The Time To Talk About The Ethics Of Quantum Computing Is Now

Quantum computing will radically change the world

It's up to all of us to ensure it does so ethically

The age of quantum computers is nearly upon us. Although the capabilities of quantum computers do not yet surpass those of classical computers, this will soon change.

The significance of this cannot be overstated. Quantum computers—and quantum computing technologies, more generally—will have revolutionary impact in areas like computer security and cryptography, quantum systems modeling, optimization, and health, with potential commercial applications across nearly all industries. While such technological innovation is both welcome and exciting, quantum computing also raises extremely serious ethical concerns—concerns we must begin to address immediately.

It is often tempting, when focusing on technological innovation, to set aside the ethical concerns raised by a technology until that technology has been developed and deployed. In the case of quantum computing, however, we believe this is a serious mistake[1]. We believe, further, that we must identify and address the ethical issues raised by quantum computing now. But why, one might ask? Why impede current innovation and growth to focus on ethics, when *large-scale* quantum computers are likely years away? There are two important reasons.

First, while quantum computers will be capable of doing things that may benefit the world in immeasurable ways, they also have the potential to cause great harm, perhaps in ways we do not yet even recognize. Seizing the unique opportunity to identify and address the important ethical issues raised by quantum computing now is our best means of mitigating or eliminating these significant future risks to people, businesses and society.

Second, while large-scale universal quantum computers are indeed many years away, smaller quantum computers (so-called "Noisy Intermediate-Scale Quantum" (NISQ) computers), as well as some quantum computing-related technologies such as advanced sensors, are likely to

[1] It's likely a mistake for *any* technology, but that's another topic.

begin appearing in the next few years. These intermediate machines, coupled with, say, advances in quantum computing algorithms, may be able to perform computational tasks that are practically impossible for classical computers, tasks that may also have drastic, unexpected, and potentially very harmful, effects.

The field of quantum computing— both hardware and software— is active, focused, and growing. It's an exciting time for the space. But there's a danger in the field's rapid advancement as well: we know some potential risks of the technology are now in our immediate future, while the possibility of serious unforeseen risks is all too real.

Given the unprecedented scale of the impact quantum computing will have on human well-being, national security, and global society, we must marshall all available resources to address the ethical implications of this technology without delay. As we suggest in the title, the time to talk about the ethics of quantum computing is now.

What Are Quantum Computers?

Not everyone who should be thinking about the ethics of quantum computers knows about them, or, if they do, really understands them. So let's begin by explaining, in layman's terms, just what quantum computers are.

It's a common misconception to think that quantum computers will simply be faster computers. Quantum computing and classical computing are fundamentally different, and quantum computers will not be better or faster than classical computers *at everything*. Rather, quantum computers will possess certain unique features that will make them particularly well-suited to certain computational tasks, including tasks that are difficult or even effectively impossible for classical computers. Quantum computers are complicated machines involving complicated physics, but one does not require a deep understanding of either quantum mechanics or computing to understand the basic features that will make them so unique and powerful.

Most people know that classical computers—everything from the phone in your hand, to the computer at your work, to the NSA's most powerful supercomputer—are at root binary: they encode their most basic unit of information in ones and zeros. This basic unit of information is called a "bit," and more and more complex combinations of bits are used to encode more and more complex information. Combinations of transistors are then used to construct logic gates which function as input/output devices, instantiating the basic boolean operators of "And," "Not," and "Or". For example, if all of the inputs to an "And' gate are 1, it outputs a value of 1; if any of its

inputs are 0, it outputs a value of 0. Combinations of logic gates yield more complex modules, including modules that can perform addition. Once you have addition you have multiplication, and once you can multiply you can in theory perform every computational task. From these simple foundations all the wondrous capabilities of modern computers flow.

For all the incredible things they can do, there are several things classical computers cannot do, or cannot do particularly well. One such thing is *optimization*, or finding the best solution from among several possible solutions. Consider an example inspired by Dr. Talia Gershon. Suppose you are having 12 friends over for dinner, and want to consider the various ways you might arrange the seating. How many combinations do you think there are for twelve guests? The answer, shockingly, is around 479 million. And for each additional guest the number rises exponentially: so, if Kelly again shows up with an unannounced dinner date, the possible seating arrangements rises to 6.2 billion. Dinner may be cold by the time you've considered every option.

Classical computers are not particularly good at computations that involve exponential scaling like this. This is also why they have trouble modeling complex molecules or other quantum systems: to simulate a molecule—say a promising pharmaceutical drug—we have to account for every electron attraction to the nuclei and repulsion between electrons, and that number grows exponentially given the size of the molecule, because every electron affects every other electron. And so, to add one more requires you to recalculate the whole thing.

But quantum computers are different, and their difference boils down to two important quantum phenomena: *Superposition* and *Entanglement*. As already noted, classical computers are binary. But quantum computers are not. Rather than being limited to one of two states, as in a classical computer, quantum computers—or, more specifically, quantum information—can have states that exist in a *superposition* of zero and one. This basic unit of quantum information is known as a "qubit". Very roughly, for a qubit to be in a superposition of zero and one means that it is both zero and one (and, strictly speaking, all points in between) *at the same time*. So consider again the dinner party example: a classical computer has to consider each possible seating arrangement individually and then compare them all. That is incredibly time consuming. In fact, once you reach a certain scale, such optimization problems become practically impossible for classical computers. In contrast, a sufficiently large quantum computer can take its qubits and go into a superposition of all the possible configurations at the same time, reducing the time required for the computation to a manageable amount. This feature allows quantum computers to perform certain tasks exponentially faster than classical computers.

Quantum computers also exploit another feature of the quantum world, *entanglement*. Entanglement is a state that subatomic particles may be in together. Again, roughly speaking, if, say, two qubits are entangled then they are paired in such a way that measuring one qubit yields information about the second qubit. This simple feature is incredibly powerful, as it allows for the creation of correlated states of a number of qubits together, and the ability to move information around in your quantum system very efficiently. These correlated states help to provide the increased speed quantum computers enjoy for certain computations. Working together, these features allow quantum computers to process information and perform certain tasks exponentially faster than classical computers.

The Promise and Perils of Quantum Computing

Quantum computers will not be better or faster than classical computers in every respect. Instead, they will excel at *certain types* of computations. Although their application will thus be somewhat limited in scope, they will nevertheless excel in ways that will make them extremely powerful. This power may be channeled in many beneficial directions, as we will discuss next. But with such tremendous power also comes the potential for serious harm.

We'll consider a few of the more exciting potential benefits of quantum computers, before considering some of the more concerning risks that they raise. Neither list will be exhaustive, but each should provide a sense of the significant impact quantum computers might have.

The Promise

Quantum computers offer the possibility of making the world a much better place. A likely early use of quantum computers will be as "quantum simulators", which is precisely what Richard Feynman, the late, great physicist, initially conceived them for. Simulating quantum systems is something that classical computers are not very good at, but quantum computers can in principle simulate any other quantum system. As we shall see, this would be extremely useful.

Current thinking is that the best method for quantum simulation will be by using a classical/quantum hybrid machine that relies on variational quantum eigensolving (VQE). Such hybrid machines seem likely to appear soon. What will they be good for? Consider how they might be used in three areas:

1. A quantum simulator could help produce extremely sophisticated models of the Earth's climate, and thus extremely accurate models of climate change. Given the dire news from

The Intergovernmental Panel on Climate Change's (IPCC) most recent report, any help we can find in potentially helping to mitigate the existential risk posed by climate change, the better.

2. Chemical modelers working in pharmaceuticals must approximate how a novel molecule might behave and interact with other compounds, and then test it in real world trials to find out. A quantum computer could instead determine exactly how that molecule would behave and interact with other compounds with no real-world trial-and-error testing required. More generally, the computational power of quantum computers will impact all areas of healthcare that require the analysis or modeling of complex data. It is not hyperbole to suggest that quantum computers may one day help cure cancer; it's also not to suggest that they may eliminate the need for animal testing.
3. Consider what's required for a thorough analysis of the energy-intensive nitrogenase reaction used to make fertiliser. It would take a 100-qubit quantum computer mere days, if not hours, to do this. This is in contrast to a modern supercomputer, which would need billions of years. This improvement in the process of manufacturing fertilizer could yield a cut of 1-2% in global natural-gas consumption.

But simulating quantum systems is not all that quantum computers might be good at. Consider another potential use:

4. We noted earlier that quantum computers will exceed at optimization. What does this mean in practice? Improved optimization can include everything from helping to produce incredibly efficient supply chains, to more optimized search, to calculating more efficient transportation delivery routes. Optimizing delivery routes might not sound as appealing as curing cancer, but it would not only have a significant financial impact for businesses relying on it, but might reduce the time vehicles spend on the road, thus reducing global CO_2 emissions.

All of these potential uses are, of course, welcome. But it's crucial to note that even they raise important ethical questions and the possibility of harm. However, the potential harm here is the harm of injustice, the unfair distribution of the benefits of quantum computing. This is an extremely complicated set of issues that we can only hint at here. It's easy to say that the best way

to ensure that the benefits of a technology are fairly distributed is to leave it to the market. Very often, this is indeed the case. But it is not always this simple.

For example, in the case of quantum computers, much of the science that is now underpinning private development of the technology was created in publicly funded universities and research centers. If taxpayers supported the research that now makes the development of quantum computers possible, and that technology leads to benefits of widespread or universal concern, such as cures for fatal diseases, should the public not have some access to these benefits apart from their ability to purchase them in the market?

Or, what of the proper allocation of limited quantum computing resources in the near future? It's likely that for the next five to ten years there will be very limited quantum computing power available. If, for example, the gravest existential threat our species currently faces is that of global climate change, and quantum computers may play a key role in helping to mitigate this risk, then should these limited computing resources be used for much less pressing commercial uses— say, building financial trading models at a hedge fund— or should much of it instead be allocated to meet this most pressing of global risks? As a society, which should we prefer?

When considering the impact of a revolutionary technology at the scale of quantum computing, the just distribution of the benefits of that technology is an equally pressing, but incredibly complex, ethical question.

The Perils

As with most technologies, while quantum computing promises clear and welcome benefits, it also introduces the possibility of serious risks and harms. Perhaps the best known and most widely discussed risk quantum computers pose is in the area of cryptography, most especially internet security.

Consider that all your secure online activities—from sending email to uploading your personal files to the cloud to online banking—are secured through the use of certain standard encryption protocols. These internet security protocols, such as public key and electronic signature, and most modern cryptography, are based on a very simple fact: While it is easy for classical computers to determine whether any given numbers are factors of a large prime number in a short amount of time (e.g. determining that AxB = N), it is extremely difficult for those same computers to find all the possible factors of a large prime number in any reasonable amount of time (e.g. finding all the possible factors of N where N is a very large number).

But here is one area where quantum computers are different. Almost 25 years ago, mathematician Peter Shor discovered a quantum algorithm capable of solving factoring in polynomial (i.e. very fast) time. In layman's terms, this means that a quantum computer could theoretically crack all existing internet security, and quickly. In case it is not clear, that could be catastrophically bad, not only for personal, corporate, or governmental privacy, but for global finance and, really, the entire modern system of secure electronic communication. It's hard to overstate how serious this threat is. Shor's Algorithm, as it's now called, is a game changer.

Thankfully, encryption protocols not based on factorization do exist (although adopting them would be no small task), and several companies—including zy4, Isara, and others—are working on post-quantum encryption technology, creating ciphers that even quantum computers could not break. Moreover, the largest quantum computer currently available has 72 physical qubits, error rate unknown. A machine capable of implementing Shor's algorithm would require something on the order of *one million* fault-tolerant qubits. It seems, then, that quantum computers capable of cracking internet encryption are many years away, and we have many years to develop or adopt post-quantum cryptography. So perhaps the risk is overblown?

Not so fast. There are two problems.

First, even if it were true that the risk to *future* internet security is low given the continued development of post-quantum cryptography, this may not entirely avoid the issue. Often ignored in talk about the ability for quantum computers to decrypt existing data is that the decryption need not take place *today*. Rather, it is possible for someone to store currently encrypted data in order to decrypt it *later*, once quantum computers have become available. It is easy to imagine data—be it personal, corporate, or governmental—that we would not want decrypted by third parties at any time, even decades from now. This is thus a live issue, and one that highlights strikingly how disruptive quantum computers may be.

It gets worse. It may no longer even be the case that we have to wait for a quantum computer capable of implementing Shor's algorithm in order to break factorization-based cryptography. Just a few months ago, Zapata published a new algorithm called Variational Quantum Factoring. Complicated technical details aside, this algorithm can theoretically be implemented on what's called a NISQ (Noisy Intermediate Scale Quantum) computer of the kind that will be available in the near to medium term. NISQ computers are much smaller quantum computers—perhaps only 50-100 qubits in size—that can nevertheless perform tasks not feasible for classical computers. If it turns out that the Variational Quantum Factoring algorithm can be

deployed on a NISQ-era computer, then the threat to internet security might be much closer than we thought. If we wish to address this extremely serious risk, we must begin now.

Quantum mechanics itself may provide post-quantum encryption options capable of mitigating (though perhaps not eliminating) some of these worries. Quantum key distribution (QKD) is a secure communication method which leverages an important feature of quantum systems to implement a cryptographic protocol. Suppose two parties wish to communicate securely with one another. QKD enables them to produce a random secret key known only to them. This key is used to encrypt and decrypt messages sent between them. Now, a unique feature of quantum systems is that the very act of measuring that system disturbs it, and this disturbance can be detected. Someone trying to eavesdrop on the communication between the parties using the QKD must in some way measure the key. But then they will necessarily disturb the quantum system. So the parties using the key will know they are being listened in on, and can abort the communication. QKD is already in use, and China, for one, is basing much of its quantum efforts around QKD, while there are signs the US is falling behind in this area.

It might be that the threat to internet security, the very backbone of the global financial and social order, is one problem we can solve through the use of QKD and other post-quantum cryptographic methods. It might also be the case that we're already too late, as is suggested in a new report from the US National Academies of Sciences, Engineering, and Medicine. But even if we're not too late, it is far from a done deal: moving to post-quantum encryption protocols will undoubtedly require massive corporate and governmental support, and both, at this moment, appear to be sorely lacking.

There are, of course, many additional ethical questions raised by quantum computers. For example, only the largest corporations and governments will own quantum computers, at least for the foreseeable future (you won't be having a quantum computer in your home anytime soon, if ever). So what, then, of the relationship between citizens and their government when the government has massively increased power to surveil its citizens through the use of quantum-based decryption protocols, while enjoying *completely* secure quantum-based communications themselves? At first glance, this might seem no different from the current asymmetry that exists between citizen and government, but it represents a fundamental change, as current governments do not have access to completely secure communications, as various hacks and interceptions illustrate. What about the impact of quantum computers on issues of privacy generally? What of geopolitical relations if one country develops quantum computers before all others?

Major investors in US quantum technology are the military and national security sectors. How might quantum computers impact these areas? What military uses of quantum computers might develop? What is the future of war in a world with quantum computers?

Quantum computers will also interact and augment other emerging technologies, which already raise a host of important ethical issues. Consider artificial intelligence. The potential future development and use of AI is currently an area of great ethical concern. There is reason to think that quantum computers can help improve machine learning, and thus AI, so it raises the question of just how much more powerful a quantum computer-driven AI might be. What might a quantum computer-supported AI look like? Are we at all prepared to address the ethical issues this raises? Are we even aware of all the potential ethical issues?

And although we have focused mainly on the ethical implications of large-scale quantum computers, there are several technologies that are being developed as components of quantum computers that will likely be available soon, and which also raise ethical concerns.

Consider, for example, the incredibly precise and powerful sensors quantum computers require in order to function. These sensors, the development of which is perhaps only a few years away, have potential applications in several areas: they might improve how MRIs function, or weather prediction for airlines, or the detection of liquid explosives. But they may also be used in advanced surveillance technologies which threaten personal privacy. Even the technological precursors to a full-blown quantum computer raise important ethical concerns.

Quantum computers have the potential to do impressive things, all of which will have significant social, economic, or (geo)political implications, and all of which raise important, and difficult, ethical questions that need to be addressed. But perhaps the biggest threat posed by quantum computers is the unknown: the unexpected use, or misuse, of quantum computing; the game-changing algorithm we didn't see coming; the unfair distribution of essential quantum-based technologies or services; or the unexpected interaction of quantum computing with other emerging technologies. It is in many ways the *unanticipated* risks of quantum computers that underpin our profound sense of urgency to identify and address the ethical issues raised by them as soon as possible.

We have a tendency to wait too long to address potentially serious or even existential threats—climate change is a case in point— and quantum computing might be yet another example. But it needn't be, and it shouldn't be, as it would be wrong, both morally and prudentially, to wait to address the ethical issues quantum computing raises. We discuss this next.

The Moral Case

The moral argument for why we ought to begin identifying and addressing the many ethical issues raised by quantum computing immediately should be clear by now, but let's be more explicit and concrete.

The most appropriate and effective time to consider the ethical implications of a technology is when that technology is still in development. Call such technologies *emerging* technologies. An emerging technology is one that is currently in the design and development phase, is new and innovative, and is expected to have large socioeconomic impact, with the potential to affect significantly or transform one or more social, economic, or technological domains[2].

The benefit of considering the ethics of an emerging technology over that of a mature, entrenched technology is that it allows for early intervention in the design, innovation, or development phase of the technology. This provides a unique opportunity to mitigate potential harms to people, society, or shared human values, or to identify and plan for the fair distribution of future benefits that might result from the technology. It offers the possibility of intervention when choices can still be made, designs modified, and options about the social embedding of the new technology discussed[3].

Waiting to discuss the ethics of a technology until after the technology is released introduces significant ethical risks and can result in missed opportunities.

Although considering the ethical implications and full impact of an emerging technology is more difficult than when the technology is fully developed and deployed, it nevertheless offers the best chance to identify and mitigate or prevent harm before its deployment.

We need not look far for examples of companies that have acted improperly and caused harm through, or suffered—either financially, culturally, or reputationally—from a failure to consider the ethical implications of their technology early on. Two prominent examples come to mind.

Consider Facebook. One of the world's largest technology companies, it has recently found itself at the center of a number of ethical scandals, the most prominent of which has been the

[2] Philip Brey, "Ethics of Emerging Technologies" in *The Ethics of Technology: Methods and Approaches*. Ed. S. O. Hansson. London: Rowman & Littlefield International, Ltd., 2017. 175-190. *Kindle Edition.*
[3] *ibid.*

revelation that Russian propagandists, through the use of fake accounts and fake news information, manipulated the platform to undermine American democratic processes, including the 2016 Presidential election.

This is to say nothing of the Cambridge Analytica scandal, nor the deeply concerning internal leak of a sales pitch to advertisers, which apparently showed Facebook declaring that they could identify teenagers who felt "insecure", "worthless" and who "need a confidence boost". The pitch also claimed that they could pinpoint those who felt "defeated" or "nervous" or a "failure" through the analysis of posts and pictures, which is something Facebook apparently actively monitors. If a selling point of your product to potential advertisers is that it can be leveraged to exploit vulnerable teenagers, this suggests there is something deeply problematic about your business model and corporate culture. Facebook denies the claims and states that this was merely research and not put into practice.

Mark Zuckerberg, founder and CEO of Facebook, in response to this, stated in his congressional testimony that Facebook would begin implementing procedures for ensuring greater transparency into who has paid for an advertisement, and what other ads that person or entity is running on the platform. It has also vowed to crack down on fake news appearing on its platform more generally. But this all raises the question of why these steps are only now being taken, after the harm has been done?

Part of the explanation is likely that the founders saw only the potential upside of their technology, and thus saw themselves as doing good by developing and deploying it. There is a lesson here. This example highlights the danger of a naïve approach to dealing with the ethical implications of emerging technologies, and suggests that early, comprehensive consideration of the possible uses and misuses of a technology is the best means of ensuring that it not lead to harm.

In Facebook's case, manipulation of the underlying algorithm that determines which content gets shown to whom on Facebook was a possible misuse of the technology that arguably would've been identified with even a cursory analysis during development or early deployment. Of course, Facebook's basic business model, that of offering a free service, which allows them to gather personal data and then sell that data to advertisers and other organizations, raises several ethical concerns on its own, all of which ought to have been considered and addressed early in its development process. Their apparent failure to do so has now cost them both financially and reputationally, perhaps permanently.

Or consider Google. It was recently revealed that Google was working with the Pentagon on Project Maven, a project to develop artificial intelligence for use in military drones. Note first

that the potential military uses of this technology—autonomous drones as weapons—and the ethical issues this raises, should have been easily foreseen by Google. Note also that it was only after an outcry from both inside and outside the company (many of its employees strongly opposed participation in this project) that Google drafted a set of ethical principles to guide their work in the field of AI. But this seems backwards. Why were these principles drafted only *after* the company found itself embroiled in an ethically-charged situation? Surely, the time to draft guiding principles for the development and use of a technology is *before* that technology is developed or used.

The point is that we're currently at a stage in the development of quantum computers that affords us the opportunity to learn the important lessons of these failures. Indeed, we argue that undertaking the work now of anticipating and analyzing the possible ethical implications of quantum computing is not only prudentially wise (more on this below), but the morally right thing to do. We offer two arguments for this conclusion.

A widely recognized, though not uncontroversial, principle in the ethics of technology is the Precautionary Principle. There are several formulations of the Precautionary Principle, but the UN's World Commission on the Ethics of Scientific Knowledge and Technology (COMEST), noting that there is a strong link between ethical responsibility and the precautionary principle, defines it as follows:

> When human activities may lead to morally unacceptable harm that is scientifically plausible but uncertain, actions shall be taken to avoid or diminish that harm. Morally unacceptable harm refers to harm to humans or the environment that is
>
> - threatening to human life or health, or
> - serious and effectively irreversible, or
> - inequitable to present or future generations, or
> - imposed without adequate consideration of human rights of those affected.

We can further make explicit something implicit in this formulation of the Precautionary Principle: that if the precautionary principle applies to a technological intervention, then the *morally right action* is to take reasonable steps to avoid or diminish harm that might result from the technology.

There are good reasons for thinking the precautionary principle applies to quantum computers (and quantum technologies more generally): We detailed above some of the many possible uses of quantum computers and the possible harms that might result from them. These

harms, while uncertain, are plausible, and they are certainly capable of meeting the criteria for morally unacceptable harm found in the statement of the precautionary principle above, perhaps most especially being "serious and effectively irreversible" and "inequitable to present or future generations".

At this stage in the development of quantum computers, actions designed to "to avoid or diminish that harm" necessarily include the anticipatory work of identifying the possible harms that may result from the development of quantum computers, ethically analyzing those harms, and making design and development decisions to avoid or diminish those harms. Some of these harms, as we noted earlier, are likely to arise in the near to medium term. Thus, one way to begin to satisfy the Precautionary Principle with respect to quantum computers is to begin this anticipatory work now.

But perhaps you are skeptical about the Precautionary Principle, believing, as many critics do, that it represents an irrational fear of unproven risks. So it is important to note that our argument needn't trade on the Precautionary Principle itself, but can instead proceed via the more general moral prohibition against causing foreseeable but unjustified harm to others.

We all have a general moral duty to not cause foreseeable and unjustified harm to others. This plausibly entails a more specific duty on the part of developers of technology to prevent or mitigate the reasonably foreseeable and unjustified harms that might result from their technology. A plausible corollary of both duties is a further duty to undertake the epistemic work necessary to identify the potential harms that may result from your actions or technology: one cannot act blindly and then avoid responsibility by claiming that the harmful consequences of your actions or technology were unforeseen.

So, together these basic claims entail that we are under a duty to take reasonable steps to identify and address the foreseeable and unjustified harms that might result from the development and use of a technology. Specific to our case here, this duty requires that all those involved in, or responsible for, quantum computing and its effects, including engineers, CEOs, Boards of Directors, investors, and policymakers, undertake the work of attempting to discover the possible ways their technologies might harm people, or society more broadly, and then take effective and urgent steps to help mitigate that harm.

Here's an example of what we mean. We noted earlier that quantum computers may be capable of cracking those encryption protocols, such as RSA, which exploit a classical computer's difficulty with factoring large numbers sooner than we thought. We've thus identified a clear danger posed by quantum computers, one that could result in significant harm to individuals,

organizations, and global society. However, there currently exist encryption schemes that are not based on classical computer's difficulty with factoring large numbers, and quantum computers would thus not threaten them. There are also other encryption schemes in the works.

Given the availability, or imminent availability, of these alternative security protocols, and the identifiable and significant harm that could result by failing to replace existing protocols, there are strong moral reasons for moving to new encryption schemes now, despite the cost, before internet security is further threatened by development of quantum computing. Such a change would require a massive and coordinated effort, not only amongst corporations, but also governments and international organizations. This is an example of the kind of ethical analysis and consideration of the steps required to mitigate harm that we are recommending more generally.

Our conclusion, therefore, is that we must begin considering the ethical issues raised by quantum computers now, even if it is not possible to identify all such issues given our epistemic position. In fact, it is precisely because we cannot easily see all the possible risks that the technology might entail that we must no longer delay trying to determine what they are. Doing so is our best means to both satisfy the precautionary principle and fulfill our moral obligation to prevent foreseeable and significant harm to others. Waiting to discuss the ethics of these emerging technologies until they are developed and deployed is not only problematic, it's unethical.

The Prudential Case

We would hope that the ethical case alone is sufficiently persuasive, but perhaps you are equally (or more) concerned about the possibility of balancing innovation with the need to undertake the kind of anticipatory ethical analysis we suggest. And so, like many others, you believe it is best to postpone consideration of the ethical issues until later, after development and deployment of quantum computing technology. If history is any guide, this is a common mindset[4].

In fact, believing that the ethical concerns raised by a technology can be bracketed out until later so as to maximally free up intellectual and financial resources for innovation appears to be the default position in technology circles. But we think this is a mistake, and not merely for the moral reasons noted above. Even for purely prudential reasons it is a mistake. To see why, let's properly state this issue for what it is: the balancing of innovation against future reputation, legal, and financial risk.

[4] Consider: "When you see something that is technically sweet, you go ahead and do it and argue about what to do about it only after you've had your technical success. That is the way it was with the atomic bomb" That's the physicist Robert Oppenheimer, father of the atomic bomb, speaking.

Establishing cause and effect is often difficult, no less so when trying to determine how soft elements can impact organizational success. But we have good reason to believe (and many examples to suggest) that unethical behaviour can lead to significant risk or outright losses for an organization.

Consider again the example of Facebook. On July 25th, 2018, Facebook's CFO, in a second-quarter earnings call, noted that the scandals Facebook was facing would take significant toll on their revenue growth. The response from the market was swift and punishing: Facebook's share price dropped almost 23% in after-hours trading, representing a drop of some $150 billion off the company's market capitalization.

Source: https://qz.com/1340290/facebooks-market-cap-just-lost-150-billion-thanks-to-privacy-scandals/

That's a steep price to pay for avoidable unethical conduct. And it is certainly not the only example. Enron is no longer remembered for its innovation or business success, after all.

So, even if you reject the moral case presented above, there are strong prudential reasons for beginning the work we're recommending now. There is genuine reputational and financial risk for those firms either developing or intending to use quantum computers if their use is objectionable to the public or seen as harmful to society.

We recognize that treating quantum computers, or any emerging technology, as subject to the precautionary principle, or a more general ethical demand to identify and address potential future harms, can seem like a drag on innovation. We can acknowledge this, while recognizing that future ethical risks also represent a risk to long-term business success. This warrants taking the time to consider the issues now, and make the required changes to forestall or mitigate any negative effects. Doing so is just good business.

What should we be doing now?

How do we get ahead of this? How do we actually begin to fulfill our ethical obligation to identify and address the ethical issues raised by quantum computing now? What can we do, and how do we start? It can feel overwhelming, given the sheer scope of the task and the potentially massive impact of the technology. But there are some clear next steps.

First, the ethics of quantum computing needs to become an area of focus for academics, global professional organizations such as the United Nations, IEEE and NIST, policy makers, Boards of Directors, journalists, and concerned citizens. There is an understandable global focus on the ethics of AI right now, and this is, of course, absolutely essential. But quantum computers will almost certainly have at least as serious an impact as AI, let alone how quantum computers may affect the development of AI itself. We simply cannot afford to ignore the ethics of quantum computing—the risk is too great.

We also need to work collaboratively, across disciplines, countries, and professions. No individual person or organization is going to be able to effectively undertake this work alone: philosophers and others working on the ethics of technology need to consider the issue, help us navigate the nuances, and create the necessary ethical and conceptual frameworks; professional organizations, such as the IEEE, need to turn their attention to quantum computing, in much the same way that they have done for AI with their excellent Global Initiative on Ethics of Autonomous and Intelligent Systems; journalists need to popularize an understanding of quantum computing and the serious ethical issues they raise, while policymakers need to consider the kind of government regulation that might be required to govern this technology.

We need our best minds working on this issue now. Individual thinkers and organizations can of course begin work on this issue, and we encourage them to do so. But as noted, this also needs to be a global, collaborative effort. We believe there are two ways to help facilitate this work.

First, we are announcing the formation of the **QC Ethics Initiative**, an effort committed to helping identify and address the ethical issues raised by quantum computing and associated technologies. It is intended to provide a valuable forum for dialogue, sharing ideas, collaborative discussions, papers, and suggestions regarding the ethics of quantum computing. We hope that you will join us in our efforts.

We also have a unique, pressing, opportunity to make meaningful progress on these critical issues. Seizing it requires urgently bringing together the best and brightest minds in a multi-disciplinary conference focused on addressing the ethics of quantum computing in creative ways. To this end, The QC Ethics Initiative will host the first annual Ethics of Quantum Computing Conference in 2019. Our hope is to bring together a diverse group of academics, engineers, researchers, business leaders, policymakers, and others to make clear, substantive progress on the most urgent issues within ethics of quantum computing. It represents a singular opportunity to get ahead of the technology and help minimize harms in a way no other effort can.

As we begin to think more tactically, there are several practical tools and methodologies that can help us identify and address the ethical issues raised by quantum computing.

At root, all such approaches take the only real form that is possible for moving forward on these issues: try and identify the ways that the technology might develop, the ways it might be used, the ways it might be misused, and the ways it might interact with other technologies or embed itself in society. With this analysis in hand, subject the various uses to ethical analysis and work to identify ways to mitigate harms and fairly distribute benefits.

An example of such a process is what might be called *Ethical Foresight Analysis* (EFA). EFA is what's known as an *Anticipatory* approach to the ethics of emerging technology. Such approaches combine the techniques of foresight analysis, forecasting, scenario planning, Delphi panels, and other methods of technology assessment with ethical analysis[5]. Foresight analysis is a well-known technique for identifying and analyzing possible futures, including the possible future development of a technology and its uses. EFA adds in the identification of ethical issues at various stages of development and across various time horizons, and subjects them to analysis as well. It can then propose ways of addressing the identified ethical issues in advance[6].

There are four phases to such analyses: (1) Begin with foresight analysis; (2) Identify ethical issues discovered in that analysis; (3) evaluate those ethical issues; (4) Utilize the ethical evaluations to determine appropriate actions to address those issues. These latter actions might

[5] Brey
[6] Brey

include a design feedback stage, where the results of the ethical analysis are integrated into the design and development of the technology, or other morally justified risk management strategies[7].

During this process we can begin to consider various general ethical questions regarding quantum computers and associated quantum technologies:

a. How might they harm or benefit people in unexpected ways?
b. What might be unforeseen consequences or uses of quantum computing?
c. How might the benefits or burdens of quantum computing be unfairly or unequally distributed? What might constitute unjust access to the benefits of quantum computing?
d. What are the rights and duties attached to various roles within the field of quantum computing?

Smart, dedicated people coming together and undertaking some version of this process is likely the only way forward.

For those that might consider it a benefit, there's an additional one to beginning this work now: it's an opportunity for the quantum computing community to show that it can regulate itself sufficiently so as to avoid or minimize more formal governmental regulation. However, we would also suggest to those working in quantum computing that they need not oppose government regulation in principle: given the revolutionary potential of quantum computers, it may be unavoidable, and even preferable, to have clear government oversight of the technology. Moreover, the history of industries successfully self-regulating is not an encouraging one, and it may just be that the scope of the potential misuses and harms that may result from quantum computing are such that government regulation is not only required, but obligatory.

That said, one simple action that can help show a genuine commitment to self-regulation, **is to work towards a shared governing statement of ethical principles for the development and use of quantum computing.** We believe that this could be a natural focus area of the QC Ethics Initiative and one outcome of the Ethics of Quantum Computing conference. Such a statement can constitute a powerful moral touchstone for future work in the area, and signal to society that the quantum computing community is committed to the principled development and use of this technology. As we noted in our discussion of Google's new AI principles, the time to draft and

[7] Brey

adopt principles designed to guide the development and use of a technology is before that technology is developed and used.

Finally, it's important to note an additional benefit to beginning all of this work now, even if the analyses turn out to be wrong or incomplete in some ways. Call it the *dispositional benefit*. We recognize that we might not be capable of here, today, identifying all of the possible ethical risks of quantum computing. But the very process of discussing and undertaking the analyses *will* help us identify the risks going forward.

The process *itself* helps to establish a habit of continuously and consciously evaluating the potential ethical impact of quantum computing. As Aristotle (and others) have noted, virtue is a skill that we improve through habitual practice. A recognition that our analyses will not be perfect is no argument for waiting to begin them, as undertaking the process itself can help ensure that it continues and is ultimately successful.

The Road Ahead

We have a singular opportunity to get ahead of the significant, and often troubling, ethical issues raised by quantum computing and begin to address them now. This is true whether we're actively developing quantum computers ourselves, supporting that work through funding, are a political representative, academic, journalist, or just concerned citizen.

This is a technology that will change the world forever, and it's up to us to ensure it does so ethically. This work begins with commitment and dialogue, and we encourage everyone to join us in the hard work ahead. Let's move forward together.

Works Cited

Brey, Philip. "Ethics of Emerging Technologies" in *The Ethics of Technology: Methods and Approaches*. Ed. S. O. Hansson. London: Rowman & Littlefield International, Ltd., 2017. 175-190. *Kindle Edition.*