies can be realized only by a complete “agent-tool coupling.” In other words, a tool without an agent isn’t going to destroy the world. Engineered pandemics require engineers, just as a nuclear missile launch requires a nuclear missile launcher. Thus, it’s crucial to study the various properties special to every type of agent. Without a careful examination of *both sides* of the agent-tool coupling, existential risk scholars could leave humanity vulnerable to otherwise avoidable catastrophes.

To illustrate this point, consider a world X in which a large number of species-annihilating technologies exist, and another world Y in which only a single such technology exists. Now imagine that world X contains a single dominant species of peaceable, compassionate beings who almost never resort to violence. How dangerous is this world? If one looks only at the tools, it appears to be extremely dangerous. But if one considers the agents too, it appears to be extremely safe. Now imagine that world Y contains a species of bellicose, warmongering organisms. Again, if one looks only at the tools, then Y appears far safer than X. But when the complete agent-tool complex comes into view, Y is clearly more likely to self-annihilate.

A final distinction needs to be made before moving on to the next section, namely that between error and terror. Note that this distinction is agential in nature. It concerns the *agential intentions* behind a catastrophe independent of the catastrophes’ consequences. Thus, an error could, no less than terror, bring about an existential disaster. In the case of world X, one might argue that an error is most likely to cause an existential catastrophe, whereas in Y the greatest threat stems from terror. The error/terror distinction is important in part because there appear to be far more token agents who might induce an extinction or stagnation disaster by accident than are likely to bring about such an outcome on purpose. The next two sections discuss agential terror and agential error in turn.

3. Agential terror

Many existential riskologists identify terror involving advanced technologies as the most significant threat to our prosperity and survival. But upon closer examination, there are fewer types of agents who would want to cause an existential catastrophe than one might suspect. Consider another thought experiment: imagine a future world in which there exists a profusion of “doomsday buttons” that are accessible to every citizen of Earth. The question then arises: what sort of individual would intentionally push this button? What kind of agent would purposively cause an existential catastrophe?

If the intention were to actualize an extinction risk, the agent would need to exhibit at least two properties. First, it would need to devalue its own post-catastrophe survival. In other words, the agent would have to be *suicidal*. This immediately disqualifies a large number of entities as potential agential risks, since states and political terrorists tend to value their own survival. Neither North Korea nor al-Qaeda, for example, is suicidal. Their goal, in each case, is to change rather than destroy humanity. Even in the case of suicide bombers and kamikaze pilots, the aim is to ensure group survival through the altruistic sacrifice of one’s own life. And second, the agent would need to want every *other* human on the planet to perish. In other words, he or she would have to be *omnicidal*. (We can coin the term “true omnicide” to refer to circumstances that combine both suicide and omnicide, as just defined, resulting in the irreversible termination of our evolutionary lineage.)

In contrast, if the aim were to actualize a stagnation risk, the agent could be suicidal, omnicidal, or neither, but not both (according to the above definitions). A terrorist could, for example, attempt to permanently cripple modern civilization without harming anyone, including him or herself. Alternatively, a terrorist could attempt to cripple civilization through a suicide attack or an attack directed at others. Either way, the relevant agent would be motivated by an ideology that is incompatible with our species reaching

a posthuman state. In the following discussion, we will consider extinction and stagnation possibilities separately.

A typology of agential risks

With these properties in mind, let's examine five categories of agents that, when coupled with sufficiently destructive tools, might purposively bring about an existential catastrophe.

(1) *Superintelligence*. This is one of the most prominent topics of current existential risk studies, although it's typically conceptualized – on my reading of the literature – as a technological risk rather than an agential risk. To be clear, a variety of agent types could use narrow AI systems as a *tool* to achieve their ends. But once an AI system acquires human-level intelligence or beyond, it becomes an agent in its own right, capable of making its own decisions in pursuance of its own goals.

Many experts argue that superintelligence is the greatest long-term threat to human survival, and I concur. On the one hand, a superintelligence could be malevolent rather than benevolent. Call this the *amity-enmity conundrum*. Roman Yampolskiy (2015) delineates myriad pathways that could lead to human-unfriendly superintelligences. For example, human programmers could intentionally program a superintelligence to prefer enmity over amity. (The relevant individuals could thus be classified as agential risks as well, even though they wouldn't be the proximate agential cause of an existential catastrophe.) A malevolent superintelligence could also arise as a result of a philosophical or technical failure to program it properly (Yudkowsky 2008), or through a process of recursive self-improvement, whereby a “seed AI” augments its capacities by modifying its own code.

But it's crucial to note that a superintelligence need not be malevolent to pose a major existential risk. In fact, it appears more likely that a superintelligence will destroy humanity simply because our species happens to be somewhere between it and its goals. Consider two points: first, the relevant definition of “intelligence” in this context is “the ability to acquire the means necessary to achieve one's ends, *whatever those ends happen to be*.” This definition, which is standard in the cognitive sciences, is roughly synonymous with the philosophical notion of *instrumental rationality*. And since it focuses entirely on an agent's means rather than its ends, it follows that an intelligence could have any number of ends, including ones that we wouldn't recognize as intelligible or moral. Scholars call this the “orthogonality thesis” (Bostrom 2012).

For example, there's nothing incoherent about a superintelligent machine that believes it must purify Earth of humanity because God wills it to do so. Nor is there anything conceptually problematic about a superintelligent machine whose ultimate goal is to manufacture as many paperclips as possible. This goal may sound benign, but upon closer inspection it appears just as potentially catastrophic as an AI that wants us dead. Consider the fact that to create paperclips, the superintelligence would need a source of raw materials: atoms. As it happens, this is precisely what human bodies are made out of. Consequently, the superintelligence could decide to harvest the atoms from our bodies, thereby causing our extinction. As Eliezer Yudkowsky puts it, “The AI does not hate you, nor does it love you, but you are made out of atoms which it can use for something else” (Yudkowsky 2008). Scholars categorize resource acquisition, along with self-preservation, under the term “instrumental convergence.”

Even more, our survival could be at risk in situations that initially appear favorable. For example, imagine a superintelligence that wants to eliminate human sadness from the world. The first action it might take is to exterminate *Homo sapiens*, because human sadness can't exist without humans. Or it might notice that humans smile when happy, so it could try to cover our faces with electrodes that cause certain muscles to contract, thereby yielding a “Botox smile.” Alternatively, it might implant electrodes into the pleasure centers of our brains. The result could be a global population of euphoric zombies too paralyzed by pleas-

ure to live meaningful lives (Bostrom 2014, 146–48). All of these outcomes would, from a certain perspective be undesirable. The point is that there’s a crucial difference between “Do what I say” and “Do what I mean,” and figuring out how to program a superintelligence to behave according to the latter is a formidable task.

Making matters worse, a superintelligence whose material substrate involves the propagation of electrical potentials rather than action potentials would be capable of processing information orders of magnitude faster than humans. Call this a *quantitative superintelligence*. As Yudkowsky observes, if the human brain were sped up a million times, “a subjective year of thinking would be accomplished for every 31 physical seconds in the outside world, and a millennium would fly by in eight-and-a-half hours” (Yudkowsky 2008). A quantitative superintelligence would thus have a huge speed advantage over humanity. In the amount of time that it takes our biological brains to process the thought, “This AI is going to slaughter us,” the AI could already be halfway done the deed.

Another possibility concerns not speed but capacity. That is, an AI with a different cognitive architecture could potentially think thoughts that lie outside of our species-specific “cognitive space.” This is based on the following ideas: (a) to understand a mind-independent feature of reality, one must mentally represent it, and (b) to mentally represent that feature, one must generate a concept whose content consists of that feature. Thus, if the mental machinery supplied to us by nature is unable to generate the relevant concept, the corresponding feature of reality will be unknowable. Just as a chipmunk can’t generate the concepts needed to understand a boson or the stock market, so too are the concept-generating mechanisms of our minds limited by their evolutionary history. The point is that a *qualitative superintelligence* could come to understand phenomena in the universe that are permanently beyond our epistemic reach. This could enable it to devise ways of manipulating the world that would appear to us as pure *magic*. In other words, we might observe changes in the world that we simply can’t understand – that are as mysterious as the science behind cellphones or the atomic bomb is to a chipmunk scientist.

In sum, not only would a quantitative superintelligence’s speed severely disadvantage humanity, but a qualitative superintelligence could also discover methods for “commanding nature,” as it were, that would leave us utterly helpless.

As with the other agents below, superintelligence itself doesn’t pose a direct threat to our species. But it could pose a threat if coupled to any of the tools previously mentioned, including nuclear weapons, biotechnology, synthetic biology, and nanotechnology. As Bostrom writes, if nanofactories don’t yet exist at the time, a superintelligence could build them to produce “nerve gas or target-seeking mosquito-like robots [that] might then burgeon forth simultaneously from every square meter of the globe” (Bostrom 2014). A superintelligence could also potentially gain control of automated processes in biology laboratories to synthesize a designer pathogen, exploit narrow AI systems to disrupt the global economy, and generate false signals in early-warning systems to provoke a nuclear exchange between states. A superintelligence could press a preexisting doomsday button or create its own button.

According to a recent poll, a large majority of AI researchers believe that an artificial general intelligence (AGI) will be created this century (Müller and Bostrom 2016). But such predictions have been notoriously inaccurate in the past. This fact is immaterial to the present thesis, though, which merely states that *if*, independent of *when*, humans build a superintelligence, it could pose a major, unpredictable agential risk.

(2) *Idiosyncratic actors*. This category includes individuals or groups who are driven by idiosyncratic motives to destroy humanity or civilization. History provides several examples of the mindset that would be required for such an act of terror. First, consider Eric Harris and Dylan Klebold, the adolescents behind the 1999 Columbine High School massacre. Their aim was to carry out an attack as spectacular as the Oklahoma City bombing, which occurred four years earlier. They converted propane tanks into bombs, built

99 improvised explosive devices, and equipped themselves with several guns. By the end of the incident, 12 students and one teacher were dead, while 21 others were injured. (Although if the propane bombs had exploded, which they didn't, all 488 students in the cafeteria at the time could have perished.) This was the deadliest school shooting in US history until Adam Lanza killed 20 children and 6 adults at Sandy Hook Elementary School in 2012 before committing suicide.

This leads to the question: what if Harris and Klebold had generalized their misanthropic hatred from their high school peers to the world as a whole? What if certain future anticipated technologies had been available at the time? In other words, what if they'd had access to a doomsday button? Would they have pushed it? The plausible answer is, "Yes, they would have pushed it." If revenge on school bullies was the deeper motive behind their attack, as appears to be the case,⁵ then what better way to show others "who's boss" than to "go out with the ultimate bang"? If people like Harris and Klebold, with their dual proclivities for homicide and suicide, get their hands on advanced technologies in the future, the result could be true omnicide.

History also provides a model of someone who might try to destroy civilization without intentionally killing anyone. Consider the case of Marvin Heemeyer, a Colorado welder who owned a muffler repair shop. After years of a zoning dispute with the local town and several thousand dollars in fines for property violations, Heemeyer decided to take revenge by converting a large bulldozer into a "futuristic tank." It was covered in armor, mounted with video cameras, and equipped with three gun-ports. On June 4, 2004, he climbed inside the tank and headed into town. With a top speed of a slow jog and numerous police walking behind him during the incident, Heemeyer proceeded to destroy one building after another. Neither a flash-bang grenade thrown into the bulldozer's exhaust pipe nor 200 rounds of ammunition succeeded in stopping him. After more than two hours of relentless destruction, the bulldozer became lodged in a basement, at which point Heemeyer picked up a pistol and shot himself.

The motivation of this attack was also a form of bullying, that is, as perceived by Heemeyer. A significant difference between Heemeyer's rampage and the Columbine massacre is that, according to some residents sympathetic with Heemeyer, he went out of his way not to injure anyone. Indeed, he was the only person to die in the attack.⁶ It's also worth pointing out that Heemeyer saw himself as God's servant. As he put it, "God blessed me in advance for the task that I am about to undertake. It is my duty. God has asked me to do this. It's a cross that I am going to carry and I'm carrying it in God's name."

Again, we can ask: what if a delusional person like Heemeyer were to someday hold a grudge not against the local town, but civilization as a whole? What if a future person feels abandoned or "screwed over" by society and wants to retaliate for perceived injustices? In the past, lone wolves with idiosyncratic grievances were unable to wreak havoc on society because of the limited means available to them. This will almost certainly change in the future, as advanced technologies become increasingly powerful and accessible.⁷

This category is especially worrisome moving forward, since it is arguably the type with the most potential tokens. Perhaps future advances in psychology, or brain-decoder technologies and other surveillance systems, will enable us to identify agents at risk of engaging in violence of this kind.

(3) *Ecoterrorists*. Imagine for a moment that scientists found a 52 per cent reduction in the global population of wild vertebrates between 1970 and 2010. Imagine that, on even the most optimistic assumptions, the biosphere had entered the sixth mass extinction event in life's 3.8 billion year history. Imagine further than a single species were almost entirely responsible for this environmental crisis, namely *Periplaneta americana*, the American cockroach. What would our response be? To exterminate the culprit, of course, thereby saving Earth from a planetary catastrophe. It's this basic line of reasoning that could lead a group of radical environmentalists or a lone wolf activist to attempt to relocate humanity from the category of

“extant” to “extinct.” In fact, the claims made above are scientifically accurate: the total population of wild mammals, birds, reptiles, amphibians, and fish really did halve in forty years, and we really are in the beginning stages of a new mass extinction event (see WWF 2014; Ceballos et al. 2015). But the culprit isn’t the American cockroach, it’s humanity.

To date, the vast majority of environmentalist movements have been peaceful, despite the FBI classifying some affiliates of the Earth Liberation Front (ELF) and the Animal Liberation Front (ALF) as constituting “one of the most serious domestic terror threats in the US” (Flannery 2016). Even groups that believe Gaia would be better off without *Homo sapiens* trampling the planet, such as the Voluntary Human Extinction Movement (VHEMT), reject violence or coercion as legitimate means for achieving their ideological aims. Nonetheless, there are exceptions. On September 1, 2010, James Lee terrorized a Discovery Channel building in DC with a gun and explosives. During several hours of negotiations, Lee explained to law enforcement officials that he wanted new programming on the Discovery Channel “to convey how to improve the natural world and reverse human civilization.” He also wanted humans to stop procreating. As Lee wrote, “*Children represent FUTURE catastrophic pollution whereas their parents are current pollution. NO MORE BABIES!*” (quoted in Flannery 2016). Around 4:48pm, Lee aimed his gun at a hostage in the building and SWAT snipers killed him.

Once more, the question is: if Lee had access to a doomsday button that, say, could have sterilized every human on the planet (thereby causing an extinction catastrophe), would he have pushed it? The answer is, “Yes, probably.”

The radical environmental movement is founded on the philosophy of biocentrism, or the view that “humans are no more intrinsically valuable than any other creature.” Some scholars have even suggested that this ideology and its following should be seen as a “new religious movement,” which Bron Taylor calls the “deep green religion” (Taylor 2009). Along these lines, Ted Kaczynski, discussed below, advocated what he called a “wilderness religion.” Given the fanaticism of this stance, it’s not hard to envisage a group emerging in the future that attempts to bring about human extinction through the use of advanced technologies in an effort to “save” the planet. This scenario is made even more plausible by the fact that the largest demographic of Earth Liberation Front members consists of “well educated” and “technologically literate” males (Flannery 2016). Thus, if synthetic biology techniques were to enable the synthesis of a designer pathogen that infects only *Homo sapiens*, a radical environmentalist could try to spread this germ around the globe. Or they could design and release self-replicating nanobots that selectively target *Homo sapiens* by recognizing genetic signatures unique to our DNA. Such nanobots could annihilate our species while leaving the biosphere more or less unharmed.

There might also be radical environmentalists who aren’t motivated by a death wish for humanity, but by a destruction wish for civilization. Ted Kaczynski, also known as the Unabomber, provides a compelling example. His objective was neither suicide nor omnicide, but the dismantlement of technological civilization, according to the slogan “Back to the Pleistocene” (Flannery 2016). Although Kaczynski’s own terrorist bombings weren’t intended to achieve his ambitious aims, if a doomsday button had been available at the time, Kaczynski probably would have pushed it, thereby initiating a stagnation catastrophe. While people might die in the process, this wouldn’t be the goal. As Kaczynski declares, “We therefore advocate a revolution against the industrial system. This revolution may or may not make use of violence; it may be sudden or it may be a relatively gradual process spanning a few decades. ... Its object will be to overthrow not governments but the economic and technological basis of the present society” (quoted in Flannery 2016).

This type of agential risk should be carefully studied moving forward, for reasons explicated in Section 5. If advanced technologies become sufficiently powerful and accessible, people under the spell of the deep green religion could inflict unprecedented harm on civilization.

(4) *Religious terrorists*. Terrorists motivated by nationalist, separatist, anarchist, Marxist, and other political ideologies are unlikely to cause an existential catastrophe because their goals are typically predicated on the continued existence of civilization and our species. They want to change the world, not destroy it. But this is not the case for some terrorists motivated by religious ideologies. For them, what matters isn't this life, but the afterlife; the ultimate goal isn't worldly, but otherworldly. These unique features make religious terrorism especially dangerous, and indeed it has proven to be both more lethal and indiscriminate than past forms of "secular" terrorism.⁸ According to the Global Terrorism Index, religious extremism is now the primary driver behind global terrorism, and there are reasons (see Section 5) for expecting this to remain the case moving forward (Arnett 2014).

The most worrisome form of religious terrorism is *apocalyptic terrorism*. As Jessica Stern and J.M. Berger observe, apocalyptic groups aren't "inhibited by the possibility of offending their political constituents because they see themselves as participating in the ultimate battle." Consequently, they are "the most likely terrorist groups to engage in acts of barbarism" (Stern and Berger 2015). The apocalyptic terrorist sees humanity as being engaged in a cosmic struggle at the very culmination of world history, and the only acceptable outcome is the complete decimation of God's enemies. These convictions, when sincerely held, can produce a grandiose sense of moral urgency that apocalyptic warriors can use to justify virtually any act of cruelty and violence, no matter how catastrophic. To borrow a phrase from the former Director of the CIA, James Woolsey, groups of this sort "don't want a seat at the table, they want to destroy the table and everyone sitting at it" (Lemann 2001).

There are two general types of active apocalyptic groups. First, there are movements that have advocated something along the lines of omnicide. History provides many striking examples of movements that maintained – with the unshakable firmness of faith – that *the world must be destroyed in order to be saved*. For example, the Islamic State of Iraq and Syria believes that its current caliph, or leader, is the eighth of twelve caliphs in total before the apocalypse. This group's adherents anticipate an imminent battle between themselves and the "Roman" forces (the West) in the small northern town of Dabiq, in Syria. After the Romans are brutally defeated, one-third of the victorious Muslim army will supernaturally conquer Constantinople (now Istanbul), after which the Antichrist will appear, Jesus will descend above the Umayyad Mosque in Damascus, and various other eschatological events will occur. In the end, those who reject Islam will be judged by God and cast into hellfire, and the Islamic State sees itself as playing an integral role in getting this process started (Torres 2016a).

Another example comes from the now-defunct Japanese cult Aum Shinrikyo. This group's ideology was a syncretism of Buddhist, Hindu, and Christian beliefs. From Christianity, the group imported the notion of Armageddon, which it believed would constitute a Third World War whose consequences would be "unparalleled in human history." Only those "with great karma" and "those who had the defensive protection of the Aum Shinrikyo organization" would survive (Juergensmeyer 2003). In 1995, Aum Shinrikyo attempted to knock over the first domino of the apocalypse by releasing the chemical sarin in the Tokyo subway, resulting in 12 deaths and sickening "up to 5,000 people." This was the biggest terrorist attack in Japanese history, and it was perpetrated by a religious cult that was explicitly motivated by an active apocalyptic worldview. Other contemporary examples include the Eastern Lightning in modern-day China, which believes that it's in an apocalyptic struggle with the communist government, and the Christian Identity movement in the US, which believes that it must use catastrophic violence to purify the world before the return of Jesus.

Second, there are multiple groups that have advocated mass suicide. The Heaven's Gate cult provides an example. This group is classified as a millenarian UFO religion, led by Marshall Applewhite and Bonnie Nettles. They believed that, as James Lewis puts it, ancient "aliens planted the seeds of current humanity millions of years ago, and have to come to reap the harvest of their work in the form of spiritual evolved

individuals who will join the ranks of flying saucer crews. Only a select few members of humanity will be chosen to advance to this transhuman state” (Lewis 2001). The world was about to be “recycled,” and the only possible “way to evacuate this Earth” was to leave their bodies behind through collective suicide. Members believed that, once dead, they would board an alien spacecraft that was trailing the Hale-Bopp comet as it swung past Earth in 1997. To fulfill this eschatological prediction, they drank phenobarbital, along with applesauce and vodka. Between March 24-26, 39 members of the cult committed suicide.

Other examples could be adduced, such as The Movement for the Restoration of the Ten Commandments of God in Uganda, which slaughtered 778 people after unrest among members following a failed apocalyptic prophesy (*New York Times* 2000). But the point should be sufficiently clear.

With respect to extinction risks, there are (quite intriguingly) no notable groups that have combined these two tendencies of suicide *and* omnicide. No major sect has said, “We must destroy the world, *including ourselves*, to save humanity.” But this doesn’t mean that such a group is unlikely to emerge in the future. The ingredients necessary for a truly omnicidal ideology to take shape are already present in our culture. Perhaps, for reasons discussed below, societal conditions in the future will push religious fanatics to even more extreme forms of apocalypticism, thereby yielding a group that believes God’s will is for *everyone* to perish. Whether this happens or not, apocalyptic groups also pose a significant stagnation risk. For example, what if Aum Shinrikyo had somehow been successful in initiating an Armageddon-like Third World War? What might civilization look like after such a catastrophe? Could it recover? Or, what if the Islamic State managed to expand its caliphate across the entire world? How might this affect humanity’s long-term prospects?

Zooming out from our focus on apocalyptic groups, there are numerous less radical groups that would like to reorganize society in existentially catastrophic ways. One of the ultimate goals of al-Qaeda, for example, is to implement Sharia law around the world. If this were to happen, it would destroy the modern secular values of democracy, freedom of speech and the press, and open scientific inquiry. The imposition of Sharia law on civilization is also the aim of non-jihadist Islamists, who comprise roughly 7 per cent of the Muslim community (Flannery 2014). Similarly, “dominionist” Christians in the US, a demographic that isn’t classified as “terrorist,” believe that God commands Christians to control society and govern it based on biblical law. If a state run by dominionists were to become sufficiently powerful and global in scope, it could induce an existential catastrophe of the stagnation variety.

(5) *Rogue states*. As with political terrorists, states are unlikely to intentionally cause an extinction catastrophe because they are generally not suicidal. Insofar as they pursue violence, it’s typically to defend or expand their territories. The total annihilation of *Homo sapiens* would interfere with these ends. But defending and expanding a state’s territories could cause a catastrophe of the stagnation variety. For example, if North Korea were to morph into a one-world government with absolutist control over the global population until Earth became unlivable, the result would be an existential catastrophe. Alternatively, a benevolent one-world government could emerge from institutions like the United Nations or the European Union. Once in place, a malevolent demagogue could climb to the power ladder and seize control over the system, converting it into a tyrannical dictatorship. Again, the outcome would be a stagnation catastrophe. Of all the agential risk types here discussed, historians, sociologists, philosophers, and other scholars have studied state-level politics and governmental systems the most thoroughly.

4. Agential error

The discussion to this point has focused on agential terror. But what about the other side of the error/terror coin? The danger posed by agential error depends in part on how accessible future technologies become. For example, if even a small percentage of the human population in 2050, which is projected to be 9.3 billion, were to acquire “biohacker” laboratories, the chance that *someone* might accidentally re-

lease a pathogen into the environment could be unacceptably high (Pew Research Center 2015a). After all, a significant number of mistakes have happened in highly regulated government laboratories over the years (Torres 2016a). The 2009 swine flu pandemic may have occurred because of a laboratory mistake made in the 1970s, and “a CDC lab accidentally contaminated a relatively benign flu sample with a dangerous H5N1 bird flu strain that has killed 386 people since 2003” (Zimmer and Burke 2009; McNeil 2014). If such problems occur among professionals, imagine the potential dangers of hobbyists around the world – perhaps hundreds of millions – handling pathogenic microbes with almost no regulatory oversight. The exact same logic applies to other technologies that are becoming more accessible, such as nanotechnology, robotics, AI systems, and possibly nuclear weapons, not to mention future artifacts that currently lie hidden beneath the horizon of our technological imaginations.

There could also be malicious agents that want to cause an existential catastrophe, but nonetheless end up doing this by accident rather than design. For example, in preparing for the “big day,” a doomsday cult could accidentally release a deadly pathogen or self-replicating nanobot into the environment, resulting in an unplanned disaster. This scenario could involve ecoterrorists, idiosyncratic actors, and states as well. Or an agent with no desire to cause an existential catastrophe could push a “catastrophe button” that inadvertently brings about an existential disaster. For example, a rogue state that hopes to gain regional or global power through the use of nuclear missiles must answer the following question: exactly *how many* nuclear missiles are required to bring the world’s governments to their knees without causing an extinction-inducing nuclear winter? The same question could be asked with respect to biological weapons, nanoweapons, and weaponized AI systems. How sure can one be that there won’t be unintended consequences that catapult humanity back into the Stone Age? (As Albert Einstein once said, “I do not know how the Third World War will be fought, but I can tell you what they will use in the Fourth – rocks!” Calaprice 2005.) The unpredictability and uncertainty inherent in global catastrophe scenarios could make it easy for non-existential *terror* to slide into existential *error*.

Finally, a special case of agential error worth examining on its own involves superintelligence. A genuinely superintelligent agent (coupled with advanced technologies) would wield extraordinary power in the world. This fact would make humanity especially vulnerable to any error made by such an agent. Even a *single* mistake could be sufficiently devastating to cause an existential catastrophe. One might respond by asserting that a superintelligence is surely less likely to make a mistake, given its superior intelligence (Bostrom 2014). But I would challenge this assumption. Consider that humans have the most developed neocortex and the highest encephalization quotient in the Animal Kingdom. Yet it is our species, rather than our intellectually “inferior” relatives, that is responsible for the environmental catastrophes of climate change and biodiversity loss. Even more, our species has greatly increased the total number of existential risks from a small handful of improbable natural threats to a dizzying array of anthropogenic and agent-tool dangers. Was this a mistake? In a sense, yes: we certainly didn’t intend for this to happen. Historically speaking, human ingenuity and the threat of existential annihilation have risen together.

This suggests that there isn’t a strong connection between higher intelligence and the capacity to avoid errors. Call this the “orthogonality thesis of fallibility” (Torres 2016a). If our own history is a guide to the future, we might expect the creation of a superintelligence to *further increase* the total number of existential risks, perhaps in ways that are either now, or permanently, inscrutable to us. The point is that *even if* we were to solve the “control problem” and create a friendly superintelligence, it could nudge us over the precipice of disaster on accident, rather than push us on purpose. What can we say? It’s only superhuman.

5. The future of agential risks

Neutralizing the threats posed by agential risks requires understanding not only their synchronic properties, but also how these properties might evolve diachronically. There are two sets of factors relevant to

this task, which we can organize into external and internal categories, depending on whether they originate from outside or within an agent's motivating ideology.

External factors

As previously mentioned, climate change and biodiversity loss are “context risks” that will frame, and therefore modulate, virtually every other threat facing humanity. According to our best current science, these phenomena – appropriately dubbed “threat multipliers” – will become more severe in the coming decades, and their effects will “extend longer than the entire history of human civilization” (Clark et al. 2016). This will significantly elevate the probability of future struggles and conflicts between state and nonstate actors. A simple thought experiment illustrates the point. In which of the following two worlds are wars more likely: one beset by megadroughts, extreme weather, scorching heat waves, desertification, sea-level rise, and the spread of infectious disease, or one without these tragedies? In which of the following two worlds are terrorist attacks more likely: a world in which food supply disruptions, mass migrations, social upheaval, economic collapse, and political instability are widespread, or one in which they're not? One could even ask, in which of the following two worlds is a malevolent superintelligence more likely to emerge: one crushed by environmental catastrophes or one in which civilization is functioning properly?

Environmental degradation could also increase the likelihood of incidents involving idiosyncratic agents, in part because it could increase the prevalence of “bullying”-type behavior. When people are desperate, moral considerations tend to be occluded by the instinctual drive to meet our biological needs. Even more, climate change and biodiversity loss could significantly fuel ecoterrorism. To quote Flannery, “As the environmental situation becomes more dire, eco-terrorism will likely become a more serious threat in the future” (Flannery 2016). Not only will the deleterious effects of industrial society on the natural world become more salient, but sudden changes in the environment's stability could prod activists to consider more aggressive, even violent, tactics. Scientists have, for example, argued that Earth could be approaching an abrupt, irreversible, catastrophic collapse of the global ecosystem. A planetary-scale “state shift” of this sort could unfold on the timescale of decades and cause “substantial losses of ecosystem services required to sustain the human population.” The result would be “widespread social unrest, economic instability, and loss of human life,” and these phenomena could inspire fierce rebellions against civilization.

There are also reasons for expecting climate change and biodiversity loss to nontrivially increase the size and frequency of apocalyptic movements in the future (Torres 2016b; Juergensmeyer, forthcoming). In fact, we already have at least one example of this happening, according to a 2015 study published in the *Proceedings of the National Academy of Sciences*. This study argues that one can draw a straight line of causation from anthropogenic climate change to the record-breaking 2007-2010 Syrian drought to the 2011 Syrian civil war (Kelly et al. 2015). And the Syrian civil war was the Petri dish in which the Islamic State consolidated its forces to become the wealthiest and most powerful terrorist organization in human history. The link between environmental havoc and terrorism has also been confirmed by the current Director of the CIA, John Brennan, the former US Defense Secretary, Chuck Hagel, and the US Department of Defense (Torres 2016c).

As Mark Juergensmeyer (forthcoming) observes in detail, apocalyptic ideologies tend to arise during periods of extreme societal stress. When a group's basic identity and dignity is threatened, when losing one's cultural identity is unthinkable to those in the group, and when the crisis isn't solvable through ordinary human means, people often turn to supernatural frameworks to make sense of their suffering and give them hope for the future. In Juergensmeyer's words, “The presence of any of these three characteristics increases the likelihood that a real-world crisis may be conceived in cosmic terms,” and “cosmic terms” form the language of apocalyptic activism.

Because of climate change and biodiversity loss, these are precisely the conditions we can expect in the future, as societies inch toward the brink of collapse. It's also worth noting that floods, earthquakes, droughts, famines, and disease are prophesied by many religions as harbingers of the end. Consequently, environmental degradation could actually *reinforce* people's prior eschatological convictions, or even lead nonbelievers to convert.⁹ There is, in fact, a strong preexisting base of widespread apocalyptic belief within the Abrahamic traditions. For example, a 2010 Pew poll finds that 41 per cent of Americans believe that Jesus will either "definitely" or "probably" return by 2050 (Pew Research Center 2010), and a 2012 Pew poll reports that 83 per cent of people in Afghanistan, 72 per cent in Iraq, 68 per cent in Turkey, and 67 per cent in Tunisia believe that the Mahdi, Islam's end-of-days messianic figure, will return in their lifetime (Pew Research Center 2012). One should expect these percentages to rise moving forward.

There are two additional reasons for anticipating more apocalyptic movements in the future. First, a statistical point. According to a 2015 Pew study, the percentage of nonbelievers is projected to shrink in the coming decades, despite the ongoing secularization of Europe and North America (Pew Research Center 2015a; see chapter 3, "Articles of Faith"). By 2050, more than 60 per cent of humanity will identify as either Christian or Muslim, in roughly equal proportion. As Alan Cooperman puts it, "You might think of this in shorthand as the secularizing West versus the rapidly growing rest" (Pew Research Center 2015b). This is disconcerting because religion normalizes bad epistemological habits, and thinking clearly about big-picture issues is the only hope our species has of navigating the wilderness of existential risks before us. In addition, not only is superstition rising as advanced technologies become more powerful, but if the relative proportion of extremists at the fringe remains fixed, the *absolute number* of religious fanatics will undergo a growth spurt. This alone suggests that the future will contain a historically anomalous number of terrorists (although we should note that it will contain more "good guys" as well).

Furthermore, the inchoate GNR (genetics, nanotech, and robotics) Revolution will result in a wide range of fundamental changes to society. It could introduce new forms of government – or, as Benjamin Wittes and Gabriella Blum (2015) argue, undercut the social contract upon which modern states are founded – and even challenge our notion of what it means to be *human*. These changes could be profound, pervasive, and quite rapid, given the exponential rate of innovation. If this is the case, it could also fulfill the conditions specified by Juergensmeyer, thereby fueling apocalyptic extremists to declare an imminent end to the world. (In a sense, this might be true, since the transition from the human era to a posthuman era would mark a watershed moment in our evolutionary history.) The fact is that past technological revolutions have inspired religious fanaticism, and by nearly all accounts the GNR Revolution will be far more disruptive than any previous revolution. As Juergensmeyer puts it, "radical change breeds radical religion," and radical change is exactly what we should expect.¹⁰

Tying this all together: a confluence of environmental degradation, demographic shifts, and disruptive technologies could significantly exacerbate the threat of apocalyptic terrorism, as well as idiosyncratic agents and ecoterrorists, in the future. The recent unrest in the Middle East is, arguably, only a preview of what's to come.

Internal factors

But there are also factors *internal* to the ideologies espoused by different agents that are no less important for existential riskologists to study.

For example, the year 2076 will likely see a spike in apocalyptic fervor within the Islamic world (Cook 2008, 2011). One can only know this, and therefore prepare appropriately, if one understands the relevant Islamic traditions. The reason 2076 will be especially dangerous is that it roughly corresponds to 1500 in the Islamic calendar (AH), and eschatological enthusiasm has risen in the past at the turn of the century.

Consider the fact that the Iranian Revolution, which was widely seen as an “apocalyptic occurrence” by Shi’ites, happened in 1979 (Cook 2011). So did the Grand Mosque seizure, during which a group of 500 insurgents took approximately 100,000 worshipers hostage. This group claimed to have the Mahdi with them and believed that the Last Hour was imminent. The point is that 1979 corresponds to 1400AH, a date that fueled the apocalypticism behind these events.

Scholars should also keep an eye on 2039, since it is the 1200th anniversary of the Mahdi’s occultation in the Twelver Shia tradition. As Cook writes, “the 1000-year anniversary of the Mahdi’s occultation was a time of enormous messianic disturbance that ultimately led to the emergence of the Bahai faith. ... [A]nd given the importance of the holy number 12 in Shiism, the twelfth century after the occultation could also become a locus of messianic aspirations.” He adds:

In one scenario, either a messianic claimant could appear or, more likely, one or several movements hoping to “purify” the Muslim world (or the entire world) in preparation for the Mahdi’s imminent revelation could develop. Such movements would likely be quite violent; if they took control of a state, they could conceivably ignite a regional conflict. (Cook 2011)

Looking forward, who knows what powerful technologies might exist by 2039 or, even more, 2076? If a messianic movement with violent proclivities were to arise in the future, it could have existential implications for humanity.

Another example involves apocalyptic US militias influenced by Christian Identity teachings. On April 19, 1995, Timothy McVeigh pulled up to the Alfred P. Murrah Federal Building in Oklahoma City and detonated a bomb that killed 168 people. As Flannery notes, this event unfolded “just as the Christian Identity affiliated Covenant, Arm, and the Sword (CSA) militia had planned a decade earlier while training 1,200 recruits in the Endtime Overcomer Survival Training School.” The date of April 19 “was no accident.” Exactly two years earlier, the government ended its confrontation with the Branch Davidians in their Waco, Texas compound, resulting in 74 deaths. And exactly 8 years before this event, there was a similar standoff between the government and the Covenant, Arm, and the Sword. And centuries before, in 1775, the Battles of Lexington and Concord that inaugurated the American Revolutionary War against Great Britain took place on April 19. Consequently, the date of “April 19 has come to resonate throughout a constructed history of the radical Right as a day of patriotic resistance” (Flannery 2016, 144).

More generally, some experts refer to April as the beginning of “the killing season.” While Harris and Klebold reportedly planned their massacre on April 19 (being inspired by McVeigh), they ended up delaying it one day to coincide with Adolf Hitler’s birthday (Rosenwald 2016). Another date to watch is April 15, the deadline for income tax filings in the United States. This has meaning to certain anti-government groups. As the Anti-Defamation League (2005) warns,

April is a month that looms large in the calendar of many extremists in the United States, from racists and anti-Semites to anti-government groups. Some groups organize events to commemorate these April dates. Moreover, there is always a certain threat that one or more extremists may choose to respond to these anniversaries with some sort of violent act.

It adds: “Because of these anniversaries, law enforcement officers, community leaders and school officials should be vigilant.”

Existential risk scholars too should be especially mindful of heightened risks in April. If a doomsday or catastrophe button were to become available to Christian Identity terrorists motivated by an active apocalyptic ideology, April 19 might be the day on which they would decide to press it.

6. Conclusion

Most states and terrorists are unlikely to intentionally cause an existential catastrophe, although certain high-impact instances of catastrophic violence could accidentally realize an extinction or stagnation risk. The primary danger posed by states and terrorists concerns their capacity to press a catastrophe button, if it were to become available. There are, however, at least five types of agents who could be motivated by various goals to bring about a cataclysm of existential proportions. I do not intend for this list to be exhaustive. Indeed, the agential threat horizon could expand, shift, or transmogrify in unanticipated ways as humanity is thrust forward by the invisible hand of time. It's nonetheless important to specify a typology of agential risks based on the best current research – a task that no one has yet attempted – because the agents of each category have their own unique properties, and must therefore be studied as unique threats in exactly the same way that nuclear weapons, biotechnology, and molecular manufacturing are studied separately.

While much of the existential risk literature focuses on the various tools, both present and anticipated, that could bring about a secular apocalypse, we must give the agents equal consideration. A key idea of this paper is that, first, *advanced technologies will provide malicious agents with bulldozers, rather than shovels, to dig mass graves for their enemies*. And second, the *risk potential* of these technologies cannot be realized without a complete *agent-tool coupling*. This is why the field of existential risk studies desperately needs a subfield of agential riskology, which this paper aims to establish. No doubt much of what I have said above will need to be refined, but such is to be expected when there are few shoulders upon which to stand.

Notes

1. See Leslie 1986.
2. For example, a world thrown into chaos by environmental degradation might be less prepared to deflect an incoming asteroid or coordinate on stopping a global pandemic.
3. Originally, “dual-use” referred to entities with both civilian and military uses, but the term has acquired a more promiscuous signification in recent scholarship.
4. With the exception of superintelligence, discussed below.
5. According to one study, over 66 per cent of premeditated school shootings have been shown to be connected to bullying (Boodman 2006).
6. Although luck might be partly to blame, as Heemeyer fired his rifle at propane tanks that could have killed someone in the vicinity if they had exploded
7. As I've written elsewhere about similar examples: these may seem too anecdotal to be scientifically useful. But drawing this conclusion would be wrong. Given the immense power of anticipated future technologies, single individuals or groups could potentially wield sufficient power to destroy the world. The statistically anomalous cases of omnicidal lone wolves or terrorist groups are precisely the ones we should be worried about, and therefore should study.
8. See, for example, Hoffman 1993.
9. History provides numerous examples of natural disasters leading to a spike in religious belief, such as the Plague of Cyprian.

10. Personal communication. But see Juergensmeyer 2003.

References

- Anti-Defamation League. 2005. Extremists look to April anniversaries. April 6. <http://www.adl.org/combating-hate/domestic-extremism-terrorism/c/extremists-look-to-april.html?referrer=https://www.google.com/-V3qrZpMrJSw> (accessed July 15, 2016).
- Arnett, G. 2014. Religious extremism main cause of terrorism, according to report. *Guardian*. November 19. <https://www.theguardian.com/news/datablog/2014/nov/18/religious-extremism-main-cause-of-terrorism-according-to-report> (accessed July 15, 2016).
- Barnosky, A., E. Hadley, J. Bascompte, E. Berlow, J. Brown, M. Fortelius, W. Getz, J. Harte, A. Hastings, P. Marquet, N. Martinez, A. Mooers, P. Roopnarine, G. Vermeij, J. Williams, R. Gillespie, J. Kitzes, C. Marshall, N. Matzke, D. Mindell, E. Revilla, and A. Smith. 2012. Approaching a state shift in Earth's biosphere. *Nature* 486: 52–58. Available at <http://www.nature.com/nature/journal/v486/n7401/full/nature11018.html> (accessed July 15, 2016).
- Boodman, S.G. 2006. Gifted and tormented. *Washington Post*. May 16. <http://www.washingtonpost.com/wp-dyn/content/article/2006/05/15/AR2006051501103.html> (accessed July 15, 2016).
- Bostrom, N. 2002. Existential risks: Analyzing human extinction scenarios and related hazards. *Journal of Evolution and Technology* 9(1). Available at <http://www.nickbostrom.com/existential/risks.html> (accessed July 15, 2016).
- Bostrom, N. 2005. A history of transhumanist thought. *Journal of Evolution and Technology* 14(1). Available at <http://jetpress.org/volume14/bostrom.html> (accessed July 16, 2016).
- Bostrom, N. 2012. The superintelligent will: Motivation and instrumental rationality in advanced artificial agents. *Minds and Machines* 22(2): 71–85. Available at <http://www.nickbostrom.com/superintelligentwill.pdf> (accessed July 16, 2016).
- Bostrom, N. 2014. *Superintelligence: Paths, dangers, strategies*. Oxford: Oxford University Press.
- Bostrom, N., and M. Čirković. 2008. Introduction. In *Global catastrophic risks*, ed. N. Bostrom and M. Čirković, 1–27. New York: Oxford University Press.
- Calaprice, A. 2005. *The new quotable Einstein*. Princeton, NJ: Princeton University Press.
- Ceballos, G., P. Ehrlich, A. Barnosky, A. Garcia, R. Pringle, and T. Palmer. 2015. Accelerated modern human-induced species losses: Entering the sixth mass extinction. *Science Advances* 1(5) (June 5). <http://advances.sciencemag.org/content/1/5/e1400253> (accessed July 15, 2016).
- Clark, P.U., J.D. Shakun, S.A. Marcott, A.C. Mix, M. Eby, S. Kulp, A. Levermann, G.A. Milne, P.L. Pfister, B.D. Santer, D.P. Schrag, S. Solomon, T.F. Stocker, B.H. Strauss, A.J. Weaver, R. Winkelmann, D. Archer, E. Bard, A. Goldner, K. Lambeck, R.T. Pierrehumbert, and G. Plattner. 2016. Consequences

of twenty-first-century policy for multi-millennial climate and sea-level change. *Nature Climate Change* 6: 360–69.

Available at <http://www.climate.unibe.ch/~stocker/papers/clark16natcc.pdf> (accessed July 16, 2016).

Cook, D. 2008. Iraq as the focus for apocalyptic scenarios. *CTC Sentinel* 1(11) (October).

<https://www.ctc.usma.edu/v2/wp-content/uploads/2010/06/Vol1Iss11-Art8.pdf> (accessed July 15, 2016).

Cook, D. 2011. Messianism in the Shiite Crescent. Hudson Institute. April 8.

<http://www.hudson.org/research/7906-messianism-in-the-shiite-crescent> (accessed July 15, 2016).

Flannery, F. 2016. *Understanding apocalyptic terrorism: Countering the radical mindset*. New York: Routledge.

Hoffman, B. 1993. “Holy terror”: The implications of terrorism motivated by a religious imperative. RAND Corporation.

<https://www.rand.org/content/dam/rand/pubs/papers/2007/P7834.pdf> (accessed July 15, 2016).

Juergensmeyer, M. 2003. *Terror in the mind of God*. Los Angeles: University of California Press.

Juergensmeyer, M. Forthcoming. Radical religion in response to catastrophe. *Exploring emerging global thresholds: Toward 2030*. New Delhi: OrientBlackswan.

Kelly, C.P., S. Mohtadi, M.A. Cane, R. Seager, and Y. Kushnir. 2015. Climate change in the Fertile Crescent and implications of the recent Syrian drought. *Proceedings of the National Academy of Sciences* 112(11): 3241–3246.

Available at <http://www.pnas.org/content/112/11/3241> (accessed July 15, 2016).

Landes, R. 2011. *Heaven on Earth: The varieties of the millennial experience*. Oxford: Oxford University Press.

Lemann, N. 2001. What terrorists want. *New Yorker*. October 29.

<http://www.newyorker.com/magazine/2001/10/29/what-terrorists-want> (accessed July 28, 2016).

Leslie, J. 1996. *The end of the world*. London: Routledge.

Lewis, J. 2001. *Odd gods: New religions and the cult controversy*. Amherst, NY: Prometheus.

WWF. 2014. Living planet report. WWF Global website.

<http://bit.ly/1sxx5m> (accessed July 15, 2016).

McNeil, D., Jr. 2014. C.D.C. closes anthrax and flu labs after accidents. *New York Times*. July 11.

http://www.nytimes.com/2014/07/12/science/cdc-closes-anthrax-and-flu-labs-after-accidents.html?_r=0.

Müller, V.C., and N. Bostrom. 2016. Future progress in artificial intelligence: A survey of expert opinion. In *Fundamental Issues of Artificial Intelligence*, ed. V.C. Müller, 553–70. Berlin: Springer.

An online version available at <http://www.nickbostrom.com/papers/survey.pdf> (accessed July 15, 2016).

New York Times. 2000. Cult in Uganda poisoned many, police say. July 28.

<http://www.nytimes.com/2000/07/28/world/cult-in-uganda-poisoned-many-police-say.html> (accessed July 15, 2016).

- Pew Research Center. 2010. Jesus Christ's return to Earth. July 14.
<http://www.people-press.org/2010/06/22/public-sees-a-future-full-of-promise-and-peril/>
 (accessed July 16, 2016).
- Pew Research Center. 2012. The world's Muslims: Unity and diversity. August 9.
<http://www.pewforum.org/2012/08/09/the-worlds-muslims-unity-and-diversity-executive-summary/> (ac-
 cessed July 16, 2016).
- Pew Research Center. 2015a. The future of world religions: Population growth projections, 2010-2050.
 April 2.
 Available at http://www.pewforum.org/files/2015/03/PF_15.04.02_ProjectionsFullReport.pdf (accessed
 July 16, 2016).
- Pew Research Center. 2015b. Event: The future of world religions. April 23.
<http://www.pewforum.org/2015/04/23/live-event-the-future-of-world-religions/> (accessed July 26, 2016).
- Rosenwald, M. 2015. The strange seasonality of violence: Why April is “the beginning of the killing sea-
 son.” *Washington Post*. April 4.
[https://www.washingtonpost.com/local/the-strange-seasonality-of-violence-why-april-is-the-beginning-
 of-the-killing-season/2016/04/03/4e05d092-f6c0-11e5-9804-537defcc3cf6_story.html](https://www.washingtonpost.com/local/the-strange-seasonality-of-violence-why-april-is-the-beginning-of-the-killing-season/2016/04/03/4e05d092-f6c0-11e5-9804-537defcc3cf6_story.html) (accessed July 16,
 2016).
- Stern, J., and J.M. Berger. 2015. *ISIS: The state of terror*. New York: HarperCollins.
- Taylor, B. 2009. *Dark green religion: Nature spirituality and the planetary future*. Los Angeles: Univer-
 sity of California Press.
- Torres, P. 2016a. *The end: What science and religion tell us about the apocalypse*. Charlottesville, VA:
 Pitchstone.
- Torres, P. 2016b. Apocalypse soon? How emerging technologies, population growth, and global warming
 will fuel apocalyptic terrorism in the future. *Skeptic* 21(2): 56–62 .
 Available at <http://goo.gl/Xh9JqO> (accessed July 15, 2016).
- Torres, P. 2016c. We're speeding toward a climate change catastrophe – and that makes 2016 the most
 important election in a generation. *Salon*. April 10.
[http://www.salon.com/2016/04/10/were_speeding_toward_a_climate_change_catastrophe_and_that_mak
 es_2016_the_most_important_election_in_a_generation/](http://www.salon.com/2016/04/10/were_speeding_toward_a_climate_change_catastrophe_and_that_makes_2016_the_most_important_election_in_a_generation/).
- Wittes, B., and G. Blum. 2015. *The future of violence*. New York: Basic Books.
- Yudkowsky, E. 2008. Artificial Intelligence as a positive and negative factor in global risk. In *Global
 catastrophic risks*, ed. N. Bostrom and M. Ćirković, 308–35. New York: Oxford University Press. Avail-
 able at <https://intelligence.org/files/AIPosNegFactor.pdf> (accessed July 15, 2016).
- Zimmer, S.M., and D.S. Burke. 2009. Historical perspective – Emergence of influenza A (H1N1) viruses.
New England Journal of Medicine 361: 279–85.
 Available at <http://www.nejm.org/doi/full/10.1056/NEJMra0904322> (accessed July 15, 2016).