5 The Responsible City: Avoiding Technology's Undemocratic Social Contracts

Throughout the book, we have explored how technology's social impacts are influenced by far more than simply its capabilities: social and political conditions constrain the outcomes that technology can generate, people can wield similar technologies to achieve distinct outcomes, and machine learning models derive insights from data that reflects a legacy of bias and discrimination. By overlooking these factors, tech goggles encourage the adoption of smart city technologies that generate unintended and undesirable social consequences.

This chapter introduces one more component to the discussion: the technical and political arrangements of technology, also known as its architecture. Questions related to architecture go beyond a technology's purpose and output—What should this algorithm do? Is it accurate?—and encompass its structure: By what means should this technology achieve its goals? Who should control it? How should we pay for it?

The answers to these questions can have significant ramifications. The ways in which technology structures social and political relations can be even more consequential than the explicit function that it serves. When developing and adopting new technology, we must therefore "examine the social contract implied by building that [technological] system in a particular form," asserts Langdon Winner. "[A]s our society adopts one sociotechnical system after another it answers some of the most important questions that political philosophers have ever asked about the proper order of human affairs," writes Winner. "Just as Plato and Aristotle posed the question, What is the best form of political society?," today we must ask, "What forms of technology are compatible with the kind of society we want to build?"[1]

Whether we recognize it or not, the technologies we implement in cities today will play a significant role in defining the social contract of the next century. And as it currently stands, the architecture of the smart city is a fundamentally undemocratic one: many technologies operate by collecting unchecked data about individuals and using opaque, often proprietary, algorithms to make life-altering decisions. In the process, they create massive information and power asymmetries that favor governments and companies over those they track and analyze, breeding impotence and subjugation within society. In this way, the smart city is a covert tool for increasing surveillance, corporate profits, and social control.

City governments eager to take advantage of new technologies must act as responsible gatekeepers and public stewards in structuring their technology to protect equity and fundamental rights. The Smart Enough City need not accept the unsavory compromises of smart cities—more democratic approaches to utilizing technology are possible.

* * *

Equitable access to high-speed internet has become a necessary staple of democratic society. Without internet access, it is difficult if not impossible to apply for jobs, access healthcare, and connect with other people. Yet because of the high price of internet subscriptions, many low-income individuals and families are unable to afford reliable broadband access. In Detroit, for example, 40 percent of residents lack broadband.[2] In New York City, that number is 23 percent.[3]

In 2016, it appeared that New York City had found the solution to this digital divide that could be a model for every city: LinkNYC, a program to provide free public Wi-Fi via more than 7,500 internet-connected kiosks placed throughout the city (as of November 2018, 1,742 had been installed).[4] "LinkNYC brings us a couple steps closer to our goal of leveling the playing field and providing every New Yorker with access to the most important tool of the 21st century," proclaimed Mayor Bill de Blasio at the program's launch. Perhaps most amazingly, providing this service will not cost the government a cent—in fact, NYC expects the program to bring in more than $500 million in revenue to the city by 2025.[5]

This appears, like many smart city technologies, to be a benevolent technical solution for an important social problem. But under the surface, where LinkNYC's architecture resides, lurks a more insidious reality.

The benefits and finances of LinkNYC sound too good to be true. So how is the program funded? The kiosks are owned and operated by Sidewalk Labs, a subsidiary of Alphabet (Google's parent company), which plans to pay for the initiative by collecting and monetizing data about everyone who uses the service. As Sidewalk founder and CEO Dan Doctoroff told a public audience in 2016, the company expects to "make a lot of money from this."[6]

LinkNYC kiosks are equipped with sensors that gather an enormous quantity of data about every device that connects to the Wi-Fi network: not just its location and operating system but also its MAC address (a device's unique identifier that helps it connect to the internet).[7] Sidewalk Labs claims that this data is merely "Technical Information," in contrast to the "Personally Identifiable Information" it collects such as your name and email (which are required to register for network access).[8] This distinction follows traditional standards of privacy, which focus on the presence of personally identifiable information (PII)—features such as name, address, and Social Security number that, on their own, identify individuals. Data containing PII is considered sensitive, while data without PII is not.[9]

To the human eye, this appears to be a sensible distinction. After all, MAC addresses are twelve-digit alphanumeric strings that look like indecipherable, machine-processing gobbledygook. But just because data does not contain names and is difficult to decipher does not mean it lacks information about people. Yes, one single data point—a phone's MAC address at a particular place at a particular time—is unlikely to reveal someone's identity or anything sensitive about them. But when millions of data points are collected and combined with modern analysis techniques, such data can be used to track people's movements and infer intimate details of their lives.

This data is so sensitive in aggregate, despite each record seeming so benign in isolation, because human behavior is incredibly unique. Data collected on a massive scale captures these particularities. Research led by the computer scientist Yves-Alexandre de Montjoye demonstrated this phenomenon by analyzing two datasets that contained mobile phone location traces and credit card transactions about more than one million people.[10] Even though both datasets lacked PII—they included just anonymous personal IDs (à la MAC addresses), locations, and times—de Montjoye revealed that it was possible to identify individuals and learn about their behavior. Remarkably, more than 90 percent of people could be uniquely

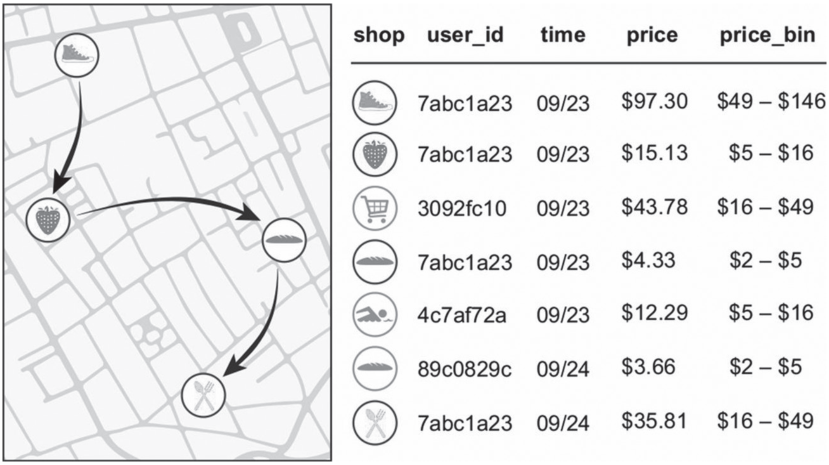

shop	user_id	time	price	price_bin
👟	7abc1a23	09/23	$97.30	$49 – $146
🍓	7abc1a23	09/23	$15.13	$5 – $16
🛒	3092fc10	09/23	$43.78	$16 – $49
🥖	7abc1a23	09/23	$4.33	$2 – $5
🏊	4c7af72a	09/23	$12.29	$5 – $16
🥖	89c0829c	09/24	$3.66	$2 – $5
🍴	7abc1a23	09/24	$35.81	$16 – $49

Figure 5.1
A stylized example of behavioral data that includes no personally identifiable information (PII) but nonetheless contains information about individuals. By examining all of the records pertaining to a given person, it is possible to infer their behavior. The map above traces the activity of the person represented by user_id=7abc1a23.
Source: Yves-Alexandre de Montjoye, Laura Radaelli, Vivek Kumar Singh, and Alex "Sandy" Pentland, "Unique in the Shopping Mall: On the Reidentifiability of Credit Card Metadata," *Science* 347, no. 6221 (2015): 537. Reprinted with permission from AAAS.

identified with just four data points of where they have been and when they were there.

Although de Montjoye's analysis shed new light on the privacy risks of granular behavioral data, it was not the first demonstration of how supposedly anonymous data can reveal a great deal about people. That came in 1997, when Massachusetts Governor William Weld released state employee medical records for research purposes, promising that the information was anonymous. A few days later, however, Weld received a letter: it contained his own health records, culled from the released data.[11] The envelope came from Latanya Sweeney, then a graduate student at the Massachusetts Institute of Technology, who had identified Weld's file by linking the medical records with publicly available voter lists via information contained in both datasets (such as birth date).[12]

Many other datasets come with similar risks of reidentification. When New York City released supposedly anonymous data about every local taxi trip in 2013, a data scientist analyzed patterns of where those rides began and ended to identify the names of several patrons of a Manhattan strip club.[13] The same method could be used to learn who prays at mosques, works late nights, visits gay bars, or undergoes chemotherapy. Another data scientist used similar data about trips on London's bike share program to reveal the travel patterns of several individuals, deducing where they live and work.[14]

But the danger is not merely that one's identity and behavior can be revealed from seemingly anonymous data—when data is combined with artificial intelligence, it is possible to infer a great deal of personal information that is not explicitly contained in a dataset. With detailed information about where you have been, for instance, machine learning algorithms can predict whom you know and where you will go next.[15] Algorithms can detect whether someone is depressed on the basis of photos she posts on Instagram.[16] Data about seemingly routine behaviors such as Facebook Likes can reveal sexual identity, race, political affiliation, and even whether one's parents are married.[17]

The potential for algorithms to identify and learn about people using supposedly anonymous and benign data highlights the remarkable invasion of public privacy presented by LinkNYC and reveals the definitional trick that enables it: Sidewalk Labs's "Technical Information" is presented as being anonymous, but it is in fact far more sensitive than the "Personally Identifiable Information" that the company so magnanimously promises to safeguard. In other words, says one privacy lawyer, the LinkNYC privacy policy is designed "to make you believe that something is being promised, when actually it lets them do anything they want."[18] The motivation is profit: the more detailed the data, the better Sidewalk Labs can monetize it.

Recognizing these privacy risks, many New Yorkers have raised concerns about LinkNYC.[19] The New York Civil Liberties Union's executive director has argued, "Free public Wi-Fi can be an invaluable resource for this city, but New Yorkers need to know there are too many strings attached."[20]

In *The Digital Person*, the privacy scholar Daniel Solove highlights two fundamental ways in which such widespread data collection and concomitant knowledge threaten society. The most salient concern is widespread surveillance, as governments and companies are allowed to watch your every action, expose secrets, or even catch everyone who runs a red light. Such fears tap into deep-seated cultural notions about privacy that are drawn from images of Big Brother, the totalitarian government in George Orwell's 1949 novel *1984*. By observing the most intimate details of everyone's lives and punishing even the slightest dissent, Big Brother controls the society's behavior. Following Orwell's influence, writes Solove, we typically conceive of privacy following the "secrecy paradigm": the idea that privacy is invaded when one's secrets are observed or exposed, leading people to self-censor (via "chilling effects") or suffer the consequences.[21]

1984-inspired fears capture a great deal of why privacy is essential for maintaining civil liberties. "The government has a long history of spying on activists in an effort to dampen dissent and we see those same actions present today, in an age where most people have substantial digital footprints," explains the activist DeRay McKesson.[22] The FBI in the 1960s, for example, focused surveillance on civil rights activists such as Martin Luther King, Jr., in order to intimidate and harass them.[23] This history appears to be repeating itself, as federal and local law enforcement officials have been tracking the identities and activities of those protesting police violence in the Black Lives Matter movement.[24]

Yet Big Brother cannot explain every risk of declining privacy. As we have already seen, a significant portion of data collection today relies on information that is neither secret, illegal, nor embarrassing—in fact, many individual data points appear meaningless and anonymous. The secrecy paradigm thus fails to explain the harms of someone's bike share trips or Facebook Likes being collected, aggregated, and analyzed. As Solove explains, nowadays many uses of data "aim not at suppressing individuality but at studying it and exploiting it."[25]

Solove likens much of today's collection and use of data to the themes of another twentieth-century novel: Franz Kafka's 1925 *The Trial*. The book's protagonist, Josef K., wakes up on his thirtieth birthday to find two men in his room declaring that he is under arrest. He is given no indication of what he has done or what agency is arresting him. The novel depicts Josef's unsuccessful attempts to uncover the identity of the mysterious court and what data it possesses about him. He is murdered by the court's agents on his thirty-first birthday without ever having learned their true nature. "*The Trial* captures an individual's sense of helplessness, frustration, and vulnerability when a large bureaucratic organization has control over a vast dossier of details about one's life," Solove explains. In doing so, it "captures the scope, nature, and effects of the type of power relationship created by databases."[26]

Just like Josef, people today have little knowledge of or control over what personal data is collected, who owns it, and how they exploit it. As more data is gathered and used by governments and companies, privacy becomes defined less by the secrets that any single piece of information reveals and increasingly by the inferences that large amounts of relatively nonsensitive data make possible—and the power that those inferences grant. For example, Facebook can calibrate its News Feed algorithm to influence a user's mood and likelihood to vote.[27] OkCupid can alter profile match scores to affect certain people's chances of getting dates.[28] Healthcare services can deny coverage if they learn that someone recently visited websites associated with having cancer.[29]

Although data collection touches everyone, the most severe impacts of diminishing privacy are suffered by the poor and minorities. Despite being more concerned than their more well-off counterparts about privacy, most lower-income individuals lack the knowledge of privacy settings and policies to sufficiently reduce the extent to which they are tracked.[30] And given that activists in racial justice movements like Black Lives Matter are targeted for surveillance and undocumented immigrants face deportation, minorities are prone to suffer the most damaging consequences of being identified and tracked by the government.

Moreover, those with the lowest socioeconomic status often have no choice but to accept government oversight in exchange for social services. Welfare offices use electronic benefits transfer (EBT) cards to monitor recipients' behavior in increasingly pervasive and intimate ways. As the political scientist Virginia Eubanks explains, these technological practices "significantly limit clients' autonomy, opportunity, and mobility."[31]

This suffocating oversight has long been a feature of government services. In his 2001 book *Overseers of the Poor*, the political scientist John Gilliom documents how welfare recipients are closely monitored by the government through endless paperwork and meetings with fraud control agents to ensure that they are eligible for services and comply with the many requirements. The government strictly constrained the parameters of daily life of the Appalachian "welfare mothers" whom Gilliom studied, its watchful and persistent eye leading to "hassle and degradation" that "hindered [the mothers'] ability to meet the needs of their families." These women were forced to adhere to restrictive rules while simultaneously finding it necessary to skirt those rules in order to survive. They thus experienced surveillance not as an invasive uncovering of their secrets but as a loss of "a great deal of the autonomy and control that they could otherwise have over their personal lives." As one explained, "All the time you are on welfare, yeah, you are in prison."[32]

The poor and minorities are also most susceptible to harms caused by a lack of privacy in dealings with the private sector. As companies increasingly make decisions about individuals based on data drawn from their online behavior and social networks, lower socioeconomic groups can be unfairly excluded from credit, jobs, housing, and healthcare in ways that circumvent anti-discrimination laws.[33] Low-wage workplaces monitor their employees' keystrokes, location, emails, and online browsing to detect unsanctioned behavior, which can result in firing.[34] The profiles created by data brokers (such as "suffering seniors" and "urban scramble") make it possible for companies to target those susceptible to predatory loans and scams with precision.[35]

Thus, as a 2018 book by Eubanks suggests, these privacy infringements and algorithms conjure up a new meaning for the acronym AI: not "artificial intelligence," but "automating inequality."[36]

* * *

The smart city represents the vast expansion of both government and corporate data collection. Embedding sensors, cameras, software, and an internet connection in everyday objects from streetlights to trashcans—creating what is known as the "Internet of Things"—makes it possible to collect remarkably precise data about what is happening in a city. This data could be used to facilitate beneficial outcomes: reducing traffic, improving infrastructure, and saving energy. But it also includes detailed information about the behavior of everyone within the city.

Smart city technologies make it easier than ever for municipalities to identify and track individuals. Sensors on streetlights and other forms of "smart" infrastructure (like LinkNYC kiosks) can track internet-connected devices in their vicinity, making it possible to follow people's movements throughout the city. Cameras paired with software that can identify people or objects generate additional surveillance threats. In Los Angeles, for example, automatic license plate readers (ALPRs) record the location of three million vehicles every week, collecting information that often finds its way into the hands of U.S. Immigration and Customs Enforcement

(ICE).[37] The push for police-worn body cameras, supported by many as a tool to hold police accountable, creates the potential for widespread surveillance by police of all public space: given that body camera manufacturers are developing facial recognition software to analyze this footage, and given that only one police department in the United States has a policy governing body cameras that "sharply limits the use of biometric technologies,"[38] it is likely that body cameras will soon be used by police wherever they go to track the movements of civilians, identify participants at protests, and scan crowds to see who has outstanding warrants.[39] Making similar use of technology, police in Orlando are using Amazon's facial recognition service to monitor in real time everyone who appears in the video feeds from traffic cameras.[40]

Meanwhile, for companies eager to expand their data collection beyond your browser and into physical space, the smart city is a dream come true. Many companies already possess the knowledge and influence necessary to restrict individual autonomy and exploit people, but if companies like Sidewalk Labs have their way, smart city technologies will vastly increase the scale and scope of data they collect. Companies that place cameras and MAC address–sniffing sensors on Wi-Fi kiosks, trashcans, and streetlights will gain heretofore unattainable insights about the behavior of individuals. And given the vast reach of hard-to-trace data brokers that gather and share data without the public's knowledge or consent, one company's data can easily end up in another's hands.[41]

Once these smart city technologies are installed, it is practically impossible for someone to avoid being tracked. Many defend online companies' mass data collection by pointing to the opportunity opt out: if you don't like data about you being collected, don't use the websites or apps that collect it. But since it is almost impossible to communicate, travel, or get hired without email, search engines, smartphones, and social media, this is an unreasonable choice. In the emergent smart city, with sensors and cameras on every street corner—remember, New York is slated for more than 7,500 LinkNYC kiosks—this argument reaches an even more perverse conclusion: if you want to avoid being tracked, you must opt out of public space.

This position places urban residents in an untenable bind. On the one hand, eschewing modern technology would mean not just forgoing public announcements and conversations that occur online but also losing out on services that governments distribute by analyzing data.[42] For example, if cities analyze people's movements with MAC address sensors to determine where to place bus stops, they will overlook the needs of those without smartphones (and those who turn off their phones to avoid being tracked). On the other hand, those with smartphones and other wireless technologies must suffer the consequences of being tracked; and in places where cameras are used to identify people, it is impossible to escape being tracked even by abandoning one's personal digital technology.

Such conditions would most significantly hurt the urban poor, who are already the most vulnerable to online tracking:[43] while well-off New Yorkers who do not want LinkNYC to track them can forgo free Wi-Fi in favor of a personal data plan, lower-class residents have no alternative to free Wi-Fi (indeed, the whole point of LinkNYC is to provide internet access to those who cannot afford to pay for it) and must accept being tracked in exchange for internet access. Thus, the inevitable outcome of accepting pervasive data collection in smart cities—and believing the myth that opting out is a legitimate option—is the creation of "a new type of societal class system: a higher class of citizens free from fear of manipulation or control and a lower class of citizens who must continue to give up their privacy in order to operate within the prevailing economic system, losing their ability to control their own destiny in the process."[44]

Smart cities thus provide welfare offices, police, employers, data brokers, and others who use data to control the lives of the urban poor with a new tool for surveillance and exploitation. A single mother could be flagged by an algorithm to lose welfare benefits because she was identified at a protest by body camera footage. A black teen could be identified for surveillance by the police because he connects to a public Wi-Fi beacon that is often used by people with criminal records. An elderly man could be targeted for predatory loans because his car was recently identified by automatic license plate readers as it was driven out of an impound lot.

The grave risks to equity, autonomy, and social justice introduced by data collection in smart cities raise new challenges for and impose new responsibilities on city governments. Beyond determining what data to collect for themselves, municipalities must also act as gatekeepers for private companies eager to access new environments for data collection. Many smart city projects are similar to LinkNYC and involve public–private partnerships through which municipalities procure technology from companies in order to offer new or improved services. For governments, working with companies makes it possible to take advantage of private-sector technologies that would be difficult to develop internally. For companies, partnerships with cities present a rare and incredibly valuable opportunity to place data-gathering sensors throughout public spaces. City governments must therefore thoughtfully consider whether the benefits of new services are worth the price of allowing companies to collect untold amounts of data about the public; if not, they must find ways to obtain the benefits of new technology without incurring those costs.

But even if cities collect data for a benevolent purpose or trust the private vendor whose technology they procure, they must grapple with the numerous ways in which sensitive information can be exposed to the public or to groups with malicious intentions. Once data is collected, it is liable to be released and abused. And even within government, as the example of Los Angeles ALPR data being shared with ICE indicates, data collected by one agency can end up in the hands of another. New technology that enables the collection of more granular and sensitive information magnifies these risks.

Over the past decade, many city governments have embraced "open data" initiatives, which involve releasing municipal datasets online in an effort to make government more transparent and accountable as well as to foster civic innovation. These efforts have led to thousands of datasets being released in cities across the country, paving the way for transit apps,[45] user-friendly tools to explore municipal budgets,[46] and countless hackathons. But because much of the data that municipalities collect relates to the people within those cities, open data also occasionally reveals sensitive information about individuals. By releasing open data, cities have inadvertently revealed the identities of sexual assault victims and people who carry large sums of cash at night,[47] as well as people's medical information and political affiliation.[48] Although there are strategies that cities can employ to reduce the risks of such disclosures, they must grapple with the inevitable tension between open data's utility (more detailed data provides greater transparency and can be used for more purposes) and risks (more detailed data contains more sensitive information), a dilemma that will only worsen as the scope of municipal data collection expands.[49]

Even when governments do not proactively release data, they often have few means to protect their information from becoming public knowledge. Federal and state public records laws, designed to enforce government transparency and accountability, compel governments to release data they control when requested to do so by a member of the public. Although these laws contain exemptions restricting the release of sensitive information, their reliance on the outdated PII and secrecy frameworks severely limits the scope of such exemptions. Thus, as cities gather and store supposedly anonymous data about people's behavior, they will be sitting on increasingly large piles of information that could easily be exposed to reveal sensitive information about people. Case in point: the NYC taxi trip data that was analyzed to infer the identities of strip club patrons was initially released through a public records request and then posted online by the requestor for anyone to use.[50]

Finally (and this is a concern for companies as well as governments), any data that is collected and stored can be released through hacks and security breaches. In 2017, cyber attackers stole the names, addresses, and credit card information of 40,000 residents of Oceanside, California, which they used to make unsanctioned online purchases.[51] The previous year, a hack of Uber exposed the personal information (including names, email addresses, and phone numbers) of 57 million users.[52] Moreover, the new sensors being installed in countless Internet of Things deployments to collect detailed data about urban conditions are egregiously insecure.[53] To the security technologist Bruce Schneier, cases like these dispel the prevailing narrative that all data is good and more is always better. Instead, he says, "[D]ata is a toxic asset and saving it is dangerous."[54]

The imperative to deploy new technology thus thrusts municipalities more strongly than ever into the role of being stewards of urban life. They must decide what data can be collected and who gets to access it (while also accounting for the fact that once data has been collected it may be exposed to others). Cities are therefore confronted not merely with technical judgments about how to operate municipal services but with deeply political decisions that will determine the future of urban life. Will cities increase their control over their inhabitants and provide similar control to companies without any public dialogue? Or will they ensure that the social contract they create through technology provides people with a right to the city free from being monitored and manipulated by corporate and governmental entities?

Through this lens, what is remarkable about LinkNYC is not that a Google-related company would provide a free service in exchange for collecting user data—monetizing such data is their fundamental business model, after all—but that the City of New York would allow them to, would sell off the public's privacy for what one local paper calls "chump change."[55] As the media theorist and author of *Throwing Rocks at the Google Bus* Douglas Rushkoff puts it, LinkNYC represents "a deal with the devil we really don't need."[56]

* * *

It is not just increased data collection that threatens to create an undemocratic social contract in smart cities. As we saw in the previous chapter, cities are increasingly using algorithms to inform core functions like policing and social services. In New York City, for example, algorithms are used to assign students to schools, evaluate teachers, detect Medicaid fraud, and prevent fires.[57] And despite the seemingly sophisticated nature of these algorithms, they are neither foolproof nor neutral: bias can arise both in the training data on which they rely and in how they are deployed.

Yet even though the decisions that algorithms inform are potentially life-altering, auditing their design and impacts is remarkably difficult. Chicago's Strategic Subjects List provides an instructive example: the Chicago Police Department has resisted countless calls to publicly share the details of how the algorithm works and what attributes of people it considers.[58] Police are therefore showing up unannounced at people's homes without ever explaining what led them there.[59]

Municipal implementation of algorithms thus raises grave concerns for urban democracy: cities typically provide the public with little or no insight into how their algorithms were developed or how they work. Cities rarely release the source code governing the algorithm or the data it learns from. The public may not even know when algorithms are being used.

In many cases, municipal algorithms are concealed because they are developed and owned by private companies with a financial interest in secrecy. Because public agencies typically lack the resources and technical sophistication necessary to develop algorithms on their own, they often contract with companies to procure algorithmic systems. And although there is value in relying on technical experts to develop algorithms, these new relationships shift decision-making power away from the public eye.

Through nondisclosure agreements and broad assertions of trade secrecy,[60] technology companies prevent the governments that utilize their services from revealing any information about those tools or their use. These companies include Intrado (which has developed Beware, software that police departments use to calculate people's "threat scores") and Northpointe (which has developed COMPAS, an algorithm that predicts one's likelihood to engage in future criminal activity, recently rebranded as equivant).[61] Even public records laws have been ineffective tools for shining light onto these proprietary algorithms. For example, two lawyers submitted public records requests for information about PredPol to eleven police departments reported to be using the software. Only three responded, and none provided substantive information about the algorithm or its development.[62]

As a result, governments may make consequential decisions about people without providing any transparency regarding how those decisions are made or any due process. Consider the case of Eric Loomis, who was arrested in 2013 in the city of La Crosse, Wisconsin, for driving a car that was involved in a shooting. Loomis pled guilty to fleeing the police, and the state used Northpointe's COMPAS algorithm to inform his sentencing process. When giving Loomis a six-year sentence, the judge explained, "The risk assessment tools that have been utilized suggest that you're extremely high risk to reoffend."[63] Because Northpointe claimed that its algorithm was a trade secret, Loomis was not permitted to assess how the algorithm made this prediction. His challenge of the judge's use of this opaque system was unsuccessful.

As cases like Loomis's become commonplace, Solove's reference to *The Trial* begins to appear prescient. Just as Kafka's Josef K. faced a trial in which he knew neither his crime nor his accuser, here was Loomis being given a sentence that was influenced by an algorithm that neither he nor the judge could inspect.

The deeper danger of such algorithmic decision making is that when governments use proprietary algorithms like COMPAS, the unelected and unaccountable developers of these systems are granted significant power to dictate municipal practices and priorities. We saw in chapter 4 how algorithms—and thus their effects—are shaped by judgments about what data to use, what input factors to include, and how to balance false positives and false negatives. These seemingly technical choices influence public policy; as governments increasingly make choices based on algorithms developed by private companies, they will increasingly make decisions based on the values and assumptions that those companies embed within the algorithms. For example, Northpointe's choice to predict people's likelihood to commit crimes in the future places criminal justice adjudication within the prosecutorial and racialized context of crime risk.[64]

One of the most important decisions that Northpointe made when developing COMPAS was how to ensure that its predictions were not racially biased. On even a purely technical level, making unbiased predictions is more complicated than it may first appear, as there are several technical criteria for "fairness" from which to choose. The company strove for what is known as "calibrated predictions," meaning that their model should be equally accurate for both black and white defendants. This is, on its face, a sensible choice. Yet in 2016, ProPublica revealed that COMPAS was twice as likely to misclassify a black person than a white person as "high risk," potentially leading to criminal sentences for black defendants that, without justification, are longer and more punitive.[65] This was seen as evidence that COMPAS was racially biased. Perhaps Northpointe should have optimized instead for "balanced classes," meaning that its algorithm would have equal false positive rates for black and white defendants. Doing so would have addressed the issue that ProPublica highlighted, but at the price of raising a new one in its stead: the new algorithm would no longer make calibrated predictions (i.e., it would be more accurate for one group than for another). In any attempt to make fair predictions about two groups for whom a phenomenon of interest—here, recidivism rates—occurs at unequal rates, this trade-off in unavoidable: it is impossible to attain both calibrated predictions and balanced classes.[66]

The point is not that COMPAS was wrong in prioritizing calibrated predictions: neither of its options for defining fairness was clearly superior to the other. Indeed, many policy decisions involve complex trade-offs of this sort. The problem is that this decision—which shapes a fundamental aspect of the criminal justice system, and hence people's lives, in every jurisdiction that adopts COMPAS—was made by the staff at Northpointe with no input from public officials or the broader populous.

Reliance on municipal algorithms thus represents a drastic change in how policies are developed and implemented. In the past, although such decisions were by no means entirely transparent and accountable, they were presumed to be political and to require democratic input, oversight, and justifications. Decisions made by computational systems (especially ones developed by private companies) eschew these obligations: it is in many cases impossible for the public to have any input into or wield any control over algorithmic decisions, if they know about the algorithms at all. And even when they do assert opinions, questions about how algorithms should be designed are often seen as technical matters best left to the “experts.”

Compounding the dangers that emerge as governments use opaque and unaccountable algorithms is the vastly expanded data collection that smart city technologies enable. As previously unattainable data about individuals becomes available, much of that information may be incorporated into algorithms that influence criminal sentences and other significant decisions. In the court cases of the smart city, factors such as where you spend time, how late you stay out at night, and whether you participated in a particular protest—data that you may never have even known was being collected, and whose collection you certainly never consented to—could influence the sentence you are given.

As these manifold dangers come to light, one city has made a valiant attempt to alter how it deploys algorithms. In August 2017, James Vacca, a member of the New York City Council, introduced a bill that would require city agencies to release the source code of every algorithm they use to target social services, impose penalties, and inform the actions of police.[67]

Vacca, who had been involved in New York City government for almost four decades, was familiar with the public’s lack of access to algorithms: his many attempts over the years to learn about the algorithms that dictate staffing in the police and fire departments had all been thwarted.[68] In a public hearing on his proposed legislation, Vacca explained his motivation for the bill. “I strongly believe the public has a right to know when decisions are made using algorithms, and they have a right to know how these decisions are made,” he said. “For instance, when the Department of Education uses an algorithm to assign children to different high schools and a child is assigned to their sixth choice, they and their family have a right to know how that algorithm determined that their child would get their sixth choice. They should not merely be told that they were assigned to a school because an algorithm made the most efficient allocation of school seats. What is considered to be most efficient? Who decided this?”[69]

The final version of Vacca’s bill, approved by the City Council in December 2017 and signed by Mayor de Blasio in January 2018, represents a curtailed enactment of his initial vision. The legislation established an "automated decision systems task force" to examine how city agencies use algorithms. The group will recommend procedures that would achieve several outcomes: increase public transparency about the algorithms being used, determine which algorithms should be subject to oversight, provide the public with opportunities to receive explanations about algorithmic decisions, and evaluate whether algorithms unfairly impact certain groups of people. The task force will present these recommendations to the mayor in a public report.[70]

Although the new law leaves much to be desired—the task force can merely make recommendations and has little power to compel public disclosure, especially when dealing with companies eager to protect their trade secrets[71]—it is a productive start, building momentum for policies that hold cities accountable for how they develop and deploy algorithms. Just as importantly, Vacca's efforts helped shift public discourse toward considering algorithms not as unassailable oracles but as socially constructed and fallible inputs to political decisions. Such a change is essential to the development of Smart Enough Cities.

* * *

When it comes to algorithms as well as data collection, municipal decisions must be grounded in democratic deliberation that provides the public with a meaningful voice to shape development, acquisition, and deployment. Such work will, perhaps counterintuitively, aid rather than hinder the adoption of technologies that improve life in Smart Enough Cities.

The Array of Things (AoT) in Chicago highlights the value of engaging the public to protect privacy. The product of a collaboration (launched in 2014) between the City of Chicago, the University of Chicago, and Argonne National Laboratory, the AoT is designed to be "a 'fitness tracker' for the city."[72] It will eventually consist of several hundred sensors installed throughout Chicago to track environmental conditions such as air quality, pedestrian and vehicle traffic, and temperature. In one neighborhood that has a highway running through it, for example, city officials hope to reduce asthma rates in children by using the data as a guide for where trees are most needed and where bus stops should be located.[73]

On the surface, the Array of Things appears quite similar to LinkNYC: it is a large-scale deployment of sensors that can collect vast quantities of data. The AoT could have led to public backlash if too few people trusted that the data was being collected and managed responsibly. Chicago took a vastly different approach from New York in deploying its sensor network, however: it designed the sensors to avoid collecting sensitive personal data while also directly engaging the public to share how it had done so and to collectively develop priorities. This is, in part, a benefit of how the Array of Things is owned and operated: whereas LinkNYC is managed by a private company and thus structured to maximize profit, the AoT is run by government and academic institutions and therefore focused on generating public benefits rather than revenue.

Chicago has made protecting public privacy an essential element in the AoT's architecture. While developing initial plans for deploying the AoT, the city convened a committee of local privacy and security experts to provide independent oversight regarding how the system collects and stores data. It then organized public meetings to explain to the broader Chicago population how the Array of Things worked, how the program protected privacy, and how the sensors could improve living conditions.[74] To fully incorporate public concerns about privacy into the policies that govern the AoT, Chicago also released a draft of its privacy policy for public comment. This draft garnered more than fifty inquiries, all of which the city publicly responded to and incorporated into the final privacy policy that governs the program.[75]

This outreach helped Chicago identify the public's privacy concerns and hold itself accountable to addressing them. For example, one major worry was that the sensor's cameras would collect images that could be used to track people's movements over time. Such tracking was not the city's intention—it sought to collect the images to obtain traffic counts (via analysis from computer vision software)—but it was certainly possible that images gathered from the cameras could be abused. Collecting that footage could violate privacy expectations and fuel residents' opposition to the entire AoT project.

In response, Chicago devised a thoughtful solution that draws on a practice known as "data minimization," which involves collecting and storing the minimal amount of information needed to achieve a project's goals. Data minimization can take several forms; two common tactics are ignoring extraneous features entirely (say, not collecting someone's location) and storing data in a deliberately imprecise format (recording someone's location simply as a zip code).[76] Because Chicago desired only the traffic counts that could be calculated from the images, there was no need to store the camera footage itself. The AoT sensors were altered to calculate this number, transmit only that value to the project servers for retention, and then immediately delete the images.[77]

Chicago's development of the AoT exemplifies how cities can synthesize public engagement and data minimization to make possible the deployment of cutting-edge technology in Smart Enough Cities. By allowing public input into how it implemented the Array of Things, Chicago ensured that the social conditions mediated by its new technology would be democratically determined and desired. When the public raised concerns, the AoT team found a way to collect the data that it required and no more, in such a way that the city could achieve its analytical goals without compromising the will or privacy of its constituents. Had it not followed these practices, the entire project might have been stymied by public opposition.

During the same years that Chicago was developing the Array of Things, Seattle was learning the hard way why it is necessary to protect privacy when deploying new technology. In November 2013, backed by $2.7 million from the Department of Homeland Security, the City of Seattle installed a network of sensors and cameras to monitor potential threats coming from its harbor. The public—who had received little warning about the new technology's installation or use—quickly became concerned when they noticed that the wireless sensors appeared capable of tracking individuals by recording the movements of their wireless devices. In response to questions about how the new sensors would be used, the Seattle Police Department replied that it was "not comfortable answering policy questions when we do not yet have a policy," which further inflamed tensions by suggesting that the city was being cavalier about individual privacy and not taking the necessary precautions to prevent undue surveillance.[78] As the public controversy mounted, even despite indications that the sensors could not in fact track people's devices, the police department shut down the program.[79]

The botched rollout "became a really good learning lesson," reflects Seattle Chief Technology Officer Michael Mattmiller.[80] In part because of a lack of familiarity with the technology and its risks, the city had deployed these new sensors without considering whether their architecture aligned with public priorities. "For those who aren't educated on how technologies work and potential privacy harms, it's very easy to just focus on the outcomes of a technology, not the means by which they achieve those outcomes," Mattmiller notes. Moreover, deploying technology with meager public outreach meant that the city had a poor understanding of how members of the public value their privacy. As existing privacy laws based on PII become obsolete, social conceptions of privacy evolve, and new technology makes it possible to collect ever-increasing amounts of data, cities have no reliable guide to tell them what level of data collection is appropriate. So as Seattle shut down a new technology program over public concerns about invasive data collection, the message became clear: without technological expertise, strong privacy policies, and public dialogue to address privacy risks, it would be almost impossible to improve municipal operations and urban life with technologies that involve collecting data.

Eager "to move forward on these missteps we've had with the public," Mattmiller created a Privacy Advisory Committee consisting of local technologists, lawyers, and community representatives. The group's task was to share public priorities and concerns related to data collection and to guide the city in developing practices that protect privacy. Through a series of public meetings, the committee developed six Privacy Principles—collectively affirming a commitment to transparency, accountability, and minimal data collection—intended to guide the city's collection and use of personal information.[81] The committee then helped the city develop a thorough privacy policy in 2015 that embodies these principles.[82] Through these efforts, says Mattmiller, the Privacy Advisory Committee "really ensured that the community's fingerprints and best practices were infused in our privacy program."

A key component of Seattle's privacy policy is a mandate to conduct a "Privacy Impact Assessment" every time it develops a new project that involves collecting personal data. The city must undertake a risk-benefit analysis that weighs the project's potential utility against its potential threat to individuals' privacy. The goal is to proactively highlight and mitigate expected risks without crippling the project, thereby enabling the city to balance its charge to enhance public welfare with its responsibility to protect civil liberties. These assessments help Seattle ensure that their projects adhere to the Privacy Principles; adjustments typically take the form of changing how data is gathered, stored, and shared to reduce the collection and exposure of sensitive information.

Mattmiller also emphasizes the importance of educating city staff about privacy risks and how to mitigate them. To help departments recognize and prevent privacy harms, the city nominated a member of every municipal department to act as a "privacy champion." These individuals coordinate Privacy Impact Assessments and educate their colleagues about best practices in following the city's Privacy Principles. When Mattmiller and his team need to update city staff about recent developments or new privacy risks, the privacy champions help disseminate this information throughout departments. And in 2017, Seattle further institutionalized its commitment to privacy by hiring a chief privacy officer, becoming one of the first cities in the nation to have someone with a citywide mandate to manage privacy.[83]

Adopting these practices meant that the next time Seattle prepared to adopt a new technology that required collecting the public's data, it was ready to do so thoughtfully and responsibly. The big test came when the Seattle Department of Transportation (SDOT) deployed a thousand sensors to measure traffic flow and travel time between various locations. By tracking the movement of wireless devices across town (via their MAC addresses)—data that had never before been available to public officials—SDOT hoped to spot patterns that would help it reduce congestion.

With their new privacy program in place, Seattle was able to deliberatively and openly weigh the costs and benefits of this technology. First, Mattmiller and his team consulted with SDOT and the technology vendor to ensure that excessive surveillance would not be the price of smoother traffic. After identifying several privacy risks in the technology, they pushed the company to implement data minimization approaches that would make it harder to identify and track any individuals using the data. And instead of working behind the scenes, as it had when deploying the security network around the harbor, the city proactively shared what it was doing and why. Mattmiller took the case for this technology to the public, explaining, "If you like using Google Maps and you like seeing the red, yellow, and green lines that show you traffic flow and reroute you around congestion, then we need to collect data to inform those maps, and we believe this is the least invasive way possible. If you don't like this, here's a website you can go to to opt out, so if we see your cell phone we ignore it."

By mitigating the most salient privacy risks and explaining the purpose of its efforts, the city garnered public support rather than outrage, Mattmiller notes. "Very plainly conveying the value of what you're doing, being transparent with privacy threats, and showing how you've mitigated those threats builds trust."

But it is not enough for cities to act in good faith—public oversight of municipal technology must be institutionalized. To further empower the public with control over how the city collects and uses data, Seattle enacted a surveillance oversight ordinance in 2017. The bill requires every department to hold public meetings and obtain City Council approval before acquiring any surveillance technology, to publicly describe how it uses surveillance technology, and to assess all surveillance technology for its impacts on privacy and equity.[84] The ordinance thus ensures that Seattle's decisions to acquire and deploy surveillance technology (whether hardware or software) are subject to robust deliberation by the public and elected officials rather than the shrouded decisions that are endemic across U.S. cities. In 2016, for instance, a local newspaper reported that the Seattle Police Department had been using social media monitoring software for two years without ever notifying the public.[85] Similarly, police in New Orleans used predictive policing algorithms for several years without having gone through any public procurement process, and even members of the City Council were left in the dark.[86] Joining Seattle in pushing back against this trend, dozens of cities across the United States (ranging from Hattiesburg, Mississippi, to Oakland) have passed or are developing similar surveillance ordinances,[87] building momentum for one of the most important tactics for transforming smart cities into Smart Enough Cities.

By recognizing that seemingly technical decisions about what data to collect are in fact political ones with massive implications for civil liberties and social justice, Chicago and Seattle both demonstrate how Smart Enough Cities can provide new services and enhance daily life with technology while at the same time fostering democratic social contracts. These achievements contradict the view through tech goggles, which falsely suggest a binary trade-off between privacy and innovation. According to this worldview, becoming smart means collecting and analyzing data to improve efficiency—and if protecting privacy and liberty requires gathering less data, then the smart city must be one without privacy or liberty.

In the Smart Enough City, however, where privacy is a human right essential to maintaining liberty and equity, protecting privacy enables rather than prohibits the use of new technology. "If you have a well-resourced, well-functioning privacy program, that program will promote innovation . . . and allow agencies to expand to new technologies," explains Marc Groman, senior advisor for privacy at the White House Office of Management and Budget under President Obama.[88]

While the smart city collects as much data as possible as it chases maximum efficiency, and the proverbial dumb city collects no data, the Smart Enough City collects data only after it has won public support to do so and after establishing privacy-protecting policies. The question for the Smart Enough City is not simply "What data should we collect?" or "How much data should we collect?" but "How can we accomplish our policy goals, with the aid of data, without violating public expectations or rights?"

* * *

In 2014, Nigel Jacob, co-founder and co-chair of the Mayor's Office of New Urban Mechanics in Boston, got a call from a local engineer. "I've been studying this whole parking problem, and it's actually very simple!" he excitedly told Jacob. The engineer had developed a clever solution to minimize the frustration of searching for on-street parking (and the congestion that this search creates): an app to pay for and reserve parking spots. Request a space before you leave the house, and a metal bollard would pop up and hold the space until you arrive. "It's just a question of resource allocation," the engineer explained.[89]

Whether or not this could work in practice—that is, make it easier to find parking—was beside the point, says Jacob. "We had this discussion where we explained that no one person has a right to public space. There's a social contract there. We started talking about that, and he started to see that it's a lot murkier." If increasing parking efficiency meant giving individuals the ability to reserve public space, Jacob was not interested.

The engineer in Jacob's story is typical of many technologists: they emphasize efficiency and convenience for some without considering the way in which those goals are achieved. Efficiency as an end seems to justify any means, or simply makes the means (and its by-products) irrelevant. And all too often, Jacob admits, cities buy into this logic and fail to consider the broader impacts. "We have a long track record of buying the wrong technology for a particular problem because we don't think about the politics of particular architectures," he laments.

When municipalities do not consider how the technology they acquire actually operates, the companies behind that technology dictate its architecture. As cities then deploy that technology, these design choices influence the social contract between people, companies, and the government. Under the guise of making a city "smart," companies sell technology that collects sensitive data and is opaque to the public—attributes that increase their profits—as if that were the only possible way that their products can function. Governments often then deploy that technology without sharing any details about it.

These are not inevitable outcomes dictated by the demands of new technology—they are instead the political arrangements desired by those who develop and control that technology. But as Chicago and Seattle demonstrate, alternative and more democratic architectures are available: it is possible to deploy pervasive sensors that improve urban life without collecting and abusing vast quantities of information about people. Similarly, the creation of an algorithm task force in New York and the enactment of surveillance oversight ordinances in Seattle and other cities provide a clear path for municipalities to reverse the trend toward becoming black-box cities.

Municipalities must accept the power and responsibility bequeathed to them as gatekeepers to information about the public and stewards of the public's privacy. Rather than merely embracing every new technology as manna from technological heaven, Smart Enough Cities are compelled by this new role to consider the risks of potential technological designs and reject the architectures motivated by police surveillance and extractive corporate practices.

When procuring technology from companies, municipal leaders must look beyond what a tool can accomplish and use their leverage to negotiate more democratic policies regarding privacy and transparency. After all, technology companies need cities more than cities need technology companies. Municipalities may gain some knowledge and efficiency from new tools and software, but we know by now that efficiency is not the most important ingredient for thriving urbanism; companies, on the other hand, need someone to buy their products. Given these dynamics, cities have the opportunity to assert themselves as market makers, acting both individually and collectively to shape the direction of smart city technology. Recognizing this power, in 2017 a coalition of 21 chief data officers published a set of guidelines for companies developing open data portals, and 50 mayors submitted a joint letter to the Federal Communications Commission in support of net neutrality.[90] Barcelona is also a notable pioneer in this regard: it has restructured contracts with several major technology vendors to enhance the public's ownership and control of data.[91]

If city governments do not take these actions, technology companies may continue to gain opaque and unaccountable private power over urban life. Already, companies such as Uber and Sidewalk Labs possess far more data than municipalities do about urban conditions, and companies such as Northpointe develop algorithms that inform highly consequential decisions. And as smart city companies accrue additional investment and profit, they will gain further leverage over cash-strapped municipalities: part of the reason smart city initiatives like LinkNYC are appealing to cities is that the latter lack the resources to provide public services themselves. If twenty-first-century urban residents are to have a right to the Smart Enough City with meaningful democratic control over technology production and use, municipalities must assert their authority over companies and be provided with the resources necessary to do so, while also themselves becoming more democratic.

Rushing to become a smart city may lead to new insights and efficiencies, but at the cost of creating cities in which the government and companies wield immense power over individuals by collecting data and making opaque decisions about them; poor and minority residents will be most subjugated. Yes, among the many responsibilities of cities are providing effective services and spending their funds judiciously. But to recklessly pursue technology that advances these goals without considering its full impacts is a severe dereliction of duty. For as we have seen, the benefits of new technology are often illusory and deploying it unthoughtfully can create more problems than it solves.

Cities can tread cautiously around technology without neglecting their responsibilities to care for the public, however. As we will see in the next chapter, technology can play a vital role in municipal innovations to improve urban welfare—but only when grounded in meaningful institutional and policy reforms that guide it toward the desired outcomes.