

Evaluating Future Nanotechnology: The Net Societal Impacts of Atomically Precise Manufacturing

Steven Umbrello, Seth D. Baum
Global Catastrophic Risk Institute, <http://gcrinstitute.org>

Futures, forthcoming, doi: 10.1016/j.futures.2018.04.007. This version 28 April 2018.

Abstract

Atomically precise manufacturing (APM) is the assembly of materials with atomic precision. APM does not currently exist, and may not be feasible, but if it is feasible, then the societal impacts could be dramatic. This paper assesses the net societal impacts of APM across the full range of important APM sectors: general material wealth, environmental issues, military affairs, surveillance, artificial intelligence, and space travel. Positive effects were found for material wealth, the environment, military affairs (specifically nuclear disarmament), and space travel. Negative effects were found for military affairs (specifically rogue actor violence and AI). The net effect for surveillance was ambiguous. The effects for the environment, military affairs, and AI appear to be the largest, with the environment perhaps being the largest of these, suggesting that APM would be net beneficial to society. However, these factors are not well quantified and no definitive conclusion can be made. One conclusion that can be reached is that if APM R&D is pursued, it should go hand-in-hand with effective governance strategies to increase the benefits and reduce the harms.

Keywords: atomically precise manufacturing; nanotechnology; ethics; research policy

1. Introduction

Atomically precise manufacturing (APM) is the assembly of materials with atomic precision. APM is also known as molecular assembly or molecular manufacturing, and is a form of molecular nanotechnology. APM does not currently exist, and some nanotechnology researchers doubt that it is feasible (e.g., Smalley, 2001), but those who believe it is feasible anticipate a wide range of transformative societal impacts (e.g., Drexler, 2013a). The best-case scenarios are incredible, featuring benefits on par with the industrial and computer revolutions. Worst-case scenarios are catastrophic, with harms up to and including human extinction.

The uncertain feasibility of APM and the wide range of potential impacts make for a difficult governance challenge. It is hard to know whether APM research and development (R&D) should be encouraged, discouraged, or simply ignored. Thus Marchant et al. (2008) argue that traditional risk management approaches cannot be applied to APM. For example, the precautionary principle could suggest that APM R&D should be discouraged until it is clear that APM could not cause catastrophe, or it could suggest that APM R&D should be encouraged on grounds that it could help prevent other catastrophes. To the extent that prudent risk management demands a full accounting of the risks, and, more generally, that prudent governance demands a full accounting of the relevant factors and issues, it is not readily clear what society should do about APM.

In order to bring some clarity to the issue, this paper characterizes the potential societal impacts of APM if it is developed. The paper examines impacts across the wide range of sectors that APM could affect, including the environment, the military, surveillance, computing, space travel, and general material wealth. It seeks to understand what the net impacts on society would be, in particular whether they would be positive or negative, and by how much. The impacts are not quantified with any significant precision, as would be unwarranted for such a complex and uncertain future technology. However, a broad and imprecise analysis can nonetheless shed some light on the impacts of APM.

The paper is motivated by the question of whether society should invest in APM R&D. Greater APM R&D investment could bring APM into existence sooner (if APM would have been developed anyway) or with higher probability (if APM might not have otherwise been developed). Whether society should invest in APM R&D depends in part on what the net societal impacts of APM would be. If the net impacts would be negative, then arguably there should be no R&D investment. If net impacts would be positive, then there may be a case for R&D investment, though it would also depend on other factors, including the size of the net impacts, the cost of APM R&D, and the attractiveness of alternative investment opportunities. Thus, while this paper's analysis can inform APM R&D investment decisions, it cannot give definitive answers to how such decisions should be made.

This interest in APM R&D decisions is a reason for the paper's breadth across the range of APM impacts. One might object that this breadth makes for shallow and thus lackluster research. However, decisions like whether to invest in APM R&D depend on the full breadth of impacts. Thus, there is a need for synthesis across the potential impacts. Such synthesis additionally can identify important points of uncertainty where more narrow research would be of greatest value. This paper considers prior literature on APM impacts (e.g., Altmann, 2006; Drexler, 2013a; Freitas, 2006, 2007; Hughes, 2007; Timmermans et al., 2011) and extends this literature with original analysis on each major sector of APM impacts in order to make progress on the question of net societal impacts and key areas of uncertainty. The prior literature on APM was mostly written over 10 years ago, Drexler (2013a) being a notable exception. Much has changed since then, both in the world at large and the relevant literature available from other fields. This paper therefore offers important updates from the prior literature, in addition to its original contribution from aggregating a broad spectrum analysis across multiple sectors.

More fundamentally, one might object that APM is too speculative and not imminent enough to merit attention. Indeed, while the prospect of APM was central to the establishment of nanotechnology as a research topic, the nanotechnology scientific community has more recently downplayed APM in favor of more modest and immediately achievable forms of nanotechnology (Selin 2007). Additionally, some ethicists argue that APM does not merit ethical analysis because it displaces attention from analysis of near-term nanotechnology and gives credibility to speculative future scenarios that may not come to be (e.g., Nordmann 2007; Nordmann and Rip 2009; Grunwald, 2010; Jones et al., 2011).

One response to these concerns is that APM is perhaps not so speculative after all. As proof of principle, APM already exists in nature in the form of biomolecules such as ribosomes, which "manufacture" protein molecules as per the "instructions" they read from RNA molecules

(Drexler, 1986; Jones, 2004; Freitas and Merkle, 2004). Additionally, rudimentary forms of artificial APM are now possible via scanning tunneling microscopy and atomic force microscopy, which are both capable of accurately analyzing and determining the atomic structure of samples. In fact, scanning tunneling microscopes can even move atoms to precise locations, building elemental atomic structures, an important step towards APM (e.g., Farrel and Levinson, 1985; Gomer, 1986; Møller et al., 2017). Progress towards APM can also be seen in the 2016 Nobel Prize in Chemistry, which was awarded to Jean-Pierre Sauvage, Sir J. Fraser Stoddart and Bernard L. Feringa for their work on controllable molecules. While these advances fall short of the advanced APM discussed in this paper, they do suggest its feasibility. It still remains possible that APM may never be built, but the possibility should not be dismissed.

Roache (2008) provides two additional reasons for early attention to speculative future technologies like APM. First, analysis of the merits of R&D can be of particular value while the science and technology remain immature and thus more governable. Once a field of science and technology grows large, it gains inertia in the form of institutional buildup and scientists' and technologists' professional investment. This makes it harder to restrict, even if analysis finds it to be dangerous. To be sure, early analysis is also more ambiguous, since important details of the science and technology have not yet been established. This is the well-known Collingridge dilemma between information about a technology and the power to control it (Collingridge 1980). This underscores the importance of early analysis of the potential impacts of a technology: to the extent that information can be obtained early, it is especially valuable.

Second, Roache (2008) argues that seemingly unlikely future scenarios can still be worth attention if their potential consequences are large enough. Similarly, Ćirković (2012) argues that improbable theories can merit attention when they would have high stakes if they turn out to be true. As noted above and discussed in detail throughout this paper, the potential consequences of APM could be quite large indeed. The combination of APM being potentially feasible, at an early stage of R&D, and potentially highly impactful render it a worthy focus of research, especially for research that is, such as this paper, oriented towards assessing its overall merits.

In order to determine the net social impacts of APM, this paper aggregates the net impacts that APM has to various spheres. First, the value that APM has to general society, particularly the social impacts of APM generated abundance. Secondly, how APM affects the environment, most notably how APM can be used in carbon dioxide removal and the fabrication of photovoltaics. Thirdly, the impacts of APM on the military and the manufacturing of new weapons systems will be analyzed. Fourthly, the value of APM to other surveillance technologies like cameras and mass surveillance networks will be discussed. Finally, the impact that APM has on the development of artificial intelligence (AI) as well as space travel will also be looked at in order to determine whether or not APM merits further investment and R&D.

2. Evaluative Framework

This paper's framework is essentially consequentialist, aiming to assess the various consequences of APM in terms of benefits or harms to society. This is not to say that consequentialism is necessarily the only correct framework for evaluating APM; we only claim that the consequences of APM are an important factor. The paper takes society to be the entire

global human population that currently exists or will exist in the future.¹ Furthermore, while the paper does not employ specific conceptions of “benefits” and “harms”, we believe the analysis is broadly consistent with common conceptions of welfare and related evaluative criteria (e.g., Ng 1980; Broome 2004). Likewise, the paper’s use of “value” and related terms can be interpreted accordingly.

APM is expected to have impacts across a wide range of sectors. Several sectors stand out as especially important. We propose a framework to represent and assess the potential value of APM R&D across these sectors based on the following equation:

$$V = G + \varepsilon + M + \Sigma + A + S + X \quad (1)$$

In Equation 1, V is the total value of APM; G is the value of APM for general material wealth; ε is the value of APM for environmental issues; M is the value of APM for military affairs; Σ is the value of APM for other surveillance;² A is the value of APM to artificial intelligence; S is the value of APM for space travel and colonization; and X is a catch-all category for other factors, including factors that are currently unknown. These sectors were selected through prior literature on and our original analysis of APM impacts. It is possible that other sectors may also be important, in which case X may be large. Furthermore, Equation 1 does not account for potential interconnections between the sectors, and thus may not be a complete accounting of APM impacts. Instead, it is intended as a good starting point for the aggregate evaluation of APM.

The following discussion does not attempt to quantify any of the variables in Equation 1. Each of them is too complex and ambiguous to quantify. Instead, the analysis aims to identify the most important factors for each variable, assess whether the effects are likely positive or negative, and start to formulate some understanding of what the net effect (V) might be.

3. Characterizing the Value of APM (V)

APM could affect several sectors of society, as modeled in Equation 1. Indeed, it may be difficult to restrict APM to one particular sector if that was desired. APM is a “generic” technology in the sense that will likely develop concurrently across multiple sectors, as has happened with other technologies such as computers and 3D printing (Drexler, 2013a; Koepsell, 2009).

3.1 The Value of APM for General Material Wealth (G)

Much of the proposed value of APM lies in its potential to create extensive, general-purpose material wealth, which Drexler (2013a) refers to as “radical abundance”. Drexler’s vision is utopian.³ Food shortages become a thing of the past. A wide range of consumer items, including complex technologies, become so efficient and inexpensive to produce that they could be widely distributed with ease. Medicine would be readily available to whoever needed it, thus eliminating

¹ As with any analysis of impacts across broad populations, the analysis in this paper may be sensitive to the particular perspectives that we, the analysts, bring to bear. We have nonetheless attempted to focus on impacts that are widely salient.

² Σ is for “other” surveillance because APM surveillance can also be used for military affairs. In order to avoid double-counting, the military applications of APM surveillance are classified within M .

³ Drexler’s work does also acknowledge some potential downsides to APM, but these are portrayed as small relative to potential benefits.

any treatable disease. Novel nanoscale devices would eliminate all disease and extend the healthy lifespan (Timmermans et al. 2011; Freitas, 1999, 2007, 2010). Fundamental conceptions of economics would have to change as supply and demand under scarcity is replaced by an effectively infinite supply at negligible production costs.

However, there is reason to believe that an increase in material wealth could bring little or even zero increase in human quality of life. Economic theory makes heavy use of the concept of diminishing marginal utility of consumption—essentially, a dollar is worth more to a poor person than to a rich person. Already, much of the world’s population is quite rich by historical standards. Meanwhile, global poverty is slowly but steadily diminishing (Chandy and Gertz, 2011). The fastest growing economies are in poorer countries, particularly Africa, and they are continually catching up with the rest of the world (Robinson, 2015; Worstall, 2015), though gains have been diminished by a growing concentration of wealth among the world’s wealthiest people (Jones, 2015). Barring a major economic shock, such as from climate change, or a radical concentration of wealth, the world population when APM comes into existence would be even wealthier than it is today. APM would thus be making the rich even richer, which, per the theory of diminishing marginal utility of consumption, would be of limited value.

Furthermore, some psychological research on quality of life suggests that the concept of diminishing marginal utility of consumption could actually overstate the value of material wealth. Some analysis of survey data about human happiness finds no diminishing marginal utility of consumption in at least some contexts (e.g., Easterlin et al., 2010), though this finding is disputed (e.g., Stevenson and Wolfers, 2008). Meanwhile, Lyubomirsky (2008) proposes that material conditions make up at most around 10% of personal happiness, with genetics being the largest factor and activities such as interpersonal relationships also being important. Similarly, Seligman (2012) proposes that the main factors in human flourishing are positive emotions like pleasure and gratitude, cognitive engagement in activities, interpersonal relationships, meaning from serving a greater cause, and a sense of accomplishment from completing projects or goals.

Taken in this context, an APM-induced radical abundance of material wealth would not necessarily improve quality of life, and could even worsen it. Would people find fewer activities to pursue or less meaning in their lives if material wealth is readily abundant? How would interpersonal relations be affected? Would the avoidance of disease be a blessing, or would it be a curse in disguise by depriving people of certain types of emotional challenges to overcome? Answers to questions like these could determine whether APM material wealth would be a good thing or not.

An additional factor is how the wealth from APM would be managed and distributed. All technologies evolve within the socio-cultural context that they are developed or introduced in. The introduction of new technologies often reinforces existing social norms and structures (Sandler, 2007). If APM and its products are restricted to a particular class or social order (e.g., as result of licensing costs and imposed restrictions), then the benefits of APM could exacerbate existing inequalities. Of particular concern is that APM could impoverish people who no longer have a role in the APM economy. Given the wide range of potential applications of APM across the economy, this could make for an especially acute case of the more general phenomenon of technological unemployment. Similarly, in places with weak governance structures, APM wealth

could fuel violence between competing factions, and in places with authoritarian governance, APM could fuel oppression, such that the population would have been better off without APM.

Taking these factors together, it appears that the value of APM for general material wealth depends heavily on how APM is governed as well as on other societal trends. This suggests that, to the extent that APM is pursued, it should be done in concert with efforts to improve human quality of life more generally. If this can be done successfully—and we believe that it can—then the value of APM for general material wealth is likely to be positive and may even be large. Given existing trends in economic growth and related factors, our expectation is for modest benefits of APM for general material wealth.

3.2 The Value of APM for Environmental Issues (E)

Despite decades of effort, society continues to struggle with a range of major environmental issues. Some observers argue that solving the environmental issues requires breakthroughs in technological innovation (e.g., Karlsson, 2014). APM advocates posit that APM could be precisely such a breakthrough environmental technology, reducing environmental pollution in several ways (e.g., Drexler, 1986; 2013a).

Perhaps the simplest way that APM could reduce pollution is through more efficient manufacturing. The current state of large-scale manufacturing poses a serious and growing environmental threat due to the large and growing global population, per-capita demand for material products, and the environmental impacts of existing manufacturing processes (UNEP, 2010). Barring significant changes in consumption patterns, this threat will continue unless and until new methods of production are introduced. APM could offer a vastly cleaner manufacturing process because of its potential to precisely determine its outputs, thereby substantially reducing humanity's ecological footprint (Drexler, 2013a).

Another environmental application of APM is photovoltaics. Photovoltaics produce electricity with little-to-zero greenhouse gas emissions. In recent years, the cost of photovoltaics has steadily decreased, but for most of the world, it is still not competitive with fossil fuels, especially for base load electricity production. It has been proposed that APM could dramatically lower the cost of large-scale photovoltaics production and enable a switch from rare materials (e.g., gallium, indium) to abundant materials (e.g., iron, sulfur) while maintaining high energy conversion efficiency (Battelle and Foresight, 2007; Drexler, 2013a, 2013b).⁴ This could greatly facilitate the mass phaseout of fossil fuels, thereby greatly reducing or potentially even eliminating humanity's ongoing carbon footprint, bringing large reductions to climate change risk. Fossil fuel phaseout would additionally be of significant value to public health by reducing the air pollution caused by burning fossil fuels.

A mass switch to photovoltaics would not solve the problem of the extensive greenhouse gas emissions that have already accumulated in the atmosphere. However, APM may be able to help with that via the removal of carbon dioxide from the atmosphere (CDR). With current technologies, CDR is feasible but generally expensive, and some CDR techniques also pose ecological risks (Tavoni and Socolow, 2013). A safe and promising CDR technique is direct air

⁴ Energy conversion efficiency refers to the portion of incoming solar energy that is converted into usable electricity.

capture, in which manufactured devices absorb or adsorb CO₂ from the air.⁵ Direct air capture is currently prohibitively expensive for large-scale CDR. The manufacturing requirement would be difficult but feasible, comparable to or somewhat less than the global automobile manufacturing industry; the greater challenge would be in the energy consumed by the operation of direct air capture units, which would amount to about 1% of current global energy consumption (Lackner et al., 2012). APM could potentially dramatically lower the cost of direct air capture by streamlining the manufacturing of CDR units and especially by lowering the cost of energy via inexpensive photovoltaics. In doing so, APM might effectively solve the climate change problem.

It is difficult to overstate the importance of APM potentially addressing climate change. Climate change is an extreme global risk. Already, it is disrupting ecosystems, forcing human migration, and worsening meteorological disasters, among other things. These impacts are projected to get significantly worse over upcoming decades and centuries. Indeed, they could be large enough to offset any concurrent gains from economic and technological development. Worst-case scenarios see large portions of Earth becoming uninhabitable (Sherwood and Huber, 2010). Meanwhile, the struggle to reduce greenhouse gas emissions has sparked interest in technological “solutions” that could be even riskier (Baum et al., 2013). If APM could prevent these harms from happening, it would go a long way towards putting Earth’s and humanity’s futures on more secure footing.

APM may have less impact on some other major environmental issues. For example, the loss of biodiversity is driven largely by habitat destruction and pollution from the livestock industry. As long as demand for livestock remains robust, APM would have comparatively little impact on biodiversity loss relative to climate change. Drexler (2013a) proposes using APM to enable agriculture in enclosed spaces (e.g., greenhouses), which could bring some increased efficiencies. Further opportunities may exist, such as in the production of “cultured meat”, which is meat produced from cell cultures instead of from procreated animals (Tuomisto and Teixeira de Mattos, 2011). APM may be able to help construct cultured meat facilities and enable the process to be conducted more precisely or at larger scales. These sorts of APM applications could lessen biodiversity loss and the other harms from livestock pollution. The effect may not be as large as that for climate change, but it would still be a significant benefit.

APM could pose environmental harms. These would generally be distinct from the harms from non-APM nanomaterials, which harm human and ecological health due to their small size and chemically reactive nature (UK, 2004). In contrast, APM would likely be used predominantly for manufacturing of larger objects from more traditional materials. The potential for environmental harms comes from APM enabling more manufacturing in general. While APM would in most contexts permit cleaner manufacturing and general environmental benefits, it could also be used in environmentally harmful ways. APM could make it easier to engage in deforestation, strip mining, deep sea drilling, factory farming, or a range of other environmentally harmful activities. At the same time, APM could make it easier to pursue cleaner alternatives to these activities. For the most part, the environmental impacts of APM would come down to how APM is used. And for the most part, we see humanity as striving to reduce environmental harms to the extent that

⁵ In absorption, the entire CO₂ molecule is taken up by (i.e., is absorbed by) another substance; in adsorption, the surface of the CO₂ molecule sticks to (i.e., adheres to) another substance.

doing so is reasonably affordable and easy, so we expect the environmental harms from APM would be relatively small.

In summary, APM shows strong potential for very large environmental benefits and little potential for environmental harms. Given the extreme severity of environmental risks, especially climate change, the environmental benefits of APM would appear to be significantly larger than the benefits of general material wealth. Regardless, both effects seem to point in the same direction, suggesting some merit for APM R&D.

3.3 The Value of APM to Military Spheres (M)

Whereas the value of APM to general material wealth and the environment is clearly positive, it is possible that the value to the military sphere (more precisely, the value to society due to changes to the military sphere) could be negative. This is due to the fact that APM could enhance weapons production and disrupt geopolitical stability (Grubrud, 2003). Indeed, some have emphasized the military sphere as the most important effect of APM, suggesting that this effect could make APM a major catastrophic risk (Häggström 2016; Sandberg, 2014). However, a closer look at the military effects of APM shows that this is not necessarily the case.

One significant concern is that APM could enable powerful new types of weaponry. For the military sector, APM could be used to create higher precision and higher yield weapons systems (Heller and Peterson, 2006). Developing novel weapons systems via APM techniques might increase weapon destructiveness while reducing manufacturing costs. States and other actors would be able to build large stockpiles of highly effective weapons at low cost. This could in turn affect the actors' propensity to engage in violence. Since targets can be neutralized with greater success rates and at minimal costs relative to conventional weapons, they may be more likely to engage in violence.

A related concern is that APM could initiate arms races to develop more powerful weapons (Drexler, 2013a, 2013b; Grubud, 2003). As APM lowers the cost and increases the destructiveness of weapons, more states and nonstate actors may be inclined to build large arsenals. When one actor sees an adversary expanding its arsenal, that actor may seek to respond in kind in order to avoid being outgunned. If the arms races proceed until one side can no longer afford to keep up, the net result could be extremely risky because APM would enable many actors to build arsenals capable of unleashing massive global destruction. The world would become a very perilous place.

However, these concerns may be overstated. First, arms races are not an inevitability. Out of the 200 or so nations, only 10 have ever built nuclear weapons, and only two of those waged an open-ended arms race. Comparable arms races have not occurred with other weapons classes such as biological and chemical weapons. While nations do continuously enhance their military technologies so as to seek or maintain strategic advantage, these upgrades generally do not cause war. Therefore, there may never be an APM arms race, and if there is one, it might not lead to war (Drexler and Pamlin, 2013).

One must also look at the types of weapons that APM would be used for. As a general-purpose manufacturing enhancer, APM could be used to produce a wide range of weapons. Therefore, to

see what types of APM weapons to expect, one should look at what types of weapons states and other actors generally want.

Among most states, the trend has been towards low-yield, high-precision weapons. This is the exact opposite of nuclear weapons, which for this reason have been labeled “dinosaurs”: large, sluggish, and doomed for extinction (Wilson, 2013). Precision weapons enable states to achieve their military and security goals with less death and destruction (“collateral damage”). It is thus likely that APM would be used mainly for precision weapons, not for high-yield weapons of mass destruction. This could enable states to achieve “winter-safe deterrence”: military arsenals that achieve deterrence and other security goals without risking nuclear winter or other global catastrophe (Baum, 2015).

APM could be a significant boon for nuclear disarmament. Already, the United States is diminishing the role of nuclear weapons as a military option while increasing the role of conventional precision weapons (US, 2010). Other states, including Russia, appear interested in doing the same, but conventional precision weapons are currently expensive. By lowering the cost of conventional precision weapons, APM could facilitate bilateral Russia-US nuclear disarmament and eventually also disarmament among other nuclear-armed states. To be sure, the availability of conventional precision weapons is very much not the only factor in nuclear disarmament (Deep Cuts Commission, 2010), but it could play a significant positive role.

A potential downside to the wider availability of conventional precision weapons is that these weapons are more likely to actually be used. Indeed, the massive destructive force of nuclear weapons is widely (but not universally) believed to have helped the world avoid major war since World War II. It could be argued that an arsenal of APM weapons with a low threshold of use could undermine geopolitical stability. However, this is an issue with all types of precision weapons, not just APM weapons. If anything, APM would make it easier for states to maintain successful deterrence—and thereby avoid war—without risking massive destruction if deterrence fails.

The benefits of APM for nuclear disarmament could be quite large, comparable to the benefits for climate change. Worst-case scenarios for nuclear war include nuclear winter, an environmental effect in which smoke from nuclear detonations enters the stratosphere and blocks sunlight, thereby cooling the surface (Mills et al., 2014). Nuclear winter could be catastrophic for agriculture; even a “small” nuclear war with a fraction of the current global arsenal could cause widespread global starvation (Xia et al., 2015). Secondary effects such as disease outbreaks could exacerbate the consequences. The ongoing probability of nuclear war can be difficult to quantify but prior analyses show that it is definitely not zero and could be quite large (Hellman, 2008); Barrett et al., 2013). Meanwhile, progress on nuclear disarmament has been slow, in part because of the difficulty of procuring alternative weapons. If APM could change this, thereby reducing or eliminating the risk of nuclear war, then human civilization may finally be free of the threat of mutual assured destruction that it has lived with for decades.

A more significant concern for military APM comes from the potential dangers of rogue actors, including rogue states such as DPRK as well as terrorist groups and other nonstate actors. Over the last two decades, rogue actors have been an increasingly prominent concern for the

international community. Looking ahead, some worry that advances in certain technologies, especially biotechnology, could enable rogue actors to cause outsized harm, potentially even a major global catastrophe (e.g., Rees, 2003). APM could also enable a wider range of rogue actors to create powerful arsenals. APM could further make these arsenals smaller and thus easier to conceal. In this regard, APM could be considered similar to biotechnology. This makes for a major risk: a world in which small rogue groups can cause global harm is a fragile world to live in.

However, APM could also be used defensively by the rest of the world to counter threats from rogue actors. While APM can make rogue actors' arsenals smaller, it can also enable the rest of the world to detect smaller arsenals. Already, non-APM nanotechnology is being used for a variety of military sensors (Altmann, 2006); APM may be able to significantly enhance sensor technology and other forms of surveillance, making it harder for rogue actors to go unnoticed. (For further discussion of APM surveillance, see Section 3.4.) Additionally, APM could have applications in countermeasures, such as vaccine development for mitigating the effects of bioterrorism. Such possibilities could limit the harms from APM-enhanced rogue actors.

The ultimate effect of military APM will depend on how it is governed. Military APM would likely fall under international governance regimes aimed at limiting the harms of violent conflict. Some APM weapons could be governed under existing treaties for nuclear, chemical and biological weapons (Reynolds, 2001; Altmann, 2006). New treaties could also be drafted for other APM weapons (Wilson, 2013; Hughes, 2007). Current discussions for controlling lethal autonomous weapons (LAWS) serves as a case study that could be emulated and expedited once ubiquitous nanotechnology is available. In 2014, the United Nations Convention on Conventional Weapons (CCW) and the Campaign to Ban Killer Robots raised the issue of the regulation of autonomous weapons systems and initiated negotiations with state actors with the aim of banning lethal autonomous weapons (Sauer, 2014). Once the risk and benefits of APM are better understood, a similar process of international attention and regulation could take place in order to govern the development and use of APM weapons. Effective governance of military APM would tend to make the impacts better (see Gubrud, 2003)

Similarly, APM technologies could also be used positively by regulatory bodies to govern the development and use of these weapons of mass destruction, nanoscale or otherwise. This can be done through the creation of precise biochemical sensors that could monitor biological and chemical weapons' compliance to treaties as well counteract them if they are deployed (Altmann, 2006; Choi, 2006). In essence, as the creation of APM weapons comes to be a reality, so could systems that are used to regulate them.

Taking these various factors together, we find it difficult to quantify the net impact to society from military APM. The potential benefit from nuclear disarmament is quite large, but so is the potential harm from rogue actors. The effects may well cancel each other out. Whether or not they do would appear to depend mainly on broader trends in geopolitical relations and non-APM military technologies. We would tentatively expect that the net effect would be beneficial due to the relative dearth of rogue actors who seek to cause global destruction and the potential for the rogue actor risk to be attenuated by advances in surveillance.

3.4 Value of APM to other Surveillance (Σ)

Surveillance, particularly mass-surveillance, is a controversial topic. The application of APM to surveillance is similarly controversial. APM could potentially have a dramatic impact on the methods and devices used to enhance or alter current surveillance networks (Drexler, 2013a). Issues of freedom, privacy, and security thus loom large for APM. Any assessment of the value of APM for surveillance will depend heavily on value judgments on these issues.

APM could enable several types of surveillance technologies, including audio/visual sensors, computers to store and process collected data, and equipment for networking it all together (Drexler, 2013a). APM sensors could be small, difficult to detect, and inexpensive to fabricate. These AP-surveillance technologies would make it relatively easy for a group to conduct a wide range of surveillance activities, including intrusive mass surveillance. For example, Anissimov et al. (2008) describe a scenario in which APM is used to fabricate millions of surveillance devices, first for foreign intelligence and later in a domestic capacity under the guise of homeland security. Per Anissimov et al. (2008), these devices would be small and disguised, making them nearly impossible to detect.

As with existing surveillance technologies, APM surveillance could have a range of impacts. Some impacts would increase the security of society. APM surveillance could reduce crime and other misdemeanors. More importantly, APM surveillance could help monitor rogue military actors (Section 3.3) or research groups developing risky AI (Section 3.5) or other possibly risky technologies. Potentially, APM surveillance could enhance counterterrorism so much that it would deter would-be terrorists from pursuing violence in the first place, typically by making the likelihood of their success low and too costly. This is known as “deterrence by denial”, and it can be especially effective against terrorists for whom the traditional deterrence by threat retaliation is ineffective, for example because they lack territory to retaliate against (Kroenig and Pavel, 2012). If APM enables the successful monitoring of rogue military actors or risky research groups, it could conceivably help avert major catastrophe.

A risk of APM surveillance is that it could be used as a tool of oppression. Surveillance tools have long been used by oppressive regimes to detect and suppress dissent. APM has the potential to significantly enhance this tendency. An era reminiscent of George Orwell’s *1984* (1949) could be ushered in through the monopolization of nanomaterials and systems. A totalitarian regime can further tighten their grip on their citizenry. This could be done with relative simplicity through the affordable, unobtrusive and abundant implementation of cameras (Drexler, 2007, 2013a; Anissimov et al., 2008). An Orwellian dystopia can be easily imagined if the ubiquitous use of APM surveillance technology becomes a norm.

Worst-case scenarios for oppression are catastrophic. The worst would be an oppressive global government or global coalition of oppressive governments. Such a government would lack competition and thus be difficult to dislodge, making it a threat to the long-term success of human civilization (Caplan, 2008). Bostrom (2006) proposes that APM could enable hostile actors to form an oppressive world government, aided by surveillance technology and other tools, though Caplan (2008) posits that this is unlikely as long as enough other actors would also have APM and be able to counter the hostile actors’ moves.

Conversely, APM surveillance could also be used to counter oppression. Inexpensive surveillance technology would be a boon to citizen “sousveillance”, in which citizens use surveillance technology to monitor for potential abuses of power (Mann et al., 2003). For example, in recent years, citizens have used cell phone cameras to record aggressive police tactics, sparking public debate about police brutality. APM could extend sousveillance capacity, reducing police brutality and potentially also reducing the prospects for large-scale and even global-scale government oppression.

The net effect of APM surveillance on society could depend heavily on government policy. Many countries have statutes that protect the public from unwarranted surveillance by restricting or prohibiting the use of surveillance devices that are not in public use. However, such statutes have not always been enforced. New surveillance technologies can also strain existing statutes. For example, in the case *Kyllo v. United States* (2001), it was ruled that, under the terms of the Fourth Amendment, a warrant was required by law enforcement agencies to use a thermal imaging device to search private homes from a public vantage point (i.e., aerial position). APM could create new legal challenges regarding surveillance. Outcomes of the corresponding policy debates and court cases and the enforcement of the rules could largely dictate the ways in which APM surveillance is used. Of course, this only applies to countries that have these sorts of legal restrictions on surveillance.

In sum, APM appears to exacerbate the familiar tension between freedom and privacy on one hand and security on the other. Whether it is through oppressive practices by state actors or continued increases in mass surveillance, APM could potentially make outcomes worse. However, APM could likewise increase state security via deterrence-by-denial tactics, monitoring of risky research groups, and increased citizen sousveillance. It is difficult to establish whether the net effect is beneficial or harmful to society. This depends on uncertain governance details as well as contentious value judgments.

3.5 The Value of APM to Artificial Intelligence (AI)

The field of artificial intelligence (AI) is of particular interest when considering the development of APM. AI is often considered alongside APM as a high-stakes, potentially transformative future technology (e.g., Joy, 2000). The AI being referred to here is generally not the AI that currently exists, which has human-level intelligence or greater in at most only select narrow domains such as playing chess or searching the internet. Instead, the AI technology of interest is artificial general intelligence (AGI), which has capabilities across a wide range of domains, and artificial superintelligence (ASI), which has significantly greater than human-level intelligence across these domains. Some analysts hold that ASI could bring massive benefits or harms, depending on how it is designed (e.g., Bostrom, 2014).

One factor in ASI risk is the computer technology available to the ASI and its precursor AI technologies. The precursor AI can include a “seed” AI which evolves into an ASI via a process called recursive self-improvement, in which the AI programs successively smarter AIs and “take off” into an ASI. While the AI is self-improving, humans can train the AI to behave in desirable ways. Thus, in terms of ASI safety, a key variable is takeoff speed, meaning how quickly the AI goes from sub-human-level seed to super-human-level ASI. Faster takeoff gives humans less time to train the AI and thus a higher probability that the ASI will behave in undesirable ways.

APM could increase AI risk by increasing computing speeds and making AI takeoff go faster. For this reason, some AI analysts have argued that AI risk would be increased if APM is developed before AI (Bostrom, 2014; Goertzel and Pitt, 2014; Goertzel, 2016; see also Shulman and Sandberg, 2010). AGI researcher Ben Goertzel expresses this concern as follows:

I think this [a very fast AI takeoff] is very unlikely, but if it happens, indeed, almost all bets are off. As I have argued elsewhere (Goertzel and Pitt, 2012), this sort of outcome becomes increasingly likely as the supporting technologies surrounding AGI become more advanced. So if we want to minimize the odds of it happening, we should make sure that advanced AGI is developed as soon as possible, relative to the advancement of other powerful technologies such as nanotech and 3D printing (Goertzel, 2016, 61).

This sort of increase in AI risk could be another major factor in the evaluation of APM, perhaps even comparable to the effects of APM on climate change and nuclear war risk. Some believe that an AI catastrophe could be severe, potentially involving human extinction (e.g., Bostrom, 2014). The probability of such an event is not well quantified due to the uncertain nature of the underlying technology; risk analysis of it has only recently begun and remains inconclusive (Barrett and Baum, 2016).

Another factor in evaluating the relative timing of APM and AI concerns the impacts that each technology could have on the other. For example, Bostrom (2014) argues that it would be better for AI to come before APM because AI could reduce APM risks. This is on account of AI's potential ability to coordinate, manage and make decisions that are ethical in accordance with its training (Goertzel and Pitt, 2014). However, this argument presupposes that, in the absence of an AI to guide it, APM would be net harmful to society. As this paper shows, it is mistaken to assume that APM is net harmful (or net beneficial) without doing the analysis. It may well be good to hasten AI development for other reasons, but that is beyond the scope of this paper.

Meanwhile, early development of APM might be able to reduce AI risks by improving surveillance. It is sometimes proposed that AI governance should include extensive surveillance in order to find any group that is developing AI and ensure that they are using adequate safety precautions (e.g., Shulman, 2009). This sort of proposal raises significant concerns regarding privacy, intellectual property, and geopolitical relations, and may do more harm than good (Baum, 2016). However, to the extent that surveillance reduces AI risks, it could be enhanced through early development of APM.

Estimating the net effect of APM on AI is difficult because future AI technology is so ambiguous. There appears to be a harmful effect from hastening AI takeoff and possibly a beneficial effect from improved surveillance, though the surveillance also raises other concerns. The net effect would thus appear to be negative, and potentially large if the risk of future AI catastrophe turns out to be large and if takeoff speed turns out to be an important factor.

3.6 The Value of APM to Space Travel and Colonization (S)

There is the potential for APM to contribute to space travel and space colonization. Space travel is not an urgent issue for human society in the same sense that (for instance) the environment or military affairs are. However, it is of enormous significance to the long-term fate of human

civilization. The opportunities of space travel are quite literally astronomical. Furthermore, space travel remains heavily technology-limited; hence, innovation in the spheres of small, lightweight and highly affordable manufacturing systems could be of astronomical value.

Currently, the cost of space travel is enormous, and space missions are correspondingly limited. With existing technology, the cost of bringing materials into space is far too great to rationalize using them for colonization (McKendree et al., 2005). APM could change this. APM could create materials that cost less, weigh less, and are more durable for space travel (Drexler, 2013a). APM could also enhance asteroid mining, which would further lower costs and expand opportunities by enabling space missions to harvest resources in nearby space locations rather than transporting them from Earth. Likewise, APM may enable the creation of energy sources that are more efficient, and more powerful, and better able to withstand the conditions of space travel. This could go a long way towards enabling human travel in outer space.

APM could also be of great benefit for what humanity does when it arrives at new planets and other extraterrestrial destinations. One of the more ambitious proposals is to engage in terraforming, which is the creation of livable environments on otherwise inhospitable extraterrestrial planets and other bodies of mass (Fogg, 1995,1998; Graham, 2004; McKay and Marinova, 2001). Terraforming proposals involve approaches such as using photosynthetic organisms to create an oxygen-rich atmosphere (Friedmann and Ocampo-Friedmann, 1995; Hiscox and Thomas, 1995), heating polar ice caps, such as the water-and-CO₂ ice cap on Mars, to create a greenhouse gas atmosphere to warm the planet (Sagan, 1973; Mole, 1995; but see Fogg, 1995, 1998), and putting megascale mirrors in a planet's orbit to reflect more radiation towards it and warm it (Birch, 1992).

For each of these approaches, it is generally assumed that the process would take at least several hundred years to create an adequate environment. An exception is the Birch (1992) proposal for megascale mirrors, which is claimed to require as few as 50 years, but this proposal would require an engineering capacity that is well beyond what is currently available. APM could dramatically reduce the time and costs for each of the methods. APM would achieve this via its ability to make materials of small size that can be programmed to complete complex tasks autonomously (Drexler, 1986). Likewise, small sized APM terraforming systems would make them easier and less expensive to transport.

The overall value of APM to space travel appears to be positive and quite large. APM shows promise to overcome some major technological hurdles that, at present, greatly constrain space travel. However, space travel is arguably less urgent than the threats to human survival that APM could affect, as discussed in Sections 3.2-3.5. Indeed, a single one of these catastrophes could harm human civilization so much that it would preclude any further space travel. Nonetheless, value of APM to space travel does offer a reason for APM R&D investment.

3.7 Other Value of APM (X)

Finally, it is possible that APM could have significant positive or negative value for factors not covered in Sections 3.1-3.6. Sections 3.1-3.6 would appear to cover the major impacts of APM that have thus far been identified. For any currently unknown impacts, it is impossible to comment on whether they would be positive or negative—otherwise, they would not be truly

unknown. Speculation on whether such impacts would tend to be positive or negative is beyond the scope of this paper, and indeed may only become possible as the passage of time reveals new potential APM impact areas.

There is one other scenario worth discussing: the so-called “grey goo” problem, a scenario where hypothetical self-replicating molecular machines (also called molecular assemblers) that spreads around the world consuming all organic matter, thereby causing the extinction of humans and all other Earth species, or at least all species found at the surface (Freitas, 2000; Hughes, 2007). The grey goo scenario has attracted extensive attention, especially in popular media (Diem, 2017; Torres and Boghossian, 2016; Institute of Physics, 2004; Rincon, 2004;). However, there is reason to not worry about it. While APM might be able to create such a thing, it would not be easy. APM devices could readily be designed to not manufacture grey goo. Making dedicated APM devices for it would require extensive technical capacity (Freitas, 2006). (Grey goo could potentially also be made without APM, but with even more difficulty.) Fortunately, there are few people in the world who want to destroy all life on Earth, and they are unlikely to have the requisite capacity. So while the grey goo scenario might be possible, it is too unlikely to be a significant factor relative to the much higher probability scenarios captured by Equation 1 and Sections 3.1-3.6.

4. The Overall Value of APM (V)

What is the overall societal effect of APM? Per Equation 1, this would be the sum of the effect for each sector discussed in Section 3. To summarize from Section 3:

- The value of APM for general material wealth is likely to be moderately large and positive.
- APM shows strong potential for very large environmental benefits and little potential for environmental harms.
- In the military sector, APM could have a large benefit from nuclear disarmament and a large harm from rogue actors, though the rogue actor harm could be attenuated by various countermeasures.
- It is difficult to establish whether the net effect from surveillance is net beneficial or harmful to society.
- There appears to be a harmful effect from hastening AI takeoff, the size of which is unclear but potentially large. There may also be a beneficial effect from improved surveillance, though the surveillance also raises other concerns.
- For space travel, the value of APM is quite large, but it is perhaps not as urgent as the various threats to human civilization covered in other applications of APM.

Rearranging these points, the key benefits of APM are in general material wealth, environmental protection, nuclear disarmament, and space travel. The key harms are in rogue actor violence and hastened AI takeoff. There is also a significant ambiguity in the net effect of surveillance. If general material wealth is a smaller factor, space travel is non-urgent, and surveillance is too ambiguous to factor in, that leaves a comparison between the benefits of environmental protection and nuclear disarmament and the harms of rogue actor violence and AI takeoff. Each of these factors could be quite large, so, superficially, with two factors on each side, the net effect may seem to be approximately zero.

Taking a closer look, a case can be made that the largest factor is the benefit from environmental protection. Climate change is a large and well-established risk. Without APM, humanity is on track to largely fail to reduce the risk. With APM, it could become easy to eliminate most or all of the risk. Other environmental benefits, such as to biodiversity loss, could also be significant. In contrast, the other risks—nuclear war, rogue actor violence, and especially AI takeoff—are more speculative. Furthermore, for these three risks, APM would be a relatively small factor: nuclear war and rogue actor violence depend heavily on geopolitical factors, while AI takeoff depends heavily on software development. This suggests a net positive value of APM.

But this conclusion may not hold. In particular, it is sensitive to the relative sizes of the risks from the environment, nuclear war, rogue actor violence, and AI takeoff. If rogue actor violence and/or AI takeoff turn out to be relatively large risks, then the net value of APM could be negative. Unfortunately, none of these risks are characterized well enough to enable comparison. Assessing the value of APM requires an integrated analysis of each of these risks in order to evaluate the tradeoff between them that APM poses, such as has been proposed by Baum and Barrett (2017).

Finally, the societal impacts of APM will depend on how society uses APM. This holds for each of the sectors in Equation 1 and Section 3. For example, if society uses APM surveillance to detect terrorists but without mass invasions of privacy, then APM surveillance could be net beneficial. Or, if high-performance APM computer hardware was not used for AI takeoff, then it would not factor into the impacts of this type of AI. Therefore, any evaluations of APM impacts will be contingent on future human actions. Likewise, plans for use and governance of APM should be factored into decisions on whether to invest in APM R&D. Such plans should likewise be pursued if APM is indeed being developed, in order to improve its societal impacts.⁶

5. Conclusion

This paper has outlined and discussed several areas where APM may have especially large impacts, including both positive and negative impacts. We studied potential effects of APM on general material wealth, the environment, military affairs, surveillance, artificial intelligence, and space travel. Positive effects were found for material wealth, the environment, military affairs (specifically nuclear disarmament), and space travel. Negative effects were found for military affairs (specifically rogue actor violence and AI). The net effect for surveillance was ambiguous. The effects for the environment, military affairs, and AI seem to be the largest, with the environment perhaps being the largest of these, which would suggest that APM would be net beneficial to society. However, these issue areas are not well quantified, nor are the effects of APM to the issues, so no definitive conclusion can be made about the net impacts of APM. This paper offers a framework for evaluating the impacts of APM, but it does not offer the final word.

Therefore, some caution is warranted with APM R&D. One should be conscious in opening up this veritable Pandora's Box in search of utopian technologies (Ferrari and Nordmann, 2010). The unprecedented potential benefits of APM come with the potential for commensurate or greater harms. It is important to consider both potential benefits and potential harms in making

⁶ The official journal version contains a writing mistake in which the second half of this paragraph is omitted. We sincerely regret this error.

APM R&D decisions, and to weigh them against each other. APM researchers and advocates should resist the temptation to focus exclusively on the potential benefits in order to promote R&D investment. This paper's analysis finds some reason to believe that net benefits are likely, but further study could easily tip the scales in the opposite direction.

Future research could resolve some ambiguities and bring more confidence to the overall evaluation. First, it may be possible to identify additional major impacts of APM beyond those covered in this paper. While this paper has aimed to include all major impacts, it is always possible that some have been overlooked. Assuming that nothing major has been overlooked, then perhaps the most urgent task is integrated analysis of the various risks that could be affected by APM in order to clarify their relative importance. Such analysis would of course also be helpful for other issues related to these risks. Another key topic not explored in any depth in this paper is the role of governance to increase the benefits and reduce the harms of APM. Ideally, any investment in APM R&D would coexist with investment in good governance measures so as to ensure that APM is beneficial to society. Finally, a full analysis of APM R&D decisions should also include the cost of R&D and any other significant factors, including factors of relevance to ethical frameworks other than the consequentialism employed in this paper. APM R&D cost analysis could likewise factor in assessment of the feasibility of APM; if APM is found to be fundamentally infeasible, then there is no point in investing in R&D.

Fortunately, there is some time available to take on these matters. APM is still in a largely conceptual phase. Some believe that APM will eventually be developed and widely implemented, but that point is not yet near. Meanwhile, there is opportunity for serious analysis and discussion regarding the potential governance of APM, including R&D investment. Given the potential effects of APM across many important societal sectors, this opportunity should be seized.

Acknowledgments

We thank Tony Barrett, Jacob Haqq-Misra, Eric Drexler and an anonymous reviewer for feedback on an earlier draft, and Joshua Pearce and Anders Sandberg for background on nanotechnology. Any errors are the authors' alone. The views in the paper are the authors' alone and not the views of the Global Catastrophic Risk Institute.

References

1. Altmann, J. (2006). *Military nanotechnology: Potential applications and preventive arms control*. London: Routledge.
2. Anissimov, M., Buerger, M., & Burgess, S. (2008). The center for responsible nanotechnology scenario project. *Nanotechnology Perceptions*, 1.
3. Barrett, A. M., Baum, S. D., & Hostetler, K. (2013). Analyzing and reducing the risks of inadvertent nuclear war between the United States and Russia. *Science & Global Security*, 21(2), 106-133.
4. Barrett, A. M., and S. D. Baum., (2017). "A Model of Pathways to Artificial Superintelligence Catastrophe for Risk and Decision Analysis." *Journal of Experimental & Theoretical Artificial Intelligence*, 29(2), 397-414.
5. Battelle Memorial Institute, Foresight Nanotech Institute (2007). *Productive Nanosystems A Technology Roadmap*.

http://www.foresight.org/roadmaps/Nanotech_Roadmap_2007_main.pdf

6. Baum, (S. D. 2015). Winter-safe deterrence: The risk of nuclear winter and its challenge to deterrence. *Contemporary Security Policy*, 36(1), 123-148.
7. Baum, S. D. (2016). On the Promotion of Safe and Socially Beneficial Artificial Intelligence. *AI & Society*, 1-9.
8. Baum, Seth D. and Anthony M. Barrett. (2017). The most extreme risks: Global catastrophes. In Vicki Bier (editor), *The Gower Handbook of Extreme Risk*. Farnham, UK: Gower.
9. Baum, Seth D., Timothy M. Maher, Jr., and Jacob Haqq-Misra. (2013). Double catastrophe: Intermittent stratospheric geoengineering induced by societal collapse. *Environment, Systems and Decisions*, vol. 33, no. 1 (March), pages 168-180.
10. Birch, P. (1992). Terraforming Mars Quickly. *Journal of The British Interplanetary Society*, 45, 331-340.
11. Bostrom, N. (2014). *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press.
12. Bostrom, N. (2006). What is a singleton. *Linguistic and Philosophical Investigations*, 5(2), 48-54.
13. Broome, J. (2004). *Weighing Lives*. Oxford: Oxford University Press
14. Caplan, B. (2008). The totalitarian threat. In Bostrom, N., Ćirković, M. (eds). *Global Catastrophic Risks*. Oxford: Oxford University Press, 504–519.
15. Chandy, L., & Gertz, G. (2011, May 17). Two Trends in Global Poverty. Retrieved August 04, 2016, from <https://www.brookings.edu/opinions/two-trends-in-global-poverty/>
16. Choi, C. Q. (2006, May 23). Nano World: Nano-loaded wireless sensors. Retrieved July 02, 2016, from <http://phys.org/news/2006-05-nano-world-nano-loaded-wireless-sensors.html>
17. Ćirković, M. M. (2012). Small theories and large risks—is risk analysis relevant for epistemology? *Risk Analysis*, 32(11), 1994-2004.
18. Collingridge, D. (1980). *The social control of technology*. New York: St. Martin's Press.
19. Deep Cuts Commission (2010). *Preparing for Deep Cuts: Options for Enhancing Euro-Atlantic and International Security*. First Report of the Deep Cuts Commission.
20. Diem, N. (2017, March 06). Will Gray Goo Cause The End of Civilization? Retrieved May 21, 2017, from <https://www.i4u.com/2017/03/121250/will-gray-goo-cause-end-civilization>
21. Drexler, K. E. (1986). *Engines of Creation*. Garden City, NY: Anchor Press/Doubleday.
22. Drexler, K. (2007, January). The Stealth Threat: An Interview with Eric Drexler. Retrieved October 14, 2016, from https://www.researchgate.net/profile/Eric_Drexler2/publications
23. Drexler, K. E. (2013a). *Radical abundance: How a revolution in nanotechnology will change civilization*. New York, New York: PublicAffairs.
24. Drexler, K. (2013b, October 28). Big nanotech: An unexpected future. *The Guardian*. Retrieved January 30, 2016, from https://www.theguardian.com/science/small-world/2013/oct/28/big-nanotech-unexpected-future-apm?CMP=tw_t_gu
25. Drexler, K., & Pamlin, D. (2013). *Nano-solutions for the 21st Century*. Retrieved from <http://www.oxfordmartin.ox.ac.uk/publications/view/1349>
26. Easterlin, R. A., McVey, L. A., Switek, M., Sawangfa, O., & Zweig, J. S. (2010). The happiness–income paradox revisited. *Proceedings of the National Academy of Sciences*, 107(52), 22463-22468.
27. Farrell, H. H., & Levinson, M. (1985). Scanning tunneling microscope as a structure-modifying tool. *Physical Review B*, 31(6), 3593.
28. Ferrari, A., & Nordmann, A. (2010). *Beyond Conversation: Some Lessons for Nanoethics*.

- NanoEthics, 4(2), 171-181. Retrieved January 31, 2016.
29. Fogg, M. J. (1995). Terraforming: Engineering planetary environments. Warrendale, PA, U.S.A.: Society of Automotive Engineers.
 30. Fogg, M. (1998). Terraforming Mars: A review of current research. *Advances in Space Research*, 22(3), 415-420. doi:10.1016/s0273-1177(98)00166-5
 31. Freitas Jr, R. A. (1999). *Nanomedicine, Volume I: Basic Capabilities*. Austin, TX: Landes Bioscience.
 32. Freitas Jr, R. A. (2000). Some limits to global ecophagy by biovorous nanoreplicators, with public policy recommendations. *Foresight Institute*.
<http://www.rfreitas.com/Nano/Ecophagy.htm>
 33. Freitas Jr, R. A. (2006). Molecular Manufacturing: Too Dangerous to Allow? *Nanotechnology Perceptions*, 2, 15-24. Retrieved January 30, 2016.
 34. Freitas Jr, Robert A. "Personal Choice in the Coming Era of Nanomedicine," in Allhoff, F., Lin, P., Moor, J., & Weckert, J. (2007). *Nanoethics: The ethical and social implications of nanotechnology*. New York: John Wiley, 161-172.
 35. Freitas Jr, R. A. "Comprehensive Nanorobotic Control of Human Morbidity and Aging," In, G. M. Fahy, M. D. West, L. S. Coles, S. B. Harris, (2010). *The Future of Aging*. 686-805. Dordrecht, Heidelberg, New York, London: Springer.
 36. Freitas Jr, R. A., & Merkle, R. C. (2004). *Kinematic self-replicating machines*. Landes Bioscience.
 37. Friedmann, E. I., & Ocampo-Friedmann, R. (1995). A primitive cyanobacterium as pioneer microorganism for terraforming Mars. *Advances in Space Research*, 15(3), 243-246. doi:10.1016/s0273-1177(99)80091-x
 38. Goertzel, B. (2016, February). Infusing Advanced AGIs with Human-Like Value Systems: Two Theses. *Journal of Evolution and Technology*, 26(1), 50-72.
 39. Goertzel, B., & Pitt, J. (2014, February). Nine Ways to Bias Open-Source Artificial General Intelligence Toward Friendliness. *Journal of Evolution and Technology*, 22(1), 116-131. doi:10.1002/9781118736302.ch4
 40. Gomer, R. (1986). Possible mechanisms of atom transfer in scanning tunneling microscopy. *IBM Journal of Research and Development*, 30(4), 428-430.
 41. Graham, J. M. (2004). The Biological Terraforming of Mars: Planetary Ecosynthesis as Ecological Succession on a Global Scale. *Astrobiology*, 4(2), 168-195. doi:10.1089/153110704323175133
 42. Grunwald, A. (2010). From speculative nanoethics to explorative philosophy of nanotechnology. *Nanoethics*, 4(2), 91-101.
 43. Gubrud, M.A. (2003). Nanotechnology and international security, Paper presented at the Fifth Foresight Conference on Molecular Nanotechnology, Palo Alto, 5–8 November 2003.
 44. Häggström, O. (2016). *Here Be Dragons: Science, Technology and the Future of Humanity*. Oxford: Oxford University Press.
 45. Heller, J., & Peterson, C. (2006). *Nanotechnology and Surveillance*. Retrieved July 2, 2016, from <https://www.foresight.org/policy/brief7.html>
 46. Hellman, M. (2008). Risk analysis of nuclear deterrence. *The Bent of Tau Beta Pi*, 99(2), 14.
 47. Hiscox, J.A. and Thomas, D.J. (1995) Genetic modification and selection of microorganisms for growth on Mars. *J. Br. Interplanet. Soc.* 48, 419–426.
 48. Hughes, James. "Global Technology Regulation and Potentially Apocalyptic Technological Threats," in Allhoff, F., Lin, P., Moor, J., & Weckert, J. (2007). *Nanoethics: The ethical and*

- social implications of nanotechnology. New York: John Wiley, 161-172.
49. Institute of Physics. (2004, June 9). Nanotechnology Pioneer Slays 'Grey Goo' Myths. ScienceDaily. Retrieved May 21, 2017 from www.sciencedaily.com/releases/2004/06/040609072100.html
 50. Jones, C. I. (2015). Pareto and Piketty: The macroeconomics of top income and wealth inequality. *Journal of Economic Perspectives*, 29(1), 29-46.
 51. Jones, G., Whitaker, M., & King, M. (2011). Speculative Ethics: Valid Enterprise or Tragic Cul-De-Sac?, *Bioethics in the 21st Century*, Prof. Abraham Rudnick (Ed.), InTech, DOI: 10.5772/19684. Available from: <https://www.intechopen.com/books/bioethics-in-the-21st-century/speculative-ethics-valid-enterprise-or-tragic-cul-de-sac->
 52. Jones, R. A. L. (2004). *Soft machines: nanotechnology and life*. Oxford University Press.
 53. Joy, B. (2000, April 01). Why the Future Doesn't Need Us. *Wired*. Retrieved January 31, 2016, from <http://www.wired.com/2000/04/joy-2/>
 54. Karlsson, R. (2014). Theorizing sustainability in a post-Concorde world. *Technology in Society*, 39, 1-9. doi:10.1016/j.techsoc.2014.05.004
 55. Koepsell, D. (2009). Let's Get Small: An Introduction to Transitional Issues in Nanotech and Intellectual Property. *NanoEthics*, 3(2), 157-166. Retrieved January 31, 2016.
 56. Kroenig, M., & Pavel, B. (2012). How to Deter Terrorism. *The Washington Quarterly*, 35(2), 21-36.
 57. *Kyllo v. United States* (Certiorari to the United States Court of Appeals for the Ninth Circuit June 11, 2001).
 58. Lackner, K. S., Brennan, S., Matter, J. M., Park, A. H. A., Wright, A., & Van Der Zwaan, B. (2012). The urgency of the development of CO2 capture from ambient air. *Proceedings of the National Academy of Sciences*, 109(33), 13156-13162.
 59. Lyubomirsky, S. (2008). *The how of happiness: A scientific approach to getting the life you want*. New York: Penguin.
 60. Mann, S., Nolan, J., & Wellman, B. (2003). Sousveillance: Inventing and Using Wearable Computing Devices for Data Collection in Surveillance Environments*. *Surveillance & Society*, 1(3), 331-355.
 61. Marchant, G. E., Sylvester, D. J., & Abbott, K. W. (2008). Risk management principles for nanotechnology. *NanoEthics*, 2(1), 43-60.
 62. McKay, C. P., & Marinova, M. M. (2001). The Physics, Biology, and Environmental Ethics of Making Mars Habitable. *Astrobiology*, 1(1).
 63. McKendree, T., Freitas Jr, R., Globus, A., Levit, M., Sherrill, C., Li, M., & Merkle, R. (2005). Challenges of Molecular Nanotechnology for Space Exploration. American Institute of Aeronautics and Astronautics. Retrieved January 31, 2016, from <http://arc.aiaa.org/doi/abs/10.2514/6.2005-6814>
 64. Mills, M. J., Toon, O. B., Lee - Taylor, J., & Robock, A. (2014). Multidecadal global cooling and unprecedented ozone loss following a regional nuclear conflict. *Earth's Future*, 2(4), 161-176.
 65. Mole, R. A. (1995). Terraforming Mars with Four War-Surplus Bombs. *Journal of The British Interplanetary Society*, 48, 321-324.
 66. Møller, M., Jarvis, S. P., Guérinet, L., Sharp, P., Woolley, R., Rahe, P., & Moriarty, P. (2017). Automated extraction of single H atoms with STM: tip state dependency. *Nanotechnology*, 28(7), 075302.
 67. Ng, Y. K. (1980). *Welfare Economics: Introduction and Development of Basic Concepts*.

New York: Wiley

68. Nordmann, A. (2007). If and then: a critique of speculative nanoethics. *Nanoethics*, 1(1), 31.
69. Nordmann, A., & Rip, A. (2009). Mind the gap revisited. *Nature Nanotechnology*, 4(5), 273.
70. Orwell, G. (1949). 1984. London: Secker & Warburg.
71. Rees, Martin, (2003). *Our Final Century: Will the Human Race Survive the Twenty-first Century?* Oxford: William Heinemann.
72. Reynolds, G. H. (2001). Environmental Regulation of Nanotechnology: Some Preliminary Observations. *The Environmental Law Reporter*, 31(6).
73. Rincon, P. (2004, June 09). Science/Nature | Nanotech guru turns back on 'goo' Retrieved May 21, 2017, from <http://news.bbc.co.uk/2/hi/science/nature/3788673.stm>
74. Roache, R. (2008). Ethics, speculation, and values. *Nanoethics*, 2(3), 317.
75. Robinson, J. (2015, February 25). The 20 Fastest-Growing Economies This Year. Retrieved August 04, 2016, from <http://www.bloomberg.com/news/articles/2015-02-25/the-20-fastest-growing-economies-this-year>
76. Sagan, Carl, "Planetary Engineering on Mars," *Icarus*, 20, 513-514 (1973).
77. Sandberg, A. (2014, May 29). The five biggest threats to human existence. Retrieved August 03, 2016, from <http://theconversation.com/the-five-biggest-threats-to-human-existence-27053>
78. Sandler, R. (2007). Nanotechnology and Social Context. *Bulletin of Science Technology Society*, 27(6), 446-454.
79. Sauer, F. (2014, April 24). The Road to Geneva: ICRAC and the Campaign headed to CCW Expert Meeting. Retrieved August 03, 2016, from <http://icrac.net/2014/04/the-road-to-geneva-icrac-and-the-campaign-headed-to-ccw/>
80. Seligman, M. E. (2012). *Flourish: A visionary new understanding of happiness and well-being*. New York: Simon and Schuster.
81. Selin, C. (2007). Expectations and the Emergence of Nanotechnology. *Science, Technology & Human Values*, 32(2), 196-220.
82. Sherwood, S. C., & Huber, M. (2010). An adaptability limit to climate change due to heat stress. *Proceedings of the National Academy of Sciences*, 107(21), 9552-9555.
83. Shulman C (2009) Arms control and intelligence explosions. In 7th European Conference on Computing and Philosophy (ECAP), Bellaterra, Spain, July 2-4.
84. Shulman, C., & Sandberg, A. (2010). Implications of a Software-limited Singularity. ECAP10: VIII european conference on computing and philosophy. K. Mainzer. Munich: Verlag.
85. Smalley, R. E. (2001). Of chemistry, love and nanobots. *Scientific American*, September, 76-77.
86. Stevenson, B., & Wolfers, J. (2008). Economic growth and subjective well-being: Reassessing the Easterlin paradox. National Bureau of Economic Research, Working Paper No. w14282.
87. Tavoni, M, Socolow, R, (2013). Science and Policy of Negative Emission Technologies (special issue). *Climatic Change* 118(1), 1-149.
88. Timmermans, J., Zhao, Y., & van den Hoven, J. (2011). Ethics and nanopharmacy: Value sensitive design of new drugs. *Nanoethics*, 5(3), 269-283.
89. Torres, P., & Boghossian, P. (2016). The Looming Extinction of Humankind, Explained. Retrieved May 21, 2017, from https://motherboard.vice.com/en_us/article/armageddon-comma-explained

90. Tuomisto, H. L., & Teixeira de Mattos, M. J. (2011). Environmental impacts of cultured meat production. *Environmental science & technology*, 45(14), 6117-6123.
91. UNEP (2010) Assessing the Environmental Impacts of Consumption and Production: Priority Products and Materials, A Report of the Working Group on the Environmental Impacts of Products and Materials to the International Panel for Sustainable Resource Management. Hertwich, E., van der Voet, E., Suh, S., Tukker, A., Huijbregts M., Kazmierczyk, P., Lenzen, M., McNeely, J., Moriguchi, Y.
92. United Kingdom - Royal Society and Royal Academy of Engineering (RS-RAE). (2004). Nanoscience and nanotechnologies: opportunities and uncertainties. Report available for download from: <http://www.nanotec.org.uk/finalReport.html>.
93. United States. (2010). Nuclear posture review report. Washington, DC: U.S. Dept. of Defense. (OIG). (2011). EPA Needs to Manage Nanomaterial Risks More Effectively (Rep. No. 12-P-0162).
94. Wilson, G. (2013). Minimizing global catastrophic and existential risks from emerging technologies through international law. *Virginia Environmental Law Journal* 31:307-364.
95. Worstall, T. (2015, May 7). IMF Pegs India To Be One Of World's Fastest Growing Economies. Retrieved August 04, 2016, from <http://www.forbes.com/sites/timworstall/2015/05/07/imf-pegs-india-to-be-one-of-worlds-fastest-growing-economies/#606a11bb6822>
96. Xia, L., Robock, A., Mills, M., Stenke, A., & Helfand, I. (2015). Decadal reduction of Chinese agriculture after a regional nuclear war. *Earth's Future*, 3(2), 37-48.