Taking Futures Seriously: Forecasting as Method in Robotics Law and Policy

Stephanie Ballard and Ryan Calo
University of Washington Tech Policy Lab
Seattle, WA, USA

A central challenge in setting law and policy around emerging technology is predicting how technology will evolve. In failing to consider the future of technology, we are often left with laws and policies that fall short of our technological reality.

The 99th United States Congress had no experience with the commercial internet, for example, leaving it ill-equipped to envision the future of communications technology or to understand how widespread access to citizen information would need to be regulated. Indeed the laws Congress passed in 1986, which still govern electronic communications to this day, made assumptions about the nature of remote computing that have not obtained for decades.

In the 1990's, the Department of Transportation envisioned that driverless cars would ride upon "smart" roads, similar to a trolley. The DOT issued extensive guidance along these lines—proposing, for instance, heavy investment in infrastructure. Today, autonomous vehicles are on the roads in several states, but they do not run on tracks. Instead, they are self-contained robots capable of sensing and responding to ordinary environments.

The difficulty in predicting the trajectory of new technology can give rise to a number of unfortunate consequences regarding policy. One is staleness—outdated rules, such as the 1986 law governing electronic communications privacy, that nevertheless persist through inertia or entrenched interest. A second is waste—the over-investment in a particular instantiation of a technology, such as the investment in "smart" infrastructure by the DOT in the 1990's. Yet another is policy paralysis—a phenomenon celebrated by libertarians but bemoaned by many as abdicating governmental responsibility to channel technology in the public interest.

Meanwhile, as David Collingridge famously observes, the lack of information that causes policymakers to hesitate in governing emerging technology eventually yields to difficulties in controlling that technology (Collingridge, 1980). By the time we understand what would have been the best way to regulate it, the technology at issue has become path dependent and protected by vested political and economic interests.

That's the bad news. The good news is that methods exist to help address the thorny problem of envisioning the future. Over the years, scholars and corporations have developed numerous qualitative and quantitative techniques by which to explore possible futures and plan for uncertainty. Known variously as "envisioning," "forecasting," and "future studies," these

methods are credited with assisting institutions from Shell Oil to the National Security Agency in navigating potential crises and otherwise making profitable or wise decisions. Despite their maturity and success, however, these techniques remain almost entirely unremarked within law and technology theory or practice.

The thesis of this paper, co-authored by an information scientist and a legal scholar, is that robotics law and policy as a field would benefit from exposure to rigorous methods of envisioning.

Our argument proceeds as follows: The first section introduces the reader to the field of future studies through an efficient review of the history of the field. The academic study of forecasting emerged around the 1960's and professional foresight has been applied in corporations, non-profit organizations, and governments ever since. Although typically practiced at the executive level to guide long-term strategic planning of an organization, the techniques hold the prospect to inform individual policymaking as well.

The second section isolates three specific methods—design fiction, scenario planning, and future wheels—and applies them to the case study of robotic delivery. We selected these methods for their feasibility, concreteness, and widespread deployment. Design fiction is an increasingly popular mode of envisioning the development of technology and its social impacts through narrative iteration. Scenario planning, pioneered by Herman Kahn at the RAND Corporation and further developed by economists Pierre Wack and Edward Newton at Shell Oil, consists of a planning technique by which managers can confront and assess the plausibility of various local and global developments. The futures wheel is a technique developed by social scientist and civic leader Jerome Glenn to explore the ramifications of emerging technology.

We selected robotic delivery as a case study because of its many potential configurations (e.g., drone or land-based robots) and its significant but still-unfolding legal context. In addition to guidance generated by the Federal Aviation Administration around the possible uses of drones to deliver packages, five states (Ohio, Florida, Wisconsin, Idaho, and Virginia) have already passed laws concerning sidewalk robots. Despite the application's plausibility, there is a dearth of discussion of robotic delivery in the legal literature.

In the final section, we leverage the insights from sections one and two to critique existing robotic delivery policies (or their absence). This section develops a case for broader application of forecasting techniques in and beyond robotics law and policy. We are mindful, of course, that the future is never certain. Literal "future-proofing" is a fool's errand. Nevertheless, a systematic approach to the exercise of envisioning has the potential to significantly improve policymaking across robotics and other domains.

Forecasting: A Brief Introduction

"It's tough to make predictions, especially about the future." – Yogi Berra

Forecasting, futures studies, and foresight are among the various terms used to describe a collection of methods and tools for systematically exploring a range of possible futures. While there are many forces outside of human control, human action or inaction shapes our collective future. We design technology, enforce policy, trade in markets, and manipulate natural resources every day.

These decisions, and their unintended consequences, make possible certain futures while foreclosing others. For example, in 1979 China introduced the one "one child policy" to slow population growth. In shaping that policy, they attempted to prevent a future where a growing Chinese population put stress on social, economic, and environmental resources. The unintended consequences of this policy are now known: In an attempt to prevent a future of resource constraint and overpopulation, China created a circumstance in which the long term disparity in sex ratio at birth has resulted in 32 million more males than would occur naturally and the "four-two-one problem," among others. Understanding the way technological, social, economic, and environmental forces interact with each other, the way unintended consequences manifest, and how best to mitigate them is the work of futures studies. In thinking rigorously about the future, we increase our ability to anticipate change and react to it in a strategic, foreword thinking way.

Since before the oracles at Delphi, humans have been concerned with predicting future events. The formal discipline of future studies, however, emerged at the tail end of World War II. This period of futures studies is characterized by the desire for methodological prediction and control, especially in the face of the growing tensions of the Cold War. During the 1950's and 60's, the RAND Corporation in particular devoted substantial resources to developing tools and theories for prediction. These early approaches tended to reify the idea that there is one possible future, which we can come to know through a general theory of prediction, precise models, and accurate data. This view of the single, likeliest future dominated future studies in the US and in the Soviet Union.

While futures work in the US stressed prediction, European intellectuals dealing with the aftermath of World War II identified the future with moral imperative. Hannah Arendt and Lewis Mumford, among others, viewed the quest for prediction as an attempt to colonize the next age (Andersson, 2018). As European future studies emerged, scholars emphasized the plurality of futures, that is, that many possible futures exist, which we can envision and make real together. European futurists were committed to democratizing futures studies by sharing futures methods and concepts widely rather than deploying them privately for strategic advantage. The ontological tension between prediction for means of control and a democratic envisioning of future possibilities is still reflected in the practice of futures studies today.

In the wake of the Cold War, futures work saw uptake outside of the military industrial complex, most notably with the Shell Oil Company. Futurists themselves organized conferences and professional organizations, eventually gaining formal accreditation as a part of the academy. During this time the field experienced a period of professionalization and an increase in consultancy-based practice.

Today, futures methods are applied in industry, government agencies, and across several academic fields of research. In industry, corporations like PricewaterhouseCoopers, Google, Microsoft, and others have used speculative futures methods to anticipate developments in their respective domains and to guide technological paths. For example, in 2017, PricewaterhouseCoopers published a report titled “Using science fiction to explore business innovation” in which they claim that science fiction has consistently predicted future technologies and social phenomena, the use of science fiction can spark innovation, and science fiction allows for the exploration of new products without constraints of reality hampering creativity (Gibbs, 2017). In 2015, Microsoft released Future Visions, a collection of science fiction based on projects at Microsoft Research (Bear, et al., 2015). And in 2016, Google created a piece of design fiction which explores the possibility of total data collection for the purposes of nudging users toward their goals (Savov, 2018). In addition to these speculative examples of futures work from industry, corporations regularly engage in strategic planning using futures methods. Indeed, Shell credits its use of scenario planning for its successful navigation of the 1973 oil crisis.

In government agencies, futures methods are primarily used for long term visioning and strategic planning. For example, every four years the National Intelligence Council publishes a report titled Global Trends: Paradox of Progress. Currently in the sixth installment, the document reports on global scenarios with specific implications for US foreign and economic policy generated using multinational, participatory scenario planning (National Intelligence Council, 2017). Across the US Federal government, strategic foresight is most established within the defense and intelligence agencies but is more tenuous within civilian agencies. While strategic foresight is relatively common across the US Federal Government, these efforts vary widely and are not yet institutionalized (Greenblott, et al., 2018).

In academia, futures methods have roots in several disciplines. Futures studies, design, and human computer interaction (HCI) among others all utilize, develop, and innovate futures methods and produce futures work. Within HCI, techniques like envisioning have long played a role in the development of technology and, more recently, speculative methods like design fiction have become increasingly popular as a tool for exploring the future of technology. The field of speculative design concerns itself with opening up the future to critique through design practices. Dunne and Raby, pioneers in the field of speculative design, state: “This is the bit we are interested in. Not in trying to predict the future but in using design to open up all sorts of possibilities that can be discussed, debated, and used to collectively define a preferable future for a given group of people…” (pp 6) (Dunne & Raby, 2013).

Across these disciplines, futures studies concerns itself with understanding possible futures and, when practiced in applied settings, how to respond strategically to those possible futures. Academic futures work is frequently published in Technological Forecasting and Social Change, Futures, and the World Futures Review. Examples of futures work in the field of HCI frequently appear at the ACM Conference on Human Factors in Computing Systems (CHI) and Designing Interactive Systems (DIS).

Futures methods and approaches vary substantially, from normative to exploratory, quantitative to qualitative, short-term (3-5 years) to long term (50+ years), and concrete to highly speculative. Different methods have different advantages and utilities, and best practice often requires the application of multiple methods to the same problem space.

Quantitative futures methods are good for making short-term predictions in a narrow application. Though they can be quite precise and even accurate in their predictions, quantitative methods rely on historical data and can only account for future events that are likely to continue along an existing trajectory. Unlike qualitative methods, quantitative methods cannot account for “wild card” events that are unlikely or unforeseen but fundamentally change the course of the future. Our modern social media landscape can be considered one such wild card. The architects of the early internet did not anticipate the future of social media sites like Facebook or even the widespread commercial use of email and digital storage, which have of course fundamentally changed the way we interact online and in person.

One example of a quantitative futures method we interact with every day is statistical modeling. Algorithmic advertising is delivered to us from a statistical model of our shopping behavior, social media newsfeeds are organized by statistical models of our attention, our car insurance rates are set by statistical models of our driving risk, and the list goes on. In statistical modeling, data about the past is used to make predictions about the future. This method relies on the assumptions that a) the historical data models the phenomena we are trying to predict, b) the data completely and accurately describes the phenomena we are trying to predict, and c) the phenomena we are trying to predict will arise in the same way as in the past.

Trend Impact Analysis (TIA) is a mixed-method that attempts to address the limitations of other quantitative futures methods by incorporating the potential effects of future trends into a forecast made using historical data. First, the quantitative forecast is made using historical data. Next, an expert determines which future events might affect the forecast and how, then adjusts the forecast to account for these potential effects. The TIA modified forecast merges the forecast based on historical data with the future trends identified by experts to create a forecast more resilient to future changes.

Qualitative futures methods are appropriate for moving beyond the probable future to understand the long term effects of policy or the possibilities of technological development. In this paper, we apply three qualitative methods–design fiction, scenario planning, and futures wheel–but many others exist. The Delphi method is an anonymous, iterative survey of experts used to reach consensus on a forecast. The method begins by asking a group of experts to

make a prediction (e.g., in what year will we find a cure for Alzheimer's). In the second round, the results from the first survey are anonymously returned to the experts. Experts holding opinions at the extremes of the range are asked to provide their reasoning and all experts make a second prediction. The Delphi method is completed in four rounds through which experts attempt to reach consensus through iterative presentation of results, controlled argument, and prediction.

There are a great many futures methods, such that detailing them exhaustively must fall outside the scope of this paper.[1] All futures methods have their strengths and weaknesses. Futures work should be evaluated by its ability to help decision makers in the present, not by the accuracy of the forecasts.

If not to predict the future, why then do we engage in futures work? As a practice, engaging in futures work enhances our anticipatory reflexes and makes us more adept at responding to change. In practice, rigorously thinking about the future helps us identify potential pitfalls, unintended consequences, and assumptions regarding the phenomena under study. In support of decision making, futures work focuses on delivering more knowledge into a decision making process. Ultimately, the goal of futures work is to transform "unknown unknowns" into "known unknowns" to build out a more complete picture of technology's potential trajectory and social impacts. In the next section, we demonstrate the value of engaging in futures work through a case study.

The Case of Robotic Delivery

Future studies has an involved history and a set of well-developed methods applied by renowned institutions across many domains. The thesis of this paper is that rigorous engagement with existing futures methods would improve technology policy as well. To evidence this claim, we have chosen the case study of the law and policy around robotic delivery. We made this choice for several reasons. The first is the significant stakes: robotic delivery is already an emerging industry backed by considerable resources. The second is the variety of technological form factors. Initial forays into robotic delivery take many forms, including delivery by air, street, and sidewalk, with lesser or greater reliance on automation. The third reason is the availability of regulatory benchmarks. Several states and municipalities have already promulgated rules for robotic delivery—rules we believe would have benefited from engagement with futures methods.

Household name companies have invested tens of millions of dollars in robotic delivery. Tech giants Amazon and Google have widely reported drone delivery programs named Prime Air and Wing, respectively. According to the Wall Street Journal, Uber intends to deliver food by drone as early as 2021. But many non-tech companies—from Mercedes to UPS to DHL to Airbus to Dominos to 7-Eleven—have entered this sector as well.

[1] For a more complete treatment of futures methods, see Glenn and Gordon's excellent resource *Futures Resarch Methodology Version 3.0*).

Often these companies will partner with drone delivery startups. These smaller companies include Flirtey, CyPhy, Workhouse Horsefly, Matternet, and Dialexa.

In addition to aerial delivery, a variety of well-funded startups have begun to offer robotic ground transport of consumer goods. Best known perhaps is London-based Starship, which has raised in excess of 17 million dollars and operates in at least six U.S. states. Starship and its competitors—including Kiwi, Dispatch, Marble, and Robby—navigate urban environments and college campuses by sidewalk, relaying on a varying degrees of automation and remote operation.[2] The startup Nuro, which eschews sidewalk robots in favor of miniature driver- and passenger-less cars that navigate city streets, has raised a jaw-dropping 1 billion dollars, mostly from the Japanese tech giant Softbank.

The proliferation of robotic delivery is not lost on federal, state, or local government.[3] The Federal Aviation Administration has certified a handful of drone delivery pilot programs in recent years, generally in rural areas. There are active drone delivery pilot programs in North Dakota, North Carolina, Memphis, Nevada, California, Virginia and potentially other states. At least three cities—Reno, San Diego, and Herndon, Virginia— are contemplating urban delivery at the time of writing. A majority of the programs involve the delivery of medicine and medical equipment exclusively.

The FAA has limited these programs, in part by rejecting the vast majority of applicants. In its May 2018 drone delivery initiative, the FAA accepted under 7% of the 150 proposals. Where the agency has given leave, it has imposed strict constraints on testing, such as limiting testing to daylight operation. Notably for our purposes, the FAA imposes a weight limit of 55 pounds for its small aircraft pilot program and requires that flights "be conducted within visual line of sight and not from a moving vehicle or aircraft."

As many have noted, the FAA's precautionary approach has led companies to test drone delivery outside of the United States where restrictions are fewer. There are signs, however, that the regulatory environment for drone delivery in the United States is set to shift again in 2019. In the most recent FAA appropriation bill (the statute allocating funding to the agency) in October of 2018, Congress specifically directed the FAA to "update existing regulations authorizing carriage of property" and to do so "not later than 1 year after the date of enactment."

While states have limited jurisdiction over the airspace,[4] multiple states and municipalities have enacted laws concerning ground-based delivery. The standard

[2] While these companies generally rely upon land based carts traveling on wheels, at least one company—Unsupervised.AI—has experimented with four-legged robots.

[3] This proliferation is largely lost on the academy, however, which has tended to focus on driverless cars and surveillance drones. But see Mason Marks, Delivery Robots and the Influence of Warehouse Logic on Public Spaces, We Robot 2019. ☺

[4] State statutes tend to focus on state employee conduct, such as police use of drones. At least two federal bills have been introduced that would grant additional authority over drone delivery to states—the Drone Federalism Act, S. 1272 and the Drone Innovation Act, H.R. 2930—but neither has moved out of

statutory term for a ground delivery robot is "personal delivery devices" or PDDs. Despite the generality of the term, these statutes envision highly specific technologies and practices. Thus, each bill of the seven state statutes we identified sets a weight limit between 50 and 200 pounds (with a mean of 104 and a mode of 80).[5] Each statute imposes a speed limit of 7 or 10 miles per hour, and each statute limits ground operation to sidewalks and crosswalks. Other common requirements include functional brakes, liability insurance, identification, a remote operator, and the capacity to comply with local ordinances while in autonomous mode.

In addition to state statutes, a number of cities and counties have allowed pilot programs under similar conditions. Washington, D.C., Redwood, San Carlos, Menlo Park, Sunnyvale, Mountain View, San Francisco, Palo Alto, and Los Angeles, California, Austin and Arlington, Texas, and Snohomish County, Washington have all permitted short PDD pilots by resolution. Finally, individual universities—among them George Mason and UC Berkeley—have hosted pilot programs within campus grounds.

The avowed goal of these statutes and ordinances is to clarify that PDDs are permitted within the jurisdiction, subject to certain obligations and restrictions. But not every state or local government to consider robotic delivery has been open to it. As an initial matter, the majority of states and cities are completely silent as to PDDs. According to the CEO of the largest sidewalk delivery company, Starship, this silence operates as a de facto ban because the robots are likely to be classified as vehicles and thereby prohibited from using the sidewalk.

Some jurisdictions have imposed categorical limits on robotics delivery systems. Thus, for example, California's Bureau of Cannabis Control specifically bans the use of robots to deliver marijuana. According to the regulation: "Transportation by means of aircraft, watercraft, drone, rail, human powered vehicle, or unmanned vehicle is prohibited." San Francisco voted to limit PPDs to certain low traffic zones, banning robotic delivery from most city streets.

In sum, robotic delivery in the United States constitutes a robust, if yet emerging, industry with household name companies and many millions of funding behind it. Robotic delivery today takes almost exclusively two forms: aerial drone and sidewalk drone delivery. The former is highly restricted by a single agency, the FAA, which requires licensure, weight limits, and operation during daylight hours by a stationary person within line of sight of drone. The latter is alternatively banned, whether by implication or expressly, or else regulated in highly specific and recurrent ways.

---

committee. See also Margot E. Kaminski, Drone Federalism: Civilian Drones and the Things They Carried, 4 California Law Review Circuit 57 (2013) (arguing that states may be better positioned to balance drone innovation with safety and civil liberty concerns through experimentation).

[5] The state of Washington is contemplating a similar bill as of this writing.

Application of Future Studies to Robotic Delivery

We now turn to the application of our three methods – design fiction, scenario planning, and futures wheel – to the case study of robotic delivery. We chose these methods because they are widely used in futures projects, easily applied in a low resource setting, and cover a range of functions within futures work (e.g., blue-sky envisioning, strategic planning, etc.).

In envisioning the future of technology policy, we must account for both the potential futures of the technology as well as the possibilities for future policy. In our case study, we performed the methods in a specific order so as to first envision the future of drone technology using the most open-ended, speculative method – *design fiction*. Using design fiction we identified several plausible trends that have important implications for technology policy in this space. We used these trends as the basis for our second method – *scenario planning*. The purpose of scenario planning is to generate strategic responses to possible futures. Through scenario planning we identified several themes that may affect policy outcomes and used one of those themes as the central point for the *futures wheel*. The futures wheel is a tool for exploring 1st, 2nd, 3rd and nth order effects of a trend or decision. Using the futures wheel, we surfaced a critical unintended consequence of drone regulation. In the following sections, we describe the methods in detail and apply them to the case of robotic delivery.

Design Fiction

Design fiction is an increasingly popular method for envisioning the evolution of technology and its social impacts through narrative iteration. Design fiction typically takes the form of a written artifact, often presented in a set (Baumer, et al., 2018; Wong, et al., 2018) but they can also take on multimedia forms such as videos (Savov, 2018). Frequently, the fictions themselves are tangential to the technology of focus, emphasizing the infrastructures that exist to support our technological systems. For example, the design fictions related to Brain Computer Interfaces (BCIs) written by Wong and colleagues highlight the existence of an API to access the technology and the ways in which crowdsourced workers interact with the API.

Design fiction does not attempt to predict the future, but instead allows us to consider the technology at hand “in relation to the sociocultural contexts in which it is presumed to exist” (Wong, et al., 2018, pp 1360). The speculative nature of design fiction encourages us to consider the future as a curiosity, something that can be influenced and changed. It severs the link of technological determinism by reminding us that we, as designers, technologists, and policymakers, continually interact with and influence the world around us. Most importantly for our purposes, design fiction encourages us to consider the kinds of futures we want and the kinds we do not (Dunne & Raby, 2013). This understanding better equips us to act in the present through tech policy and guide technology toward desirable futures—the subject of the final section.

As with many futures methods, design fiction is ill equipped to define which futures are preferable and for whom. However, by stimulating conversation about the future using speculative methods, we are able to create space for a conversation about which futures we prefer and why.

Here, we apply design fiction to our case study of urban robotic delivery to explore possible futures in this domain. We began with ideation about the future of robotic deliver technology considering questions such as, in what ways might technology develop in the future? What infrastructure is needed to support this technology? How would people interact with deliver robots? What would it be like to live in a future where such robots are ubiquitous? How would widespread use of robots change the ways we work and live? In what ways could these systems fail? And so on. After identifying some key elements to include, we created a narrative to show how a person might move through a daily task – retrieving something they forgot at home – in this future to support several design fiction artifacts.

Forgetting something at home, 2037

K has a very important presentation with a high profile client. She's getting ready for the day when she receives a notification that the meeting has been moved up several hours and she needs to be at work in 30 minutes to finish preparing. She rushes out the door and just barely makes the hydro-bus into the city. As the vessel pulls away from the dock, K realizes that she's made a horrible mistake. Today, of all days, she forgot the prototype and design prints for the presentation. Without them, she's missing the heart of her presentation content. K quickly downloads a personal drone delivery service, WingSail, which provides her with a menu of delivery options.

Frak. The case is too heavy to fly. She has mixed feelings about aerial drones anyway after the memorial went up outside her work. Several years ago, when drones really started to displace workers, an angry UPS driver shot a drone out of the sky in protest of losing her job. Turns out, it was someone's personal drone. The owner had come to think of the drone as their pet and set up a memorial to honor their loss.

K selfishly hopes her drone doesn't get shot today, because she really needs that damn portfolio and soon. And a water-to-land handoff could get the package to the office in under an hour.

K finalizes the delivery as the ferry docks in the city and she continues her walk to work. She has a 20 minute walk from the docks to her office building and it's the most annoying part of her commute. About two minutes in she's joined by a sidewalk drone. Companies like WingSail make money on their deliveries, but most of their profits come from ads. There are ads on the side of every flying drone and every sidewalk drone. They even project ads on the ground in front of you or on the walls beside you. Lately, K has noticed the ads to reflect her recent purchases. Targeted advertising is ubiquitous, but she's never noticed the drones using it until now. How are they sensing it's her?

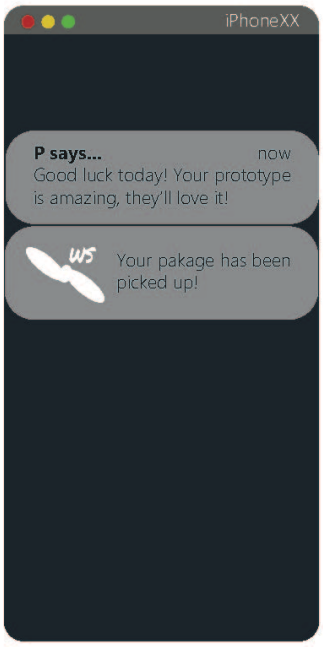

K is joined by an inevitable pack of sidewalk drones. One, name of Sparks, is covered in tiny butterfly stickers like it got into a fight with a three year old. K thinks it's kind of cute until it starts projecting ads in front of her. This ad is particularly embarrassing and seems to be based on a google search she did last night for line dancing classes as a gag gift for a friend. You can't even play practical jokes anymore, they end up as a joke on you.

Targeted ads are so ubiquitous now no one seems to notice or care, but K is still relieved when Sparks takes a sharp right turn to continue its delivery route. She wonders where it may be headed, and with what.

K makes it to work and preps the meeting room. She anxiously checks the map every few minutes as the drone makes its way to her office building. She wonders what kind of ads it played to commuters along the way.

The portfolio arrives just in time. K whips out the prototype and pours all her anxious energy and relief into the presentation. K's clients are very pleased with her work. K thanks her lucky stars to live in this age of wonder.

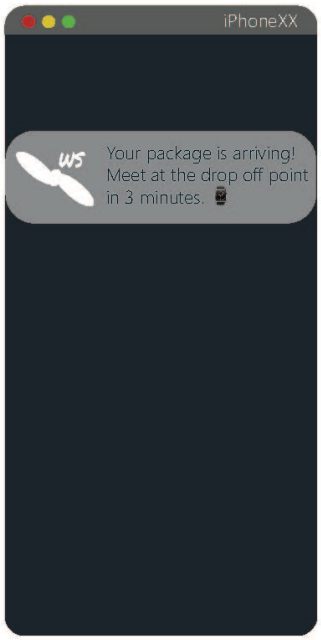

Scenario Planning

In the generative session regarding the design fiction content we identified the specific way a person might engage with this technology, a new spatial dimension (water based drone delivery) that has yet to be considered in the policy landscape, and the way two technologies (drones and personalized advertising) might exist together. By considering just the existing technology or the current policy landscape in isolation we would likely not have opened up this space for consideration. These themes, water based drone delivery, drones and personalized advertising, and consumer choice in drone delivery, were the basis for our next method, scenario planning.

Scenario planning is a method for exploring strategic responses to possible futures. It was pioneered by Herman Kahn at the RAND Corporation in the 1950's and further developed by economists Pierre Wack and Edward Newton at Shell Oil. In developing scenarios, planners consider both the future possibilities and strategic responses to those possibilities, so there is equal focus on exploring possibilities and how to manage those possibilities. Scenario planning is widely credited with helping Shell Oil successfully navigate the Iran oil crisis of the 1970's. At the time, Shell had a well-defined practice around scenario planning that included both a) rigorous work to understand the business environment and policy landscapes and b) attention to how the scenarios were received and utilized by managers (Wack, 1985). Wack describes the real work of scenarios as encouraging decision makers to question their assumptions of the future.

Scenarios are short narratives that plausibly describe a future and the series of events that could lead to the emergence of that future. They are typically presented in contrasting sets of three or four scenarios, each describing a different possible future. The narrative form of

scenarios concretizes the future in a way that makes the variety of possibilities seem more real, making them an effective tool for soliciting strategic responses.

As with most futures methods, scenario planning is not a tool for predicting the future and they should not be judged on whether or not that future comes to pass. Often, none of the precise futures described in the scenarios are realized. Scenarios should instead be evaluated on their effectiveness at soliciting strategic responses. In other words, decision makers should be able to identify strategies and responses to the range of possible futures described in the scenario set. Similarly, most of the value derived from the scenario planning process is experiential. The work of creating scenarios highlights assumptions about the future, encourages us to question those assumptions, increases our ability to anticipate change, and subsequently respond to change effectively.

When developing the scenarios for this case study, we used the branch analysis method described in (Foresight Horizon Scanning Centre, 2009). Branch analysis is one of several methods for scenario planning in which critical decisions are arranged as points along axes moving toward the future. In our branch analysis, we identified trends from the design fiction, built a branch diagram based on future events, identified a range of possible outcomes for each branch, and developed characteristics for the scenarios. In a full scenario planning project, these scenarios would be developed into narratives.

For our purposes of demonstrating the method and how multiple futures methods can work together in an analysis, we did not fully develop the narrative structure of the scenarios. Figure 1 below shows the branch diagram we created as part of our scenario planning process. Scenarios A-E and their corresponding policy and social implications are also below.

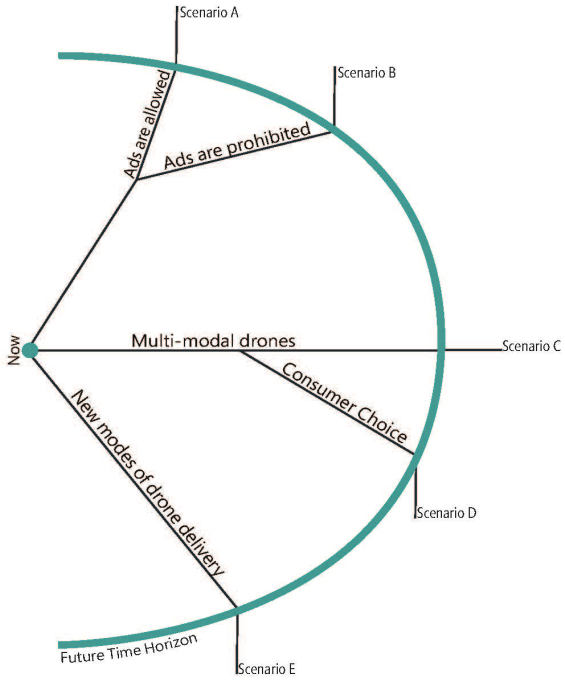

Figure 1: The branch analysis diagram for the case of personal drone delivery.

Scenario A: Ads are used to offset the cost of robotic delivery. Robots flying and driving around major cities are a source of distraction. Personalized ads are constantly violated people's privacy in public places. People have begun to move out of neighborhoods where drones delivery frequently. Some people have stated wearing screen blocking glasses so they can't see the pervasive ads.

Implications: Robotic delivery is more affordable, leading to more and equitable widespread use. Ubiquitous screen technology on robots provides a new creative medium for artists (e.g., interactive or distributed murals) and a new mode for government agencies to transmit public service announcements. The ubiquity of drones could be useful in an emergency if critical infrastructure is damaged and emergency response is limited to aerial delivery of supplies.

Scenario B: Ads are prohibited on robots in many cities because of the distraction. As a result, robotic delivery is more expensive, so fewer people can access this important delivery service for things like medical supplies. Several cities are sued for violation of free speech because of the categorical ban on advertising on robots (cf. City of Cincinnati v. Discovery Network, 1993).

Implications: Drone delivery is more expensive as the cost is not offset by advertisements, reinforcing issues of equity and access for low-income individuals. Local governments face litigation regarding prohibition of free speech which is costly and time consuming.

Scenario C: Companies have developed what they call “atomized delivery.” A large, autonomous container ship can pull up to the shore and thousands of individual flying drones can take off from the ship’s deck to delivery packages throughout the region. As a result, major job loss has occurred in the shipping industry and port cities have suffered economically.

Implications: Atomized delivery could disrupt our current customs process. Large impacts to the global economy and globalization more generally. Significant job loss across multiple sectors. Novel handoff problems regarding safety arise as robots transition between different modes of operating.

Scenario D: Robotic delivery has become ubiquitous to the point where consumers can order anything and have it delivered. Coffee, lunch, prescriptions. Anything that could be delivered to your home or workplace is now delivered using robots. Consumers can even specify which type of robot they want to delivery their items (e.g., flying, driving, ads, no ads, etc.). Individuals and families own drones almost like pets. When they need something picked up they send the drone to retrieve the items. A new government agency has been stood up to oversee drone network interoperability as numerous companies and families fly and drive drones within busy urban areas.

Implications: Need for regulated network interoperability across multiple delivery modes (air and land) as well as personal and commercial drone operations. Municipalities may limit robotic delivery areas to minimize disruption.

Scenario E: Robot delivery started out exclusively as aerial or land based mechanisms, but quickly spread to other delivery forms. In Seattle and other cities situated around waterways, underwater robots became a popular option for avoiding the strict FAA regulations. In cities with underground subway systems, drones latch on to the existing infrastructure as they move about the city delivering packages. Innovators are constantly pushing the boundaries of how we can deliver items.

Implications: Current policy may be inadequate for anything other than air or land based robotic delivery. Other robot delivery mechanisms may affect city or private infrastructure.

Futures Wheel

Through scenario planning, we identified several trends related to the future of UDD, possible outcomes of these events, and a range of policy responses. The trends we identified could arise through either policy decisions (such as the decision to prohibit ads on drones) or through markets (the emergence of unforeseen drone technologies). For policy makers, it may not be important to intervene at each decision point, but our experience demonstrates that it is important to explore the range of possibilities to anticipate where current policy efforts may fall short. One of these trends, the ubiquity of ads on drones, was the foundation for our final method, the futures wheel.

The futures wheel is a tool for exploring 1st, 2nd, 3rd and nth order unintended consequences of a specific decision, trend, or event. The futures wheel is a technique developed in the early 1970's by social scientist and civic leader Jerome Glenn as a way to structure and organize thinking about the future. It spread through workshops, futures education courses, and consultants and it is widely used in futures work.

The futures wheel is a useful tool for thinking through the possible impacts of a trend in a way that demonstrates complex interrelationships. The process is intuitive and inexpensive to implement. It facilitates thinking about future impacts in a rigorous way that allows us to identify potential issues at a much deeper level than if we were causally thinking about the future.

While concentric rings help us think deeply about the future, they can also become overwhelming. If not well organized, futures wheels can become too complex, hiding patterns that may exist. The potential impacts identified using a futures wheel are limited to the frames of reference held by the participants.

In applying the future wheel to our case study, we used the ubiquity of drone advertising as the central trend to explore. Radiating from the central trend, we identified a range of first order effects that may arise from an increase in ads on drones, followed by the second, third and fourth order effects. Figure 2 below shows the futures wheel we developed for this case study.

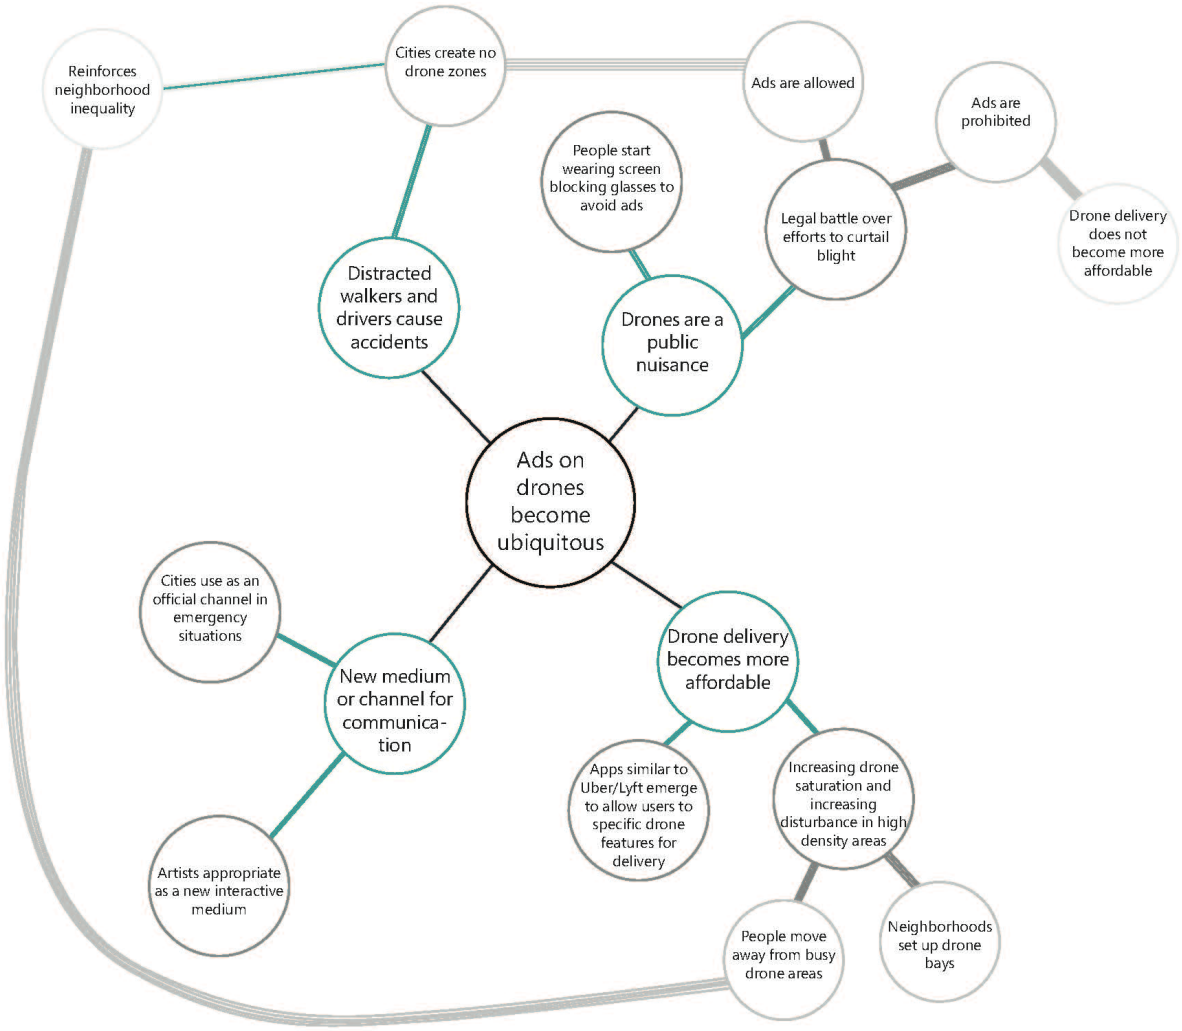

Figure 2: The futures wheel for the case of personal drone delivery.

In organizing and reorganizing the impacts in our wheel, we surfaced several important impacts that were not previously discussed in relation to PDDs including the possibility of "no drone zones," the rise personalized drones for delivery, and the potential for PDDs to reinforce neighborhood inequality. In cities already experiencing high rates of neighborhood inequality, any zoning decisions and blight have the potential to reinforce inequality. We surfaced this possible impact of three distinct primary impacts, demonstrating that there are multiple paths by which the ubiquity of advertising on drones could reinforce neighborhood inequality. In this case study, the futures wheel deepened our understanding of the potentially widespread impacts of PDDs and corresponding policies. As with all envisioning methods, impacts surfaced using a futures wheel are not prediction, but they do point to issues we should monitor in the development of policies for PDD.

So What? How Forecasting Can Improve Policymaking

Policymaking is difficult, technology policymaking perhaps especially so. A significant component of this difficulty flows from the challenges around predicting how technology might develop and how individuals and society will come appropriate it.

A key insight of David Collingridge's book *The Social Control of Technology* is that there exists a tension between how much information is available about technology and the possibility of exerting control. The basic idea—sometimes referred to as the Collingridge dilemma—is that the trajectory and social impacts of nascent technology are hard to know in advance, leading to an information problem. Yet the longer we wait and see, the harder the technology will be to control. Perhaps the technology becomes diffuse, path dependent, or people come to rely upon it; perhaps specific instantiations of the technology accrue vested economic or political interests.[6] Regardless, in waiting for sufficient information, we have a better understanding of how certain technologies affect society and our environment, but those technologies become harder to regulate.

One sees echoes of Collingridge in discussions of the "pacing problem," i.e., the popular belief that contemporary technology evolves too quickly for law to address. The pacing problem is often overstated and undertheorized, alternatively conveying a sense of urgency and futility. The former can lead to hasty legislation that waste resources, become rapidly outdated, or otherwise do more harm than good. The latter can lead to paralysis as legislatures fear to act in any way that may hinder business, stifle innovation, or upset constituents.

Sophisticated analyses of pacing exist (e.g., Marchant et al. 2011) and tend to focus on designing governance mechanisms capable of anticipating or responding to complex socio-technical interactions. The field of technology assessment also acknowledges, without resolving, the trade-offs between intervening too early and waiting too long to attempt to channel a technology (e.g., Tribe 1973). This leads to a number of insights and recommendations, including deeper analysis of technology by government and the acknowledgment that the preservation of options constitutes an important governance value.

At the University of Washington Tech Policy Lab, we collect case studies of instances of tech policy breakdowns in an effort to identify common mechanisms.[7] Several of the common mechanisms we've identified involve the failure to anticipate the direction or effects of emerging technology. For example, a law could be over or under-inclusive, "brittle," or

[6] Later authors have layered in considerations of political economy, observing that vested economic interests will push policymakers to accept the most optimistic of a range of assessments. Edward A. Parson, Social Control of Technological Risks: The Dilemma of Knowledge and Control in Practice, and ways to Surmount It, 64 UCLA L. Rev. Disc. 464, 468 (2016), citing Donald Ludwig et al., Uncertainty, Resource Exploitation, and Conservation: Lessons for History, 260 Sci. 17 (1993).

[7] This work, funded by the Hewlett Foundation, also seeks to document impartial evidence that a tech policy has broken down, since one person's breakdown could be another's triumph. These include subsequent retraction or preemption, the failure to achieve a stated objective, widespread criticism, lack of stakeholder remedies, and other criteria identified through literature review and expert surveys.

premised on assumptions about technology or practice that quickly becomes obsolete. Often these breakdowns find their source in a failure of imagination or else the reliance on the technical vision of a particular stakeholder or set of stakeholders to the exclusion of any other. For example, the State of Nevada had to repeal and redraft its pioneering driverless car law after only a year because it reflected Google's vision of complete vehicle autonomy but swept in partially autonomous features—such as auto lane correction, adaptive cruise control, and self-parking—that already existed in luxury cars being sold in the state (Calo, 2014).

We see these dynamics at play in the context of robotic delivery. In apparent response to lobbying by one or more specific startups, states have promulgated rules for robotic delivery that enshrine a particular, narrow vision of the enterprise: small, land-based carts that travel on sidewalks and crosswalks. The FAA, meanwhile, continues to maintain exclusive jurisdiction and tight controls over aerial drone delivery that include a requirement even for testing of line of sight operation from a fixed space on the ground. Any potential entrant into this market must adhere to one or other set of specific and stringent guidelines.

By applying three forecasting techniques, we were able to surface a wide variety of plausible directions for robotic delivery within a matter of hours. These included *technical variations* such as an entirely new geographic element (water), as well as multimodal delivery (e.g., drones that can both drive along sidewalks and fly); *social dimensions* (e.g., personal guilt over externalities or socio-economic impacts on city zoning); and *market possibilities* (e.g., consumer choice as to modality of transport or the expansion of advertising services).

These plausible (but again, hardly inevitable) directions for robotic delivery could be instructive to policymakers along at least three lines. First, they militate in favor of regulatory flexibility. The present scope of PDD laws are highly narrow and premised upon a unitary vision of what robotic delivery affords. New entrants to the delivery market with a different technology or vision (e.g., four legged systems) would have to re-approach a legislature—a body with competing priorities that feels like it has already dealt with the robotic delivery issue. Innovators with ideas of how to utilize alternate channels of transportation, such as underground or water based drones, are left with no regulatory pathway at all. And highly plausible instantiations of last mile delivery—such as trucks that drive through neighborhoods and release aerial drones for front door delivery—cannot even be tested under existing restrictions.

Regulators interested in greater flexibility, meanwhile, have several levers. One is a sunset provision—which only a single, wise state has adopted. Sunset provision are time limits on rules that force reexamination of the technical and social landscape and permit for iterative modifications. Another is safe harbors or exemption mechanisms that permit industry to propose and pilot alternatives—technically the FAA can do this with drones but in fact seldom deviates from its general guidelines. Yet another involves the use of standards instead of rules, including goal- or performance-based standards that incentivize safety and other specific values without prescribing a particular technical configuration or strategy (cf. Willis 2015). These are just a few examples.

Second, these techniques suggest ways that states and cities can prepare for greater automation in last mile delivery. In addition to safety considerations, governments will have to grapple with the effects of robotic delivery on information privacy, solitude, and visual blight. Crucially, robotic delivery at scale invites considerations of geographic equity, as certain neighborhoods flourish or flounder due to issues of access, affordability, or nuisance. The passage of state statutes or city ordinances that do not grapple with such considerations spend political capital unwisely. Later efforts to address these concerns will face headwinds because, again, policymakers are likely to feel as though they have already addressed the issue of robotic delivery.

Finally, failures of imagination lead to opportunity costs. Futures work is not, or should not be, limited to envision the risks of new technologies. There are also the opportunities. Significant developments in technology represent an invitation to inventory our societal goals and values and revisit whether new affordances help us achieve them. Our futures exercises surfaced a number of opportunities for municipalities regarding robotic delivery, including new modalities of public art and expression, new avenues for public service announcements, and help with disaster response and relief. Robotic delivery holds great potential for government bodies that cannot be realized simply be enshrining the vision of a for-profit start up into law.

These techniques are useful not only to policymakers in designing a wiser, more inclusive legal framework, but also for scholars, students, and others hoping to develop a better understanding of emerging technology and its social impacts. By taking futures seriously, we begin to develop a richer and more responsive technology policy in theory and practice.

Conclusion

References

Andersson, Jenny. *The Future of the World: Futurology, Futurists, and the Struggle for the Post Cold War Imagination*. First edition., Oxford University Press, 2018.

Baumer, Eric P. S., et al. "What Would You Do?: Design Fiction and Ethics." *Proceedings of the 2018 ACM Conference on Supporting Groupwork*, ACM, 2018, pp. 244–256. *ACM Digital Library*, doi:10.1145/3148330.3149405.

Bear, Elizabeth, et al. *Future Visions: Original Science Fiction Inspired by Microsoft*. Microsoft, 2015.

Collingridge, David. *The Social Control of Technology*. Frances Pinter, 1980.

Dunne, Anthony, and Fiona Raby. *Speculative Everything: Design, Fiction, and Social Dreaming*. 1st edition, The MIT Press, 2013.

Funk, David A. "Legal Futurology: The Field and Its Literature." *Law Libary Journal*, vol. 73, 1980, pp. 625–33.

Gibbs, Amy. "Using Science Fiction to Explore Business Innovation." *Digital Pulse*, 4 Sept. 2017, https://www.digitalpulse.pwc.com.au/science-fiction-explore-business-innovation/.

Glenn, Jerome C., and Theodore J. Gordon, editors. *Futures Resarch Methodology Version 3.0*. 3rd ed., http://107.22.164.43/millennium/images/FRM-V30.jpg. Accessed 20 June 2018.

Greenblott, Joseph M., et al. "Strategic Foresight in the Federal Government: A Survey of Methods, Resources, and Institutional Arrangements." *World Futures Review*, Dec. 2018, p. 194675671881490. *Crossref*, doi:10.1177/1946756718814908.

National Intelligence Council (U.S.). *Global Trends: Paradox of Progress*. 2017.

Savov, Vlad. "Google's Selfish Ledger Is an Unsettling Vision of Silicon Valley Social Engineering." *The Verge*, 17 May 2018, https://www.theverge.com/2018/5/17/17344250/google-x-selfish-ledger-video-data-privacy.

Scenario Planning. Foresight Horizon Scanning Centre, 2009, http://webarchive.nationalarchives.gov.uk/20140108141323/http://www.bis.gov.uk/assets/foresight/docs/horizon-scanning-centre/foresight_scenario_planning.pdf.

Tribe, Laurence H. *Channeling Technology Through Law*. Bracton Press, 1973.

Wack, Pierre. "Scenarios: Shooting the Rapids." *Harvard Business Review*, vol. 63, no. 6, Dec. 1985, pp. 139–50.

Willis, Lauren E. *Performance Based Consumer Law*, 82 Chicago Law Review 1309, 2014.

Wong, Richmond Y., et al. "When BCIs Have APIs: Design Fictions of Everyday Brain-Computer Interface Adoption." *Proceedings of the 2018 Designing Interactive Systems Conference*, ACM, 2018, pp. 1359–1371. *ACM Digital Library*, doi:10.1145/3196709.3196746.