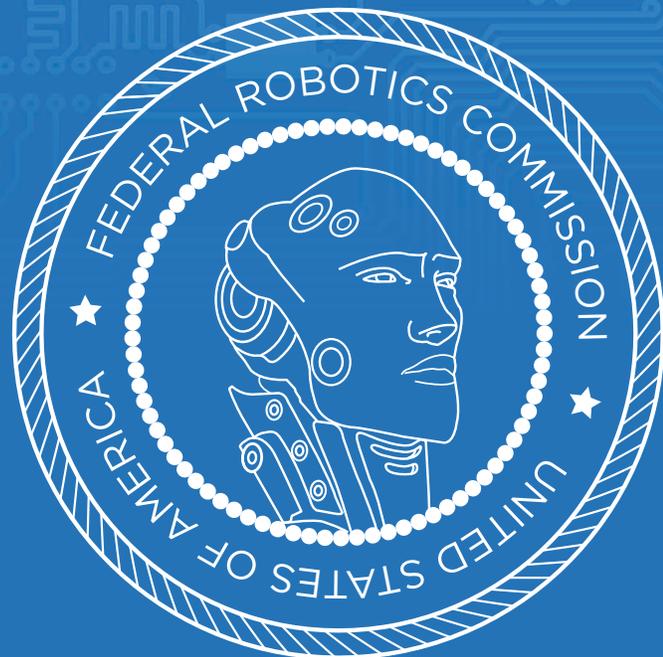


# The Case for a Federal Robotics Commission

By Ryan Calo | September 2014





**Ryan Calo** is an assistant professor at the University of Washington School of Law and a former research director at The Center for Internet and Society. A nationally recognized expert in law and emerging technology, Ryan's work has appeared in the *New York Times*, the *Wall Street Journal*, *NPR*, *Wired Magazine*, and other news outlets.

---

## Introduction

The U.S. Department of Transportation had a problem: Toyota customers were alleging that their vehicle had accelerated unexpectedly, causing death or injury. The National Highway Traffic Safety Administration (NHTSA) found some mechanical problems that may have accounted for the accidents—specifically, a design flaw that enabled accelerator pedals to become trapped by floor mats—but other experts suspected a software issue was to blame. Like most contemporary vehicles, Toyotas rely on computers to control many elements of the car. Congress was worried enough at the prospect of glitches in millions of vehicles that it directed the DOT to look for electronic causes.

NHTSA lacked the expertise to disentangle the complex set of interactions between software and hardware “under the hood.” The agency struggled over what to do until it hit upon an idea: let’s ask NASA. The National Aeronautics and Space Administration builds semi-autonomous systems and sends them to other planets; it has deep expertise in complex software and hardware. Indeed, NASA was able to clear Toyota’s software in a February 2011 report.<sup>1</sup> “We enlisted the best and brightest engineers to study Toyota’s electronics systems,” proudly stated U.S. Transportation Secretary Ray LaHood, “and the verdict is in. There is no electronic-based cause for unintended high-speed acceleration in Toyotas.”<sup>2</sup>

Under extraordinary circumstances, the best and brightest at NASA can take a break from repairing space stations or building Mars robots to take a look at the occasional Toyota. But this is not a sustainable strategy in the long run. Physical systems that sense, process, and act upon the world—robots, in other words<sup>3</sup>—are increasingly commonplace. Google, Tesla, and others contemplate widespread driverless cars using software far more complex than what runs in a 2010 sedan. Amazon would like to deliver packages to our homes using autonomous drones. Bill Gates predicts a robot in every home.<sup>4</sup> By many accounts, robotics and artificial intelligence are poised to become the next transformative technology of our time.

I have argued in a series of papers that robotics enables novel forms of human experience and, as such, challenges prevailing assumptions of law and policy.<sup>5</sup> My focus here is on a more specific question: whether robotics, collectively as a set of technologies, will or should occasion the establishment of a new federal agency to deal with the novel experiences and harms robotics enables.

New agencies do form from time to time. Although many of the household-name federal agencies have remained the same over the previous decades, there has also been considerable change. Agencies restructure, as we saw

with the formation of the Department of Homeland Security. New agencies, such as the Consumer Financial Protection Bureau, arise to address new or newly acute challenges posed by big events or changes in behavior.

Technology has repeatedly played a meaningful part in the formation of new agencies. For instance, the advent of radio made it possible to reach thousands of people at once with entertainment, news, and emergency information. The need to manage the impact of radio on society in turn led to the formation in 1926 of the Federal Radio Commission.<sup>6</sup> The FRC itself morphed into the Federal Communications Commission as forms of mass media proliferated and is today charged with a variety of tasks related to communications devices and networks.

The advent of the train required massive changes to national infrastructure, physically connected disparate communities, and consistently sparked, sometimes literally, harm to people and property. We formed the Federal Railroad Administration in response. This agency now lives within the U.S. Department of Transportation, though the DOT itself grew out of the ascendance of rail and later the highway. The introduction of the vaccine and the attendant need to organize massive outreach to Americans helped turn a modest U.S. Marine Hospital Service into the United States Centers for Disease Control and Prevention (CDC) and sowed the seeds for the Department of Health and Human Services.<sup>7</sup> And, of course, there would be no Federal Aviation Administration without the experiences and challenges of human flight.

---

*“Robotics and artificial intelligence are poised to become the next transformative technology of our time.”*

---

In this paper, I explore whether advances in robotics also call for a standalone body within the federal government. I tentatively conclude that the United States would benefit from an agency dedicated to the responsible integration of robotics technologies into American society. Robots, like radio or trains, make possible new human experiences and create distinct but related challenges that would benefit from being examined and treated together. They *do* require special expertise to understand and *may* require investment and coordination to thrive.

The institution I have in mind would not “regulate” robotics in the sense of fashioning rules regarding their use, at least not in any initial incarnation. Rather, the agency would advise on issues at all levels—state and federal, domestic and foreign, civil and criminal—that touch upon the unique aspects of robotics and artificial intelligence and the novel human experiences these technologies generate. The alternative, I fear, is that we will continue to address robotics policy questions piecemeal, perhaps indefinitely, with increasingly poor outcomes and slow accrual of knowledge. Meanwhile,

other nations that are investing more heavily in robotics and, specifically, in developing a legal and policy infrastructure for emerging technology, will leapfrog the U.S. in innovation for the first time since the creation of steam power.

This white paper proceeds as follows: The first section briefly describes some of the challenges robotics present, both specifically by technology, and in general, across technologies. The second describes what a federal robotics agency might look like in the near term. Section three addresses a handful of potential objections to the establishment of a federal robotics agency and a final section concludes. My hope for this white paper is to give readers a sense of the challenges ahead, diagnose our potentially worrisome trajectory here in the United States, and perhaps open the door to a conversation about what to do next.

## Law & Robotics

Robotics stands poised to transform our society. This set of technologies has seen massive investment by the military and industry, as well as sustained attention by the media and other social, cultural, and economic institutions. Law is already responding: several states have laws governing driverless cars. Other states have laws concerning the use of drones. In Virginia, there is a law that requires insurance to cover the costs of telerobotic care.<sup>8</sup>

The federal government is also dealing with robotics. There have been repeated hearings on drones and, recently, on high speed trading algorithms (market robots) and other topics on the Hill. Congress charged the Federal Aviation Administration with creating a plan to integrate drones into the national airspace by 2015.<sup>9</sup> The Food and Drug Administration approved, and is actively monitoring, robotic surgery. And the NHTSA, in addition to dealing with software glitches in manned vehicles, has looked extensively at the issue of driverless cars and even promulgated guidance.<sup>10</sup>

This activity is interesting and important, but hopelessly piecemeal: agencies, states, courts, and others are not in conversation with one another. Even the same government entities fail to draw links across similar technologies; drones come up little in discussions of driverless cars despite presenting similar issues of safety, privacy, and psychological unease.

Much is lost in this patchwork approach. Robotics and artificial intelligence produce a distinct set of challenges with considerable overlap—an insight that gets lost when you treat each robot separately. Specifically,

robotics combines, for the first time, the promiscuity of data with physical embodiment—robots are software that can touch you. For better or for worse, we have been very tolerant of the harms that come from interconnectivity and disruptive innovation—including privacy, security, and hate speech. We will have to strike a new balance when bones are on the line in addition to bits.

Robotics increasingly display emergent behavior, meaning behavior that is useful but cannot be anticipated in advance by operators. The value of these systems is that they accomplish a task that we did not realize was important, or they accomplish a known goal in a way that we did not realize was possible. Kiva Systems does not organize Amazon’s warehouses the way a human would, which is precisely why Amazon engaged and later purchased the company. Yet criminal, tort, and other types of law rely on human intent and foreseeability to apportion blame when things go wrong.

Take two real examples of software “bots” that exist today. The first, created by artist and programmer Darius Kazemi, buys random things on Amazon. Were Kazemi’s bot to purchase something legal in the jurisdiction of origin but unlawful where he lives, could he be prosecuted? Not under many statutes, which are formulated to require intent. Or consider the Twitter bot @RealHumanPraise. The brainchild of comedian Stephen Colbert, this account autonomously combines snippets from movie reviews from the website Rotten Tomatoes with the names of Fox News personalities. The prospect of accidental libel is hardly out of the question: The bot has, for instance, suggested that one personality got drunk on communion wine. But free speech principles require not only specific intent but “actual malice” when speaking about a public figure.<sup>11</sup>

The stakes are even higher when systems not only display emergent properties but also cause physical harm. Without carefully revisiting certain doctrines, we may see increasing numbers of victims without perpetrators, i.e., people hurt by robots but with no one to take the blame. Putting on one’s law and economics hat for a moment, this could lead to suboptimal activity levels (too much or too little) for helpful but potentially dangerous activities. What I mean is that, under current law, people who use robotic systems may not be held accountable for the harm those systems do, and hence may deploy them more than they should. Or, alternatively, people might never deploy potentially helpful emergent systems for fear of uncertain and boundless legal liability.

Finally, robots have a unique social meaning to people: more than any previous technology, they *feel* social to us. There is an extensive literature to support the claim that people are “hardwired” to react to anthropomorphic

---

*“Robots are software  
that can touch you.”*

---

technology such as robots as though a person were actually present. The tendency is so strong that soldiers have reportedly risked their own lives to “save” a military robot in the field.<sup>12</sup> The law impliedly separates things or “res” from agents and people in a variety of contexts; law and legal institutions will have to revisit this dichotomy in light of the blurring distinction between the two in contexts as diverse as punishment, damages, and the law of agency.

The upshot for policy is twofold. First, robotics presents a distinct set of related challenges. And second, the bodies that are dealing with these challenges have little or no expertise in them, and accrue new expertise at a snail’s pace. It is time start talking about whether a common institutional structure could help the law catch up, i.e., serve as a repository for expertise about a transformative technology of our time, helping lawmakers, jurists, the media, the public, and others prepare for the sea change that appears to be afoot.

---

*“Soldiers have reportedly risked their own lives to ‘save’ a military robot in the field.”*

---

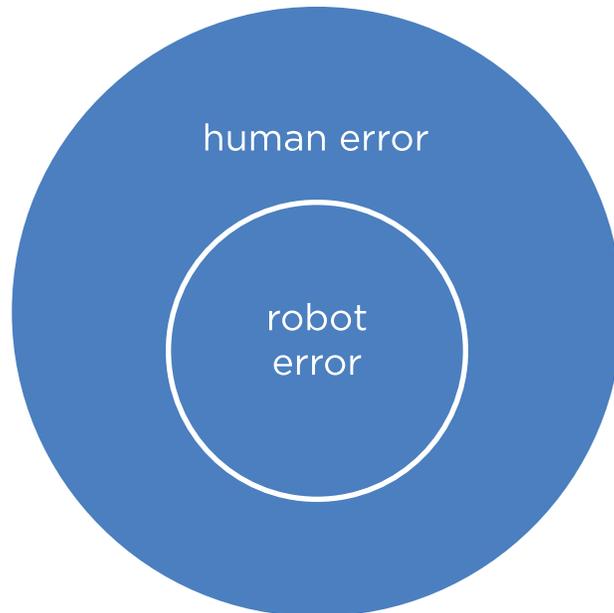
Arguably we have already seen a need for a federal robotics agency or its equivalent based on these three properties of robotics and AI organized to act upon the world. I opened with an example of Toyota and sudden acceleration. But this is just one of the many issues that embodiment, emergence, and social meaning have already raised. Some issues, like drones and driverless cars, are all over the news. Others, such as high-speed trading algorithms, is the subject of a best-selling book by Michael Lewis. Still others, however, concern high stake technologies you may have never heard of let alone experienced, in part because the problems they generate have yet to be resolved to stakeholder satisfaction and so they do not see the light of day.

### **Driverless cars**

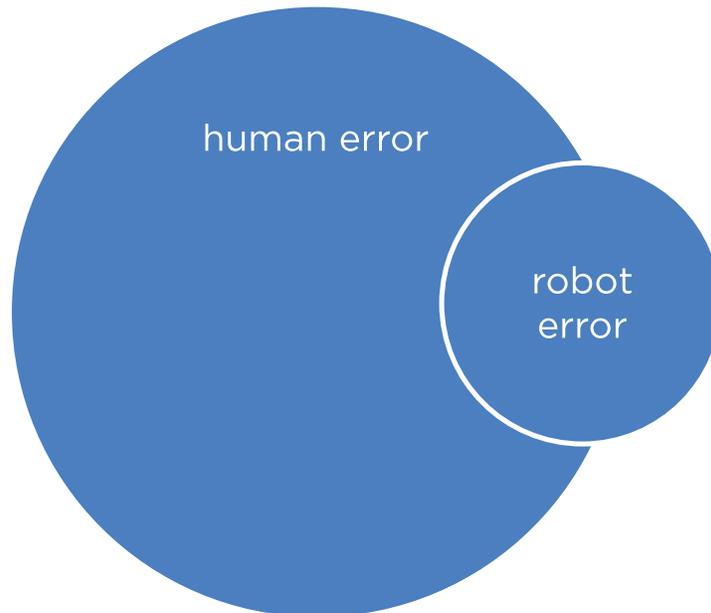
The state of Nevada passed the first driverless car law in 2011. It represented one of the first robot-specific laws in recent memory, as well as one of the first errors due to lack of expertise. Specifically, the Nevada legislature initially defined “autonomous vehicles” to refer to any substitution of artificial intelligence for a human operator. Various commentators pointed out that car functionality substitutes for people quite often, as when a crash avoidance system breaks to avoid an accident with a sudden obstacle. Nevada’s initial definition would have imposed hefty obligations on a variety of commercially available vehicles. The state had to repeal its new law and pass a new definition.<sup>13</sup>

One of the most significant challenges posed by driverless cars, however defined, is human reaction to a robot being in control. Human error accounts for an enormous percentage of driving fatalities, which number in the tens of thousands.<sup>14</sup> The promise of driverless cars is, in part, vastly to reduce these

accidents. In a “perfect,” post-driver world, the circle of fatalities caused by vehicles would simply shrink. The resulting diagram would look something like this:



But in reality, driverless cars are likely to *create* new kinds of accidents, even as they *dramatically reduce* accidents overall. Thus, the real diagram is more likely to look something like this:



The addition of even a tiny new area of liability could have outsized repercussions. A robot may always be better than a human driver at avoiding a shopping cart. And it may always be better at avoiding a stroller. But

---

*“Driverless cars are likely to create new kinds of accidents, even as they dramatically reduce accidents overall.”*

---

what happens when a robot confronts a shopping cart and a stroller at the same time? You or I would plow right into a shopping cart—or even a wall—to avoid hitting a stroller. A driverless car might not. The first headline, meanwhile, to read “Robot Car Kills Baby To Avoid Groceries” could end autonomous driving in the United States—and, ironically, drive fatalities back up. This possibility will be hard for laws, insurance, or video clips to inoculate against, requiring instead a subtle understanding of how the public perceives autonomous technologies in their midst.

### **Drones**

More immediate, because further along, is the case of domestic drones. Back in 2010, I predicted that drones would catalyze a national conversation around technology and privacy.<sup>15</sup> Unlike the Internet and other vehicles of data collection and process, I reasoned, it is easy for people to form a mental model of drone surveillance: there is a flying, inscrutable metal object, one you associated with the theatre of war, looking down upon you. Lawmakers and the public (but not yet the courts) have indeed reacted viscerally to the prospect of commercial and governmental use of drones domestically. Initially, however, the FAA through its officials attempted to distance itself from the problem. The agency observed that its main expertise is in safety, not civil liberties. It was only following tremendous outside pressure that the FAA began formally to consider the privacy impact of drones. The agency missed this issue—continues to miss it, to a degree—because it has no little to no experience with social meaning.<sup>16</sup>

---

*“There is a flying, inscrutable metal object, one you associated with the theatre of war, looking down upon you.”*

---

Law that confronts drones also tends to be underinclusive. There is little reason to target robots that can fly and take pictures over those that, say, climb the side of buildings or can be thrown into a building or over a crime scene. Arguably there is no good reason even to exclude birds with cameras attached to them—an old technique that is seeing something of a renaissance with cheap and light digital video. And yet “drone” laws almost inevitably limits themselves to “unmanned aircraft systems” as defined by the FAA,<sup>17</sup> which would leave law enforcement and private companies and individuals quite a few options for mobile surveillance.

Finally, FAA policy toward commercial drones has been roundly criticized for being arbitrary and non-transparent, including by an administrative law judge.<sup>18</sup> Here, again, the agency’s lack of experience with robotics—including what should or should not be characterized as a robot—may be playing a role. On the one hand, operators of small, low-flying drones argue that the FAA should not bother to regulate them because they do not raise issues any different than a remote control airplane flown by a hobbyist. On the other,

Amazon is concerned because the company would eventually like to deliver packages by drone autonomously and the recent FAA roadmap on drone interpretation seems to take autonomous navigation off of the table. These debates are ongoing before the agency itself and the courts.

### **Finance algorithms**

I mentioned the prospect of emergent behavior and the challenges it might pose for law and legal institutions. Michael Lewis' new book *Flash Boys* has raised awareness of the central role of one potential hazard, algorithmic trading, on Wall Street.<sup>19</sup> The SEC has been looking at the issue of high-speed trading, and the market volatility it can create, for years. The Commission seems no closer today to a solution than it was in the immediate wake of the 2010 "flash crash" where the market lost a significant percentage of its overall value in just a few minutes.

---

*"The law is still unsure how to handle the prospect of emergent behavior that ranges from benign, to useful, to potentially catastrophic."*

---

But high-speed trading could be the tip of the iceberg. Imagine, for instance, a programmer that designs software capable of predicting when a stock will make sudden gains in value—surely a gold mine for traders. This software is comprised of a learning algorithm capable of processing large volumes of information, current and historical, to find patterns. Were this software successful, traders would not necessarily understand how it worked. They might feed it data they felt was innocuous but that, in combination with publically available information, gave the algorithm what would otherwise be understood as forbidden insider information under the mosaic theory of insider trading. These traders or programmers might never be prosecuted, however, again because of the nature of criminal statutes.<sup>20</sup>

Legal scholars such as Tom Lin and Daria Roithmayer are looking at how law can adapt to the new reality of computer-driven investment.<sup>21</sup> Their insights and others in this space will inform not just high frequency trading, but any very fast and automated activity with real world repercussions. In the interim, the law is still unsure how to handle the prospect of emergent behavior that ranges from benign, to useful, to potentially catastrophic.

### **Cognitive radio**

A contemporary example you may not have heard of is the "cognitive radio," i.e., radios capable of "choosing" the frequency or power at which they will operate. Radios in general are locked down to one specific frequency so as not to interfere with other devices or networks. (A chief reason your cell phone has an FCC emblem on it is because the FCC is certifying non-

interference.) But communications bandwidth is scarce, and many believe it is not being used efficiently. Cognitive radio has the capability to modulate various parameters, including frequency and power, intelligently and in real time. These systems could operate on, for instance, emergency frequencies while they are not being used, or put out power just short of interfering with the next broadcaster a few miles away.

The upside of such technology, which is already in use in limited contexts today, is large and clear: suddenly more devices can work at the same time and more efficiently. The downside is equally large. Communications companies pay enormous sums to lease spectrum to provide services to consumers. In the case of emergency frequencies used by first responders, interference could be literally life threatening. Cognitive radios can malfunction and they can be hacked, for instance, by convincing a device it is in the mountains of Colorado instead of the city of San Francisco.<sup>22</sup> Thus, as the FCC has recognized, cognitive radios must have adequate security and there must be a mechanism by which to correct errors, i.e., where the radio uses a frequency or power it should not.

The FCC has been looking at cognitive radio for *ten years*; comments on how best to implement this technology remain open today. Current proposals include, for instance, a dual structure whereby meta algorithms identify and punish “bad” cognitive radios.<sup>23</sup> Technologists at UC Berkley, Microsoft, and elsewhere claim these solutions are feasible. But how does the FCC evaluate the potential, especially where incumbent providers or institutions such as the Navy tell the FCC that the risks of interference remain too high? It would be useful, arguably at least, for a single agency with deep expertise in emergent software phenomena help the SEC and FCC evaluate what to do about these, and many other, artificial intelligence problems.

### **Surgical robots**

I have outlined a few instances where technology policy lags behind or delays robotics and AI. We might be tempted to draw the lesson that agencies move too slowly in general. And yet, problems with robotics can as easily come from an agency moving too quickly. For instance, consider recent lawsuits involving surgical robotics. Some think the FDA moved too quickly to approve robotic surgery by analogizing it to laparoscopic surgery. The issues that arise, at least judging by lawsuits for medical malpractice, seem to stem from the differences between robotic and laparoscopic surgery. For instance, and setting aside allegations that improperly insulated wires burned some patients, robots can glitch. Glitches have not led to harm directly but rather require the surgical team to transition from robotic to manual and hence keep the patient under anesthetic longer.<sup>24</sup>

## FRC: A Thought Experiment

I have argued that present efforts to address robotics have been piecemeal in approach and, too often, unfortunate in outcome. Much of the problem turns on the lack of familiarity with robotics and the sorts of issues the mainstreaming of this technology occasions. The FAA does not know what to say to Amazon about delivering goods by drone, and so it says “no.” Even where one government body learns a lesson, the knowledge does not necessarily make its way to any other. Here I conduct a thought experiment: what if the United States were to address this problem, as it has addressed similar problems in the not-so-distant past, by creating a standalone entity—an agency with the purpose of fostering, learning about, and advising upon robotics and its impact on society?

Agencies look all kinds of ways; a Federal Robotics Commission would have to be configured in a manner appropriate to its task. Outside of the factory and military contexts, robotics is a fledgling industry. It should be approached as such. There are dangers, but nothing to suggest we need a series of specific rules about robots, let alone a default rule against their use in particular contexts (sometimes called the “precautionary principle”) as some commentators demand. Rather, we need a deep appreciation of the technology, of the relevant incentives of those who create and consume it, and of the unfolding and inevitable threat to privacy, labor, physical safety, and so on which robotics actually presents.

At least initially, then, a Federal Robotics Commission would be small and consist of a handful of engineers and others with backgrounds in mechanical and electrical engineering, computer science, and human-computer interaction, right alongside experts in law and policy. It would hardly be the first interdisciplinary agency: the FTC houses economists and technologists in addition to its many lawyers, for example. And, taking a page from NASA or the Defense Advanced Research Projects Agency (DARPA), the FRC should place conscious emphasis on getting the “best and brightest.” Such an emphasis, coupled with a decent salary and the undeniable cache of robotics in tech-savvy circles, could help populate the FRC with top talent otherwise likely to remain in industry or academia.<sup>25</sup>

What would an FRC do then? Here are some tasks for a properly envisioned Commission:

- Channel federal dollars into basic robotics research in an attempt to solve the still considerable technical challenges this technology presents.

---

*“The FAA does not know what to say to Amazon about delivering goods by drone, and so it says ‘no.’”*

---

- Attract highly skilled technologists who might be reticent to work for the government otherwise.
- Advise other federal agencies on matters having to do with robotics, including the DOT on driverless cars, the SEC on high speed trading, the FDA on robotic medical devices, the FCC on cognitive radios, the FAA on drones and, eventually, the Federal Trade Commission on increasingly sophisticated consumer products.
- Advise federal, state, and local lawmakers on robotics law and policy.
- Convene domestic and international stakeholders from industry, government, academia, and NGOs to discuss the impact of robotics and artificial intelligence on society.
- File amicus curiae (friend of the court) briefs, as many federal agencies do today, in matters involving highly complex interactions between software and hardware.

There is much room for disagreement over this list. One could question the organizational structure. The thought experiment is just that: an attempt to envision how the United States can be most competitive with respect to an emerging transformative technology. I address some deeper forms of skepticism in the next section.

## Objections

Today many people appreciate that robotics is a serious and meaningful technology. But suggesting that we need an entirely new agency to deal with it may strike even the robot enthusiast as overmuch. This section addresses some of the pushback—perhaps correct, and regardless healthy—that a radical thought experiment like an FRC might occasion.

### Do we really need another agency?

When I have outlined these ideas in public, reactions have varied, but criticism tended to take the following form: We need another federal agency? *Really?*

Agencies have their problems, of course. They can be inefficient and are subject to capture by those they regulate or other special interests. I have in these very pages criticized three agencies for their respective approaches or actions toward robotics. This question—whether agencies represent a good way to govern and, if so, what is the best design—is a worthwhile one. It is the subject of a robust and long-standing debate in administrative law that

cannot be reproduced here. But it has little to do with robotics. As discussed, we have agencies devoted to technologies already and it would be odd and anomalous to think we are done creating them.

A more specific version of the “really?” question asks whether we really want to “regulate” robotics at this early stage. I am very sympathetic to this point and have myself argued that we ought to remove roadblocks to innovation in robotics. I went so far as to argue that manufacturers of open robotics systems ought to be immunized for what users do with these platforms, product liability being a kind of “regulation” of business activities that emanates from the courts.<sup>26</sup>

Let me clarify again that I am using the term “regulate” rather broadly. And further, that there is nothing intrinsically anathematic between regulation and innovation. Copyright is regulation meant to promote creativity (arguably). Net neutrality is regulation meant to remove barriers to competition. Google—a poster child for innovation in business and, last I checked, a for-profit company—actively lobbied Nevada and other states to regulate driverless cars.<sup>27</sup> One assumes they did this to avoid uncertainty around the legality of their technology and with the hopes that other legislatures would instruct their state’s Department of Motor Vehicles to pass rules as well.

Note also that agencies vary tremendously in their structure and duties; the FTC, FDA, and SEC are *enforcement* agencies, for instance. Contrast them to, say, the Department of Commerce, DARPA, the Office of Management and Budget, or NASA itself. My claim is not that we need an enforcement agency for robotics—indeed, I believe it would highly undesirable to subject robotics and artificial intelligence to a general enforcement regime at such an early place in its life cycle. My claim is that we need a repository of expertise so that other agencies, as well as lawmakers and courts, do not make avoidable errors in their characterization and regulation of this technology.

A possible further response is that we have bodies capable of providing input already—for instance, the National Institute of Standards and Technology, the White House Office of Science and Technology Policy, or the Congressional Research Service. I would concede that these and other institutions could serve as repositories for knowledge about complex software and hardware. OSTP had a very serious roboticist—Vijay Kumar at University of Pennsylvania—serve as its “assistant director of robotics and cyber physical systems” for a time, and the Office’s mandate overlaps here and there with the possible FRC tasks I outline in the previous section.

Yet the diffusion of expertise across multiple existing agencies would make less and less sense over time. If robotics takes on the importance of,

---

*“We have agencies devoted to technologies already and it would be odd and anomalous to think we are done creating them.”*

---

for instance, cars, weather prediction, broadcast communications, or rail travel, we would want in place the kernel of an agency that could eventually coordinate and regulate the technology in earnest. Moreover, even in the short run, there would be oddness and discomfort in an institution that generally advises on a particular issue (e.g., standards), or to a particular constituency (e.g., Congress), suddenly acting as a general convener and broad advisor to all manner of institutions that have to grapple with robotics. Although I could see how existing institutions could manage in theory, in practice I believe we would be better off starting from scratch with a new mandate.<sup>28</sup>

### **How are robots different from computers?**

I will address one last critique briefly, inspired by the response science fiction writer Cory Doctorow had to my recent law review article on robotics. Writing for *The Guardian*, Doctorow expresses skepticism that there was any meaningful distinction at law or otherwise between robots and computers. As such, Doctorow does not see how the law could “regulate” robotics specifically, as opposed to computers and the networks that connect them.<sup>29</sup> “For the life of me,” writes Doctorow, “I can’t figure out a legal principle that would apply to the robot that wouldn’t be useful for the computer (and vice versa).”

In my view, the difference between a computer and a robot has largely to do with the latter’s embodiment. Robots do not just sense, process, and relay data. Robots are organized to act upon the world physically, or at least directly. This turns out to have strong repercussions at law, and to pose unique challenges to law and to legal institutions that computers and the Internet did not.

Consider, for example, how tort law handles glitches in personal computers or how law in general handles unlawful posts on a social network. If Word freezes and eats your important white paper, you may not sue Microsoft or Dell. This is due to a very specific set of legal principles such as the economic loss doctrine. But the economic loss doctrine, by its terms, is not available where a glitch causes physical harm. Similarly, courts have limited liability for insurers for computer or software glitches on the basis that information is not a “tangible” good covered by a general policy. A federal law, meanwhile, immunizes platforms such as Facebook for much of what users do there. It does so rather specifically by disallowing any legal actor from characterizing Facebook as the “publisher” of “content” that a user posts. This includes apps Facebook might run or sell. The same result would not likely obtain were someone to be hurt by a drone app purchased from a robot app store.<sup>30</sup>

---

*“We need a repository of expertise so that other agencies, as well as lawmakers and courts, do not make avoidable errors.”*

---

In any event, Doctorow's thesis does not necessarily cut against the idea of a Federal Robotics Commission. We might say that robots are just computers, but that computers today are more powerful and complex, and increasingly organized to act upon the world in a physical or direct manner without even the prospect of human intervention. Few in government, especially on the civilian side, understand this technology well. Accordingly, the latent need for a neutral government body with deep expertise on how to deal with cyber physical systems has become quite acute.

## Conclusion

I was recently at a robotics conference at the University of California, Berkeley and a journalist, who is a long-time student of robotics and one of its most assiduous chroniclers, made a remark to a colleague that struck me. He said that in recent years robotics has felt like a tidal wave, looming somewhere in the distance. But in recent months, that wave seems to have touched down upon land; keeping up with developments in robotics today is a frantic exercise in treading water.

---

*“Our government has a responsibility to be prepared for the changes robotics already begins to bring.”*

---

Our government has a responsibility to be prepared for the changes robotics already begins to bring. Being prepared means, at this stage, understanding the technology and the unique experiences robots allow. It means helping other institutions make sense of the problems the technology already creates. And it means removing hurdles to development of robotics which, if not addressed, could seriously compromise America's relevance in robotics and the development of its technology sector.

There are a number of ways our government could go about achieving these goals. I have explored one: the establishment of a federal robotics agency. We have in the past formed formal institutions around specific technologies, for the obvious reason that understanding a technology or set of technologies requires a dedicated staff, and because it can be more efficient to coordinate oversight of a technology centrally. I do not argue we should go so far as to put into place, today, a full-fledged enforcement body capable of regulating anything that touches robotics. That would be deeply inadvisable. Rather, I believe on balance that we should consider creating an institutional repository of expertise around robotics as well as a formal mechanism to promote robotics and artificial intelligence as a research agenda and industry.

The time to think through the best legal and policy infrastructure for robotics is right now. Early decisions in the lifecycle of the Internet, such as the decision to apply the First Amendment there and to immunize platforms

for what users do, allowed that technology to thrive. We were also able to be “hands off” about the Internet to a degree that will not be possible with robotics and systems like it that are organized not merely to relay information but to affect the world physically or directly. Decisions we make today about robotics and artificial intelligence will affect the trajectory of this technology and of our society. Please think of this white paper, if you do, as a call to be thoughtful, knowledgeable, and deliberate in our dealings with this emerging technology. Your thoughts and comments on how best to do so are warmly welcome. Thank you for reading and to the Brookings Institution for this opportunity to address you.

## Acknowledgements

**I would like to thank Wells Bennett and Benjamin Wittes for helpful comments and Woodrow Hartzog and Olivier Sylvain for inspiring this thought experiment.**

---

1. See National Highway Traffic Safety Administration, *U.S. Department of Transportation Releases Results from NHTSA-NASA Study of Unintended Acceleration in Toyota Vehicles*, DOT 16-11 (Feb. 8, 2011).

2. *Id.*

3. While there is no consensus definition of “robot,” many find it useful to refer to the sense-think-act paradigm to delineate robotics from previous or constituent technologies. Thus, a laptop with a camera senses and processes (“thinks”) but does not “act.” Whereas a remote control car with a camera acts and senses but does not process. Robots represent technology that senses the world, processes and integrates what it senses, and then acts upon the world in some way. Of course, each of these factors exists on a spectrum. For a longer discussion of robotics, please see Ryan Calo, *Robotics and the Lessons of Cyberlaw*, 103 CALIFORNIA LAW REVIEW (forthcoming 2015).

4. Bill Gates, *A Robot in Every Home*, SCIENTIFIC AMERICAN 58 (Jan. 2007).

5. E.g., Ryan Calo, *Robotics and the Lessons of Cyberlaw*, 103 CALIFORNIA LAW REVIEW (forthcoming 2015); Ryan Calo, *Open Robotics*, 70 MARYLAND LAW REVIEW 571 (2011).

6. See TIM WU, THE MASTER SWITCH 82-84 (2010).

7. See MICHAEL WILLRICH, Pox 77-81 (2011).

8. Va. Code § 38.2-3418.16.

9. FAA Modernization and Reform Act of 2012, § 332(a)(3).

10. National Highway Traffic Safety Administration, *U.S. Department of Transportation*

*Releases Policy on Automated Vehicle Development*, NHTSA 14-13 (May 30, 2013).

11. *Gertz v. Robert Welch, Inc.*, 418 U.S. 323, 345 (1972).

12. See P.W. SINGER, *WIRED FOR WAR: THE ROBOTICS REVOLUTION AND CONFLICT IN THE TWENTY-FIRST CENTURY* 337-43 (2009).

13. NRS § 482A.020 (repealed 2013); NRS § 482A.30 (2011).

14. See John Villasenor, *Products Liability and Driverless Cars: Issues and Guiding Principles for Legislation*, Brookings Institution White Paper (April 24, 2014).

15. Ryan Calo, *The Drone as Privacy Catalyst*, 64 *STANFORD LAW REVIEW ONLINE* 29 (2011).

16. Obviously planes themselves have cultural and societal meaning, but we do not—or no longer—react to them viscerally, let alone attribute agency to them when they appear to be piloted.

17. *Id.*

18. *FAA v. Raphael Pirker*, NTSB Docket CP-217, March 6, 2014.

19. MICHAEL LEWIS, *FLASH BOYS: A WALL STREET REVOLT* (2014).

20. SEC rules require that the trader be “aware” of the non-public or otherwise problematic nature of the information she trades upon.

21. *E.g.*, Tom C.W. Lin, *The New Investor*, 60 *UCLA LAW REVIEW* 678 (2013).

22. *Cf.* T. Charles Clancy & Nathan Goergen, *Security in Cognitive Radio Networks: Threats and Mitigation*, Proceedings of the Third International Conference on Cognitive Radio Oriented Wireless Network Communications (CrownCom) (May 15-17, 2008).

23. *E.g.*, Kristen Ann Woyach et al., *Crime and Punishment for Cognitive Radios*, Proceedings of Communications, Control, and Computing, 46<sup>th</sup> Annual Allerton Conference (Sep. 23-26, 2008).

24. See Sulbha Rai, *Robotic Surgery and Law in USA—A Critique*, available online at <http://ssrn.com/abstract=2425046> (June 1, 2013).

25. Ask yourself how cool it would be to tell people you work at the Federal Robotics Commission.

26. Ryan Calo, *Open Robotics*, 70 *MARYLAND LAW REVIEW* 571 (2011). See also Ryan Calo, *The Need to Be Open: U.S. Laws are Killing the Future of Robotics*, MASHABLE (Jan. 1, 2014).

27. John Markoff, *Google Lobbies Nevada to Allow Self-Driving Cars*, *NEW YORK TIMES* A18 (May 10, 2011).

28. One particularly helpful agency for robotics might have been the Office of Technology Assessment. Unfortunately, this agency was discontinued by Congress. A version of the thought experiment resuscitates OTA and charges it with emerging technology policy, a plan discussed in detail by Jathan Sadowski in forthcoming work. Jathan Sadowski, *Office of Technology Assessment: History, Implementation, and Participatory Critique*, manuscript on file with author.

29. Cory Doctorow, *Why It Is Not Possible to Regulate Robots*, THEGUARDIAN.COM (Apr. 2, 2014).

30. For two examples, see <http://ardrone2.parrot.com/usa/apps/> and <http://www.robotappstore.com/Pages/Model.aspx?Model=DRONE>. Disclosure: I have served on the advisory committee of the latter.

