

6 The Innovative City: The Relationship between Technical and Nontechnical Change in City Government

One of the smart city's most alluring features is its promise of innovation: it uses cutting-edge technology to transform municipal operations. Like efficiency, innovation possesses a nebulous appeal of being both neutral and optimal that is difficult to oppose. After all, who would want her city to stagnate rather than innovate?

Consider the homepage of Sidewalk Labs, which (as of October 2018) uses the word “innovation” five times. The company promises that it is “investing in innovation,” will “accelerate urban innovation,” provides “infrastructure that inspires innovation,” and will “make Toronto [the site of its most ambitious project; see chapter 7] the global hub for urban innovation.”¹ Elsewhere, the company has declared that “our mission is to accelerate the process of urban innovation.”² Even more than technology, it appears that innovation is Sidewalk's key product. In this sense, innovation is of a piece with other smart city buzzwords like “optimization” and “efficiency”: a vague but supposedly neutral and beneficial goal that is often touted by companies to advance their corporate agenda.

There is little doubt that cities could benefit from new ideas, policies, practices, and tools. Where smart city proponents like Sidewalk go astray, however, is in equating innovation with technology—or, to use Sidewalk's language, in concluding that “reimagining cities to improve quality of life” requires “digital advances to transform the urban environment.”³

We will see in this chapter just how misguided that perspective is. It is wrong not only because technology alone cannot solve intractable social and political problems but also because of an attribute of city governments that we have observed but not yet fully explored: to derive benefit from technology, they must overcome institutional barriers by reforming policies

and practices. This chapter will provide case studies of several cities—most notably, New York City, San Francisco, and Seattle—that demonstrate the painstaking processes required to improve governance and urban life with data. We will observe a very different relationship between technology and innovation than technophiles would ever recognize or praise.

* * *

In July 2015, public health officials in New York City identified an outbreak of Legionnaires' disease (an acute form of pneumonia) in the South Bronx. Seven people had already died and dozens more were infected. If not addressed immediately, the illness could spread throughout the Bronx and across New York City, threatening the well-being of millions.

The city's Department of Health and Mental Hygiene (DOHMH) quickly determined that the disease-inducing *Legionella* bacteria were incubating in the cooling towers that sit atop large buildings to support their air-conditioning systems. This is a common source for Legionnaires', especially during the summer when the use of air conditioning increases. As DOHMH cleaned the contaminated cooling towers, its staff recognized that a city-wide inspection effort was necessary to prevent the disease from incubating in others. The City Council mandated that the city form a tactical response team to rapidly register and clean every cooling tower.

In many respects, this was nothing new for the most populous city in the United States. Led by NYC Emergency Management (NYCEM), city departments are adept at coordinating responses to crises ranging from hurricanes to terrorist attacks to citywide blackouts. But in the case of Legionnaires' disease, coordinating agencies was not enough—the city also had to coordinate multiple sources of data. Several critical questions loomed over the response effort: How many cooling towers are there in New York City? Where are they? Who owns them? Which ones are incubating *Legionella*? These questions could not be immediately or easily answered. Only a small fraction of buildings have cooling towers, and the city lacked any comprehensive database of their locations or owners.

So on a Friday afternoon a week into the emerging crisis, the Mayor's Office called Amen Ra Mashariki—New York City's chief analytics officer—asking for help.

"You can imagine this was an emergency of epic proportions," Mashariki says, looking back. "At the core of what we're supposed to do as government

agencies is protect New Yorkers.” Failing to quickly identify and inspect every cooling tower could allow Legionnaires’ to spread out of control. Mashariki adds that what made this emergency particularly daunting is that addressing it required “a dataset that no one has ever considered. No one in City Hall or the Department of Buildings wakes up in the morning saying, ‘We need to make sure our cooling tower dataset is primo because there may be an emergency that involves cooling towers.’ This is a dataset that virtually didn’t exist, and we had to cull it together.”⁴

Luckily for New Yorkers, Mashariki’s unique personal and professional background had prepared him for this moment. Growing up in a middle-class family in Brooklyn, Mashariki was strongly influenced by both of his parents. His father was a Vietnam War veteran and a social activist who founded a nonprofit to assist other veterans. His mother was a human resources executive at IBM, a position that afforded Mashariki access to some of the first PCs ever made. As a child, he was obsessed with computers and video games—he could not wait to learn how to program Donkey Kong—and his mother made a habit of bringing him to the office during school vacations. She taught herself how to code in BASIC (an early programming language) and began teaching her son when he was in fourth grade.

After studying computer science in college, Mashariki took a highly coveted job at Motorola in Chicago. He was building a successful career there, developing security protocols for two-way radios, when the Twin Towers were hit on September 11, 2001. The next day, when work resumed and the office was functioning as usual, Mashariki began to question how his work affected the world. “If something happens that changes the world but my job doesn’t change, then the corollary must be true, which is my job doesn’t have an impact on the world,” he concluded.

Mashariki may have been developing technology that would help lead to the smartphones we all rely on today, but developing cutting-edge technology was no longer fulfilling for him. Following the lead of his activist father, Mashariki decided on that day that “anything I do from here on out has to explicitly have impact.”

After spending most of the next decade in medicine, creating software for surgical robots and analyzing cancer treatment data, Mashariki plunged into government in 2012 as a White House Fellow in the Office of Personnel Management (OPM). The first ever computer scientist to be a White House Fellow, Mashariki entered with swashbuckling confidence that his

technical expertise would help solve all the government's problems. "When I came in, I was like, 'I'm going to be the hottest thing,'" he recalls, cringing at the memory of one speech in which he announced his intention to use algorithms "to fix how you guys do things and blow up the way you're thinking about problems.' I thought for sure I was going to come in here and be a superhero," he says. "And I remember looking around like, 'Why aren't they really digging this?'"

Mashariki entered government, in other words, like a typical technologist: confident that cutting-edge technology was the solution to many of government's challenges and that providing technical expertise would make him a savior. But his early efforts at OPM floundered because the solutions and approaches he espoused were poorly suited to the agency's needs. He was too focused on wielding technology in every situation rather than understanding the problem.

"Needless to say, I got my ass handed to me so many times, in many different ways," Mashariki recalls with a laugh. Whenever he suggested a technology that he thought would provide a quick and obvious fix, Mashariki was shot down because his colleagues had already considered that technology and determined that it would not address their needs.

These experiences helped Mashariki remove his tech goggles and realize that solving government issues with technology was much more difficult and complex than it had initially appeared. He realized that the issues he had diagnosed as technology problems were actually related to organizational capacities and needs, and that the key to addressing them was working with people and institutions rather than building technology. Mashariki also saw how bureaucracy, so often maligned as the force that stops innovation, prevented the implementation of bad ideas. Contrary to his expectations, working within the system was more productive than blowing it up. Mashariki's preconceived skepticism of government faded, leaving him with "a high level of respect for public servants."

Mashariki was named OPM's chief technology officer in 2013 and was handed responsibility for the massive project of digitizing the federal government's retirement process. Whereas a year earlier Mashariki would have concentrated on identifying the best software and persuading his colleagues to adopt it, now he recognized that success required bringing people together and focusing on institutional needs. He lists the many factors that he needed to consider: "You have to build relationships. You have to build

consensus. You have to identify the people that you have to influence. You have to identify the people that you have to get insight from.” Mashariki also knew not to dismiss the expertise of others within the organization. In the face of extensive doubts about the project from his colleagues, he built trust throughout OPM by emphasizing, “We’re not here to tell you how to do your job. We’re here to help you, learn how you do your job, and provide some capability for you.” Mashariki’s people-first approach was highly successful: over six months, his team achieved more progress than others had in OPM in the previous fifteen years.

Mashariki left OPM in 2014 to become New York City’s chief analytics officer and director of the Mayor’s Office of Data Analytics (MODA), the fledgling municipal analytics shop that the city had established the previous year. Having learned at OPM about the limits and potential of technology to improve government, Mashariki was eager for “the challenge of being the data guy for one of the largest cities in the world. Who wouldn’t want to take that on?” He knew it would be the most demanding role of his life, but, Mashariki says now, “I had no idea just how complex it was really going to be.”

Mashariki was only nine months into the job when Legionnaires’ disease broke out. The scale of the task and the precision required to complete it were overwhelming. New York has more than a million buildings; given the city’s limited human and financial resources, visiting each one to check for a cooling tower would take years, allowing the bacteria to fester and spread. But the city had to be comprehensive in its search. “This can’t be 98 percent confidence that we got all the buildings,” Mashariki explains. “This has to be 100 percent confidence.” Mashariki’s job was to accelerate the pace of inspections by using data and analytics to identify the buildings most likely to contain a cooling tower, and thus the buildings on which the inspection and cleaning teams should focus.

Unfortunately, synthesizing all of the city’s data to form a coherent picture proved more difficult than anyone had predicted. For instance, it initially appeared that the Department of Finance (DOF) possessed the necessary data, since it tracked some buildings with cooling towers as receiving a tax write-off. But this dataset did not include every cooling tower, nor did it contain the names and contact information of building owners—information that was needed to verify the presence of a cooling tower, register it, and inspect it. And while the Department of Buildings (DOB) does

collect information about building owners, it was incredibly difficult to align the two datasets, because DOB identifies buildings by address whereas DOF does so by tax parcel. In addition, DOB recorded the number of buildings that use a cooling tower, overlooking that some cooling towers service multiple buildings and some buildings are serviced by multiple cooling towers. MODA's first responsibility was to synthesize these conflicting and incomplete datasets, but despite their painstaking work the team could piece together only an incomplete list of cooling towers and their owners.

These gaps and disparities in data are common in city governments: although many departments collect nominally related data, each typically interprets and documents that information differently. Datasets collected by different departments are rarely designed to be synthesized. Every administrative division has its own IT systems and data structures, which are tailored to its individual needs and missions. This facilitates everyday tasks but hinders efforts that require merging data from multiple departments.

"A lot of people don't realize that there are different ways to count entities in the city," Mashariki explains. "Oftentimes you think you're counting the same thing but these two agencies are counting different things and they report it to the leadership of the city two different ways. If you don't have a team like MODA there, then it can be mayhem."

There was no room for such mistakes in this crisis, which required many sources of data. DOB created a website where building owners could register their cooling towers. The city's 311 call center contacted building owners to ask if they had a cooling tower. NYC Emergency Management canvassed the city as part of a public awareness campaign. Firefighters traversed the city, inspecting buildings to see if they had a cooling tower. The Department of Health and Mental Hygiene tested and cleaned the cooling towers that were identified.

The Mayor's Office of Data Analytics became the glue holding these rapid-fire efforts together. Every morning at 7 a.m., MODA would tell each agency where its resources for outreach or inspections were most needed that day. Departments would spend the day working on these tasks, recording data along the way. By 11 p.m., MODA would receive reports about each agency's progress—at which point it would assess the response effort's progress and determine each agency's tasks for the next day. Mashariki and his team became accustomed to sleepless nights.

MODA's next step was to synthesize these disparate and imperfect streams of information to quickly yet accurately identify every cooling tower in New York City. Early in the crisis response effort, only 10 percent of buildings visited by the city's inspection and cleaning teams had cooling towers—the outreach effort was wasting an enormous amount of time. Given such a low hit rate and the more than one million buildings in New York City, it could take years to find every cooling tower. To speed up the effort, MODA began developing machine learning algorithms that identified which buildings were most likely to have cooling towers by comparing their characteristics to those of buildings already identified as having cooling towers.

Despite the advanced data analysis required, MODA could not succeed by treating this as a purely technical challenge. Fortunately, Mashariki and his team were collaborating with other municipal agencies rather than focusing only on optimizing their algorithm. MODA's first list of potential cooling tower locations contained 70,000 buildings—a good start, but still too many buildings to inspect if they hoped to win the race to prevent more people from becoming ill, or worse. While reviewing this list, however, a few firefighters picked up on a key detail that the analytics had missed: the local fire code prohibited cooling towers on buildings with fewer than seven stories. When MODA incorporated this information into its algorithm, the list of potential cooling tower locations was cut in half.

“Your machine learning algorithm would not know stuff like that,” Mashariki explains. “We would have been probably futzing around with a larger dataset if those folks didn't say, ‘No, you don't have to go to those buildings.’” Whereas his younger self would have expected to save the day with a sophisticated algorithm, by this point in his career Mashariki understood that data and technology cannot solve every problem on their own. So even at a time when the city needed accurate and precise data to save lives, he reached out beyond the realm of databases and analytics to access as much contextual knowledge as possible. “You come in with your fancy machine learning algorithm in your pocket,” Mashariki observes, “but what's always going to be your ace in the hole is the knowledge of the people who actually do the real work.”

With the aid of contextual knowledge from other agencies, MODA's machine learning algorithm identified cooling towers with 80 percent

accuracy—eight times the hit rate achieved by the city before incorporating analytics. The algorithm provided the city with the guide that it needed to identify, inspect, and clean every cooling tower in NYC within several weeks, stopping the outbreak by mid-August. The toll was significant—with 138 illnesses and 16 fatalities, this was the largest Legionnaires' outbreak in New York City's history⁵—but it would have been far worse without the efficient response effort that MODA made possible.

* * *

The Legionnaires' outbreak was a “game changer,” according to Mashariki. The challenges experienced during the response effort highlighted major gaps in data quality and utility that could paralyze NYC in future crises. The next time an emergency arose—and Mashariki knew that there would be a next time—the city would need to respond more effectively and efficiently. The fire department might not have the spare capacity to traverse the city collecting data. A day spent reconciling discontinuities between datasets could slow the response effort and allow a crisis to intensify.

Because Mashariki knew that it would be impossible to precisely predict what the next emergency would be and what information would be essential—factors that he calls the “unknown unknowns”—he realized that the city needed to do more than just clean up a particular dataset or collect a specific new type of information. Instead, municipal departments across New York City would have to improve their data infrastructure and to develop generalized data skills so that they could better access, interpret, and use data for any purpose.

Mashariki adapted a page out of the city's existing playbook. One of NYC Emergency Management's responsibilities is to conduct emergency drills (akin to fire drills, but for municipal crises), in which multiple agencies practice responding to emergencies such as heat waves, coastal storms, and blizzards. These exercises are low-risk opportunities to identify gaps in services and coordination so that city agencies are prepared to act and work together when real crises occur. Following NYCEM's lead, Mashariki developed a similar training mechanism, called “data drills,” during which departments could practice sharing data and using analytics to support the municipal response to an emergency.

The first data drill, in June 2016, brought together a dozen agencies to address a hypothetical blackout in Brooklyn. Every elevator in the area was

shut down, leaving people stuck inside buildings and in need of rescue. City departments were required to synthesize data from multiple agencies to determine the location of every elevator in the region, predict which buildings had a population that was likely to be injured in those elevators, and develop a dispatch strategy to quickly send emergency response vehicles to those locations. The next drill, a few months later, involved the aftermath of a coastal storm and prompted agencies to assess the damage by integrating new data from post-emergency inspections with existing databases. The third data drill emphasized data sharing, enabling municipal agencies to practice how they access and use data from different departments during fast-paced crises.⁶

These drills are essential because, as MODA sees it, departments will be unable to improve operations and life in New York using data until they can manage and understand the pertinent information. By creating opportunities to work with data across a variety of situations, data drills push NYC municipal staff toward more effective and impactful uses of data. Departments have learned what information is collected by other agencies and how to prepare their own data so that other agencies can use it. To further aid these efforts, MODA is developing technical tools that make data easier to interpret and access. The team's first major project is a comprehensive Building Intelligence tool kit that unifies seven agencies' data about buildings into one interactive system, freeing departments from the burden of having to painstakingly make sense of conflicting information about buildings from different agencies. Data drills have also helped departments become more skilled in analyzing and applying data, whether to handle emergencies or to improve daily operations. And as these practices, processes, and tools permeate City Hall, they enable MODA to help departments serve New Yorkers more effectively. In one project, for instance, the team used machine learning to help the Department of Housing Preservation and Development proactively prevent landlords from harassing and forcing tenants out of rent-controlled apartments.

Mashariki's data drills exemplify how a city can become Smart Enough and illustrate the benefits of doing so. It is impossible to predict precisely what data and algorithms will be needed in the future—cities are too complex. But we can certainly predict the types of problems that will arise and the challenges that will accompany the use of data: poorly managed datasets that are inaccurate or incomplete, a lack of data fluency throughout

Project types

MODA works with agencies to form analytics questions in line with their overall mission and specific goals. While MODA strives toward operational analytics in every project, most begin with a data management or decision support question.

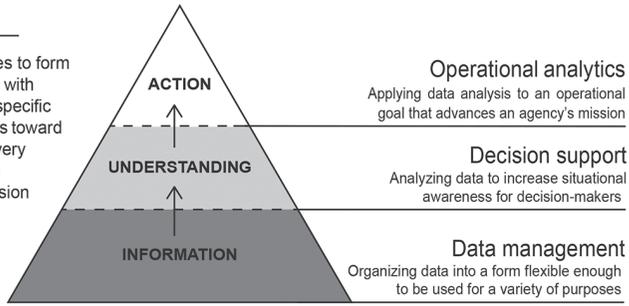


Figure 6.1

The project pyramid that New York City's Mayor's Office of Data Analytics uses to guide its strategy.

Source: NYC Analytics, "Mayor's Office of Data Analytics (MODA)," p. 1, <http://www1.nyc.gov/assets/analytics/downloads/pdf/MODA-project-process.pdf>.

departments, and the inability to synthesize information across datasets. These are not at their core technology problems, but to use new technology effectively they must be addressed.

As we similarly observed in Johnson County (see chapter 4), these issues typically arise because municipal departments and agencies operate as largely independent entities: each collects data appropriate for its particular tasks and responsibilities, without considering what data another department has. Two agencies might monitor the same aspect of the city but record information in a way that makes it difficult to match records. And because data has not traditionally been seen as a resource for analysis above and beyond its immediate purpose, there has been little reason historically to enforce data quality standards or uncover datasets that are tucked away on each department's computers. Moreover, municipal staff, who typically lack training in data analysis, are often wary of external attempts to improve or alter their work with technology.

In San Francisco, Chief Data Officer Joy Bonaguro is charged with overcoming these obstacles to make data a more valuable resource. With little patience for buzzwords and a no-frills attitude, Bonaguro effectively resists the allure of tech goggles. This resistance is necessary because, although her ultimate goal is to help the City of San Francisco use technology more effectively, the challenges Bonaguro faces are related primarily to people and policy.

Bonaguro's background in design makes her particularly attentive to the needs of technology users rather than to the capabilities of technology. And as a self-described "*Harvard Business Review* junkie," Bonaguro is comfortable dealing with complex bureaucracies such as city government. These perspectives help Bonaguro focus on using data to improve her city without getting caught up in the hype. "Smart cities are very technology-centered and technology-driven, and that's almost never a good strategy," Bonaguro says. "The reason that we're doing data science is not so we can be cool. We want to demonstrate that this is a tool we should be using."⁷

Since becoming chief data officer in 2014, Bonaguro—along with her team, DataSF—has had the mission of systematizing effective data infrastructure and governance throughout City Hall. They began by asking every department to create and share a data inventory, requiring them to catalog every data source and dataset they managed. By March 2015, thirty-six out of fifty-two departments had completed a full inventory; as of October 2018, 916 datasets have been cataloged.⁸

DataSF's next step was to make these datasets more accessible across departmental boundaries. In many cases, they could be released publicly on San Francisco's open data portal—enabling any department (as well as any member of the public) to access them without needing to go through bureaucratic channels or arrange data-sharing agreements; more than half of the inventoried datasets have been published as open data. For more sensitive datasets that cannot be released to the public, data-sharing agreements have been developed as necessary. But that process can be arduous: one such agreement, which will enable all the local health and human service agencies to coordinate delivery of their services, took more than a year to establish.⁹

The next phase of DataSF's ongoing efforts was ensuring that the city's data is of a high enough quality to facilitate analytics. This means that datasets should be accurate, up-to-date, consistent, and complete—a dataset about, say, cooling towers that omits records or has not been updated for several years would be of little use. But because city staff rarely employ their administrative data for analysis, they are not trained to consider these attributes. Bonaguro is therefore educating them in the tenets of data quality and how to achieve it. In 2017, DataSF released a guide titled "How to Ensure Quality Data," with an accompanying worksheet that walks staff through the steps necessary to assess and improve data quality.¹⁰ The team

is also training departments on data profiling, a technique for evaluating the integrity and reliability of datasets. Showing departments the limitations of their data has quickly paid dividends, Bonaguro says. “We did a test where we profiled a department’s data and we brought it to a meeting. Their eyes bugged out. They’d never looked at their data that way.” Shocked by the poor quality of their data, staff in that department have been following DataSF’s new quality guide. DataSF has also created citywide standards for commonly occurring fields such as dates and locations to make it easier to match records and aggregate statistics across datasets.¹¹

Creating an ecosystem of well-curated, accessible, and high-quality data required years of heavy lifting throughout San Francisco City Hall. Yet this was just the foundation of a larger effort: the data adds value only when it actually helps departments provide better services and governance. With that in mind, Bonaguro has recently shifted her focus toward training departments to improve their operations using data.

It is in this effort that Bonaguro’s background in management and user design—together with her humble personality—has proved most essential. If DataSF were to approach departments claiming that data has all the answers, she says, “You wouldn’t even be laughed at. You’d just be ignored.” Departments would simply choose to avoid working with DataSF. “You really need to focus on developing relationships,” Bonaguro adds. “These people know their business and so we have a lot to learn from them.”

Bonaguro also knows not to immediately push departments toward the most sophisticated uses of data. After all, making decisions in a new way can require significant operational changes. And if staff are not familiar with data and algorithms, some will be threatened or insulted by suggestions that these technologies could improve their work, in part as a response to prior bad experiences with technologists who came in with little respect for the practices and expertise of existing staff. Recognizing these barriers, Bonaguro strives to learn about departmental needs and to “meet people where they’re at.” In partnership with the city’s Controller’s Office, she created a program called Data Academy to “lift data skills and capacity across the city,” with courses that teach skills such as using databases, visualizing data, and creating information dashboards. “We think about everything in terms of gateway drugs. Data Academy is a gateway drug,” Bonaguro says. “It’s this continuous story of moving people up the ‘data-use chain.’”

These tactics have helped Bonaguro develop partnerships with almost every department across the city. While some are eager to use data, others are resistant to disruption and outside influence. To showcase the value of data and demonstrate her good-faith intentions, Bonaguro starts with small projects that address each department's priorities and needs. She asks questions such as "What are the key questions you feel like you can't answer easily, or that you're answering over and over again?" and uses the responses to create dashboards that track and visualize metrics of interest—thus demonstrating how data could improve operations by breaking departments out of inefficient reporting practices that left data difficult to access and interpret.

Bonaguro recalls one department that spent most of their first meeting yelling at her. But she remained attentive to its needs and worked with staff to develop several dashboards that helped them monitor their performance. The department quickly warmed up to Bonaguro and has since made great strides in performance through better use of data. "That's how you move them to the next level," she explains. "You've solved something so you have a basis of trust on which to build the next step. How do you discover that? Through user research and design thinking, not through technology thinking."

Even once departments have recognized data as a valuable resource, much work remains to be done. We have seen several times that choosing what metrics to monitor and optimize for is a difficult and consequential task. Many well-intentioned efforts to use data in government go wrong because they fail to synthesize the glut of available data into metrics that appropriately capture their goals. "Metrics are tricky. Most metrics are bad," Bonaguro says, adding that when you choose the wrong metric, "then you're working toward the wrong thing."

All too frequently, Bonaguro notes, departments track metrics related to the quantities and processes behind their operations yet overlook the actual impacts of those operations and the desired outcomes. Instead of asking departments how many people they served last year, Bonaguro asks, "Did you serve them well? What happened as a result?" For as the next section of this chapter will describe, social service agencies that focus on how many people they serve rather than the impacts of those services will flounder. DataSF is therefore creating a Data Academy course to help departments

design metrics that are tailored for their specific operations and goals. The course divides metrics into three categories: how much was done (quantity), how well it was done (quality), and who is better off as a result (impact).

Bonaguro sees all of this work as creating “fertile ground” for the most sophisticated stage of using data in city government: applying machine learning to improve operations. DataSF launched a program in 2017 to help departments use data science, and several quickly got on board.¹² The Department of Public Health built a predictive model to identify mothers who will drop out of WIC (Women, Infants, and Children, a federal program that provides services to low-income pregnant women, recent mothers, and young children) so that it can identify program barriers and make changes to better aid the women and their children.¹³ In another project, the Mayor’s Office of Housing and Community Development created an algorithm to flag eviction notices that appear anomalous or unlawful so that the city’s eviction prevention services can intervene and keep residents in their homes.¹⁴

Bonaguro’s work in San Francisco closely mirrors Mashariki’s efforts in New York City. What makes them both exemplary Smart Enough City leaders are not their technical skills but their ability to pair technical acumen with a firm grasp of municipal needs and operations. For although enthusiasts of smart cities typically focus on the value that machine learning algorithms can unleash, such benefits cannot be realized without a long and arduous process of governance and institutional change: creating data inventories, bridging gaps between departments, and training staff to manage and use data.

Even then, insights from data cannot be translated into social impact without the traditional government operations that are so often maligned by technophiles. In New York City, for example, the Mayor’s Office of Data Analytics provided invaluable information and analysis that aided the response to the Legionnaires’ outbreak—but analytics did not, on its own, resolve the crisis: NYC Emergency Management coordinated the activities of several agencies, the fire department inspected buildings for cooling towers, and the Department of Health and Mental Hygiene tested and cleaned cooling towers. These activities were essential to curbing the spread of Legionnaires’ disease. MODA helped direct these efforts, but success in preventing further illness ultimately depended on the work and expertise of other agencies.

“I don’t want to make it come off as if MODA was this superstar,” says Mashariki after sharing the story about the Legionnaires’ outbreak.¹⁵ “Data and analytics don’t solve the problem: they support and add value to the people in your city that do solve the problem. Finding cooling towers in New York was like looking for needles in a haystack. MODA’s job wasn’t to find the needle—our job was to burn down the haystack to make it easier for the people who are actually doing the job to find the needle.”¹⁶

* * *

Teams such as MODA and DataSF are crucial because poor data management and a lack of coordination across agencies can doom even the most well-intentioned efforts. That is exactly what happened in Seattle, where the Human Services Department (HSD) identified major gaps between its attempts to curb homelessness and the outcomes it was generating.

In 2015, Seattle’s homeless population surpassed 10,000 people. Almost 4,000 of that group were living unsheltered on the street—a 38 percent increase from 2013, and the fourth straight year in which Seattle’s unsheltered population had grown.¹⁷ Dozens of homeless people died every year, and thousands of children were homeless.¹⁸ As the city concluded a ten-year initiative (commenced in 2005) to end homelessness, it was clear that the situation facing Seattle was “worse than ever.”¹⁹ Local leaders declared a state of emergency.

For homeless mothers in Seattle like Shakira Boldin, accessing services for herself and her son was a constant struggle. “I would call programs and they would either be full or wouldn’t have space or wouldn’t be able to take me and my young child,” she recalls. Local service providers lacked the resources and coordination to give Boldin’s family the support it needed to stay safe and escape homelessness. “I had to have my son in a really volatile environment,” she says. “We were sleeping on mats on the floor, and I didn’t have anywhere to go.”²⁰

The HSD, which plays a primary role in maintaining the local social safety net, knew that drastic changes were necessary. Although the City of Seattle does not directly run any homeless shelters or other programs, it provides funds to community-based organizations—known in this context as “service providers”—to operate services such as shelters, hygiene centers, and meal programs. HSD was spending \$55 million each year to fund homeless services, and yet families like Boldin’s were falling through

the cracks. Hoping to determine how it was performing and where it was falling short, HSD undertook a detailed analysis of its investments in homeless services.²¹

“We needed to do a thorough investigation to see how those investments were performing,” says HSD Deputy Director Jason Johnson. “What we found was we couldn’t always tell. We did not always have the level of information that told us whether a program was successfully moving people out of homelessness and into permanent housing. That’s where the ‘Aha!’ happened,” Johnson says. “We weren’t able to tell the full story of the impact that these investments and programs were having on individuals. That’s what we needed to do.”²²

In discovering that it could not even ascertain how its efforts actually affected people’s lives and which of its programs were providing effective services, the Human Services Department realized that data about its homelessness programs was woefully incomplete.²³ Information was split across three separate data systems, forcing redundant data entry and preventing a cohesive portrait of services and their impacts from emerging. And because HSD did not clearly articulate the need for or value of data, service providers reported incomplete or unreliable information. This further reduced HSD’s interest in the data, creating a downward spiral that left providers feeling justified in their poor reporting practices.

Such meager data management made it difficult for HSD to answer even simple questions about homelessness. Determining how many people had been served a meal required a manager to coordinate with ten program specialists and manually add numbers from separate spreadsheets. Answering more consequential questions, such as which families were in permanent housing after receiving services, was impossible. “On both the funder side and the provider side, there was way too much energy spent simply on collecting and reporting numbers and data,” Johnson recalls.

This is not to say that HSD entirely lacked data about homeless services. “We had a lot of data” from providers, Johnson explains, but it was mostly “a bunch of ‘widget counts’ like how many people they’ve served and the demographics of those people.”

“Seattle had wound up with this patchwork quilt of outcomes,” says Christina Grover-Roybal, a Fellow from Harvard University’s Government Performance Lab who helped HSD assess and overhaul its homeless services.²⁴ The city evaluated programs unsystematically, using a mishmash of

criteria: showers taken, number of people who received services, amount of food distributed, exits to permanent housing, and so on. Even services of the same type (say, emergency shelters throughout the city) were often assessed on distinct metrics. Thus, Grover-Roybal explains, HSD “really couldn’t compare performance across programs even if they were in the same service delivery model.”

Service providers also struggled under this system, Grover-Roybal adds. “Sometimes one provider would have the same programs but run two different shelters, and those two different shelters would have different outcomes they were being held accountable to by the city. So as a service provider, they don’t know what they’re trying to achieve. We needed to get Seattle to monitor consistent outcomes across every single homeless service delivery program.”

Because Seattle had not specified to service providers what it wanted to achieve, every service provider was working toward different goals. “We were not always clear upfront about the outcome that we were trying to achieve,” reflects Johnson. “We were operating on the assumption that everyone was trying to get individuals and families into permanent housing, but in practice that was not always true. Services were helping people manage and mitigate the survival risks, but not necessarily trying to actively end their homelessness by getting them into housing.”

Remedying this situation required Seattle to reform how it structures and manages contracts with local service providers. “Honestly, I feel like contracting was the only tool we had,” Johnson says. “There were no other changes that we could implement to change service delivery, how data was collected, or how we looked at performance. Our only tool to do that were the contracts.”

Government contracts are granted through a process called procurement: when Seattle decides that it wants to provide a social service, it requests proposals (or “bids”) from companies and nonprofits. The city reviews these proposals and chooses to work with the organization that submits the best one (typically defined, whether explicitly or not, as the lowest-price proposal). Seattle then signs a contract with the winning bidder, after which it provides funding to that organization in exchange for that group’s providing the desired program or service.

Seattle’s dependence on contracts is not unique: governments across the United States rely on contracts to complete many of their most

essential tasks, explains Laura Melle, the senior procurement lead for Boston's Department of Innovation and Technology. "Contracts are an input to every single output," she says. "A lot of people don't realize that government doesn't deliver our core services from scratch"; instead, whether paving streets or designing a website, "we're actually doing that in partnership with a private-sector company. Our role is often to select partners and manage contracts, with the goods and services provided by the private sector."²⁵ One estimate suggests that on average half of a city's budget goes to procured goods and services.²⁶

Contracts, in other words, are the tools that turn the government's policy vision into reality. "A lot of really smart people have a lot of great ideas, but how do we make it happen?" Melle asks. "Whatever the great idea is, contracts are how you translate that into something that actually works for people the way that it was intended to." Effective contracts can bolster valuable government programs, while poorly constructed or managed contracts can doom even the best-designed policies.

Unfortunately, because procurement and contracts are typically seen as administrative and boring, the latter outcome is far more common.²⁷ Procurement processes are highly regulated and rigid, unattractive features that sharply reduce the quality and quantity of proposals that governments receive. And instead of being structured to incentivize desired performance outcomes, government contracts are typically managed with an emphasis on affordability and basic compliance.

Contracts had long been poorly managed in Seattle. When the Human Services Department reviewed its homeless services contracts, it found that a vast tangle of poorly defined goals and disconnected programs had accumulated over the years—more than 200 contracts with sixty homeless service providers.²⁸ Every time in the past that HSD had wanted to expand or provide new services, it received a pot of money from City Council and signed a new contract with a local service provider for that specific purpose. These contracts were never thereafter reprocured or restructured; the resulting jumble of contracts made providing and evaluating services difficult. "Service providers that had been around long enough had lots of different contracts for lots of different pieces, based entirely on how City Council had been divvying up money for the last ten to fifteen years," explains Grover-Roybal. Some contracts were more than a decade old. Some providers had numerous contracts for closely related services.

This “administrative nightmare,” as Grover-Roybal calls it, made it difficult for service providers to effectively meet people’s needs. “Even though the service providers didn’t necessarily think of all of these programs as separate,” she explains, they had to rigidly allocate staff time and other resources according to their specific contracts, regardless of what services would actually have the most impact. Moreover, because contracts were never adjusted after their initial procurement, providers could not adapt their services to meet the community’s changing needs. This created a situation, Grover-Roybal says, in which “there are some shelters that are frequently underutilized and some shelters that are frequently overutilized. But the way that it’s set up right now, each shelter is limited to the size it was when HSD did the initial request for proposals, and that could have been five to ten years ago.”

Johnson points to the local Young Women’s Christian Association (YWCA) “as emblematic of this issue.” Over the years, the YWCA had collected nineteen separate contracts for addressing homelessness. Managing these contracts took three dedicated staff members at the YWCA and four more at the city. More importantly, the artificial barriers raised by these contracts prevented the YWCA from most effectively serving those in need. Shakira Boldin and her son might walk in and be the perfect fit for a program designed for families—but if the YWCA had already spent the money allocated to that program, it could not pull unspent funds from another program because the two were governed under separate contracts. Boldin and her son would be left without help.

To counteract this problem, the Human Services Department developed a novel approach: “portfolio contracts” that consolidate their previously separate contracts with service providers so that funding could be allocated more flexibly. Instead of having distinct contracts for each of their programs, providers now have a single contract with one pool of funding to cover a larger portfolio of services; the first pilot of this approach merged twenty-six contracts (totaling \$8.5 million per year) into just eight.²⁹ This is the “biggest win,” according to Johnson, because it “allows agencies the flexibility to move city money to where the individual that they’re trying to serve needs that money to be.”

The introduction of portfolio contracts solved the problem of service providers being burdened with and constrained by too many separate contracts. But the city still had to ensure that service providers were working

toward the common goal of moving homeless people and families into permanent housing. HSD planned to do so through incentives in the contracts that reward service providers for meeting performance benchmarks.

Despite its intentions, however, the city can achieve only so much acting independently: not only does Seattle not provide social services directly, but it is just one of several social service funders in the region. King County (which encompasses Seattle) and the local United Way are also major funders of the same service providers. Even if the city designed contracts with clear performance measures, almost half of the service providers' work would still be commissioned by those two other funders; if they continued to promote separate objectives, social services would remain disjointed and ineffective. Thus, to effectively create a coherent agenda for every local service provider, Seattle had to align its goals with those of the other key stakeholders.

Over the course of a year—"I don't even want to think about how many meetings," Johnson remarks—the city, county, United Way, and political leaders formulated a common set of goals for homeless services that emphasize long-term desired outcomes. Chief among these is getting the homeless into permanent housing and preventing them from returning to the streets. Another key goal, because the African American and LGBTQ populations have historically been underserved, is ensuring that every homeless demographic receives services proportional to its needs. Finally, these stakeholders wanted service providers to collect more accurate and comprehensive data about their operations. All three social services funders now include performance incentives tied to these outcomes in their contracts.

With the new contracts in place, HSD has instituted monthly meetings with service providers to ensure that they achieve adequate progress toward these goals. In the past, providers were rarely monitored beyond ensuring that they complied with local regulations. Whether performing well or not, HSD had little insight into or influence over their activities. Now, says Johnson, "If providers are underperforming, every month there's going to be a discussion about how to remedy that." These conversations (aided by the newly collected data) have already helped the city and service providers align their resources and priorities—for example, by identifying families who are falling through the cracks and devising plans to provide them with tailored aid.

By creating flexible portfolio contracts, setting clear goals, and gathering better data, Seattle has vastly expanded its ability to decrease the local

homeless population and mitigate the harms that homeless individuals face. “Now that they have the actual performance data, they can figure out what is working for people and what is not,” Grover-Roybal explains, adding that HSD has already learned a great deal about how people are actually moving through services and which providers are most effective.

Although much work remains—more resources and new policies are necessary to address the underlying issues—these gains are beginning to translate directly to improvements in the lives of homeless people and families in Seattle: in the first quarter of 2018, more than 3,000 households were moved into permanent housing or maintained their housing through city investments in homeless services—a 69 percent increase over the first quarter of 2017.³⁰ In fact, only six months after HSD began piloting portfolio contracts and performance-based pay, Shakira Boldin’s family was placed into permanent housing. “I can’t really explain the feeling,” she says. “Every day I wake up and I just feel blessed that me and my child have a roof over our heads. I feel like my future is bright.”³¹

* * *

One of the smart city’s greatest and most pernicious tricks is that it misappropriates the role and meaning of innovation. First, it puts innovation on a pedestal by devaluing traditional practices as emblematic of the undesirable dumb city. Second, it redefines innovation to simply mean making something more technological.

This chapter, which provides our most extensive look at what actually enables and sustains Smart Enough Cities, defies that logic: the most important innovations occur on the ground rather than in the cloud. Technological innovation in cities is primarily a matter not of adopting new technology but of deploying technology in conjunction with nontechnical change and expertise (of course, innovation need not involve technology at all). Cities must overcome numerous institutional barriers just to make data meaningful and actionable. MODA and DataSF did not need to find the optimal machine learning algorithm—instead, they had to painstakingly break down departmental silos, create new practices to manage data repositories, and train staff in new skills.

Seattle most clearly illustrates the benefits of recognizing that innovation means more than just “use new technology.” Municipal governments operate within a remarkably complex structure: their powers and capabilities are limited, and they must engage with numerous other institutions.

Yet no smart city technologies are designed with any such structure in mind; a focus solely on technology would have left HSD powerless to improve homeless services. When Jason Johnson identifies contract reform as the city's "only tool" to address endemic homelessness and highlights the year of meetings required to unite several social services funders behind shared goals, we see clearly that technology is impotent to address many of the pressing challenges that cities actually face. Data is helping Seattle to evaluate programs and identify where resources are needed, but it would have little impact without these more systemic reforms.

Technology also cannot provide answers—or even questions—on its own: cities must first determine what to prioritize (a clearly political task) and then deploy data and algorithms to assess and improve their performance. *Wired* famously promised that "the data deluge makes the scientific method obsolete" and represents "the end of theory,"³² but in today's age of seemingly endless data, theory matters more than ever. In the past, when they collected minimal data and had little capacity for analytics, cities had few choices about how to utilize data. Now, however, cities collect extensive data and can deploy cutting-edge analytics to make sense of it all. The magnitude of possible analyses and applications facing them is overwhelming. Without a thorough grounding in urban policy and program evaluation, cities will be bogged down by asking the wrong questions and chasing the wrong answers. Seattle had lots of data about homeless services, for example, but lacked a strategy to guide data collection and analysis toward its ultimate goals.

"The key barrier to data science is good questions," observes Joy Bonaguro. Using data effectively in city governments requires determining which issues, among the many that cities confront, can effectively be addressed with data. Furthermore, improving operations with data often hinges not on developing a fancy algorithm but on thoughtfully implementing an algorithm to serve the precise needs of municipal staff. Bonaguro therefore seeks far more than technical expertise when building her team. "When we hire data scientists," she explains, "I really want someone who does not want to just be a machine learning jockey. We need someone who is comfortable and happy to use a range of techniques. A lot of our problems aren't machine learning problems."

In Chicago, Chief Data Officer Tom Schenk has the same priorities when hiring for his team. "The challenge is trying to find the data scientist and

the researcher who can work well with departments, because that's the key aspect," he notes. "And there are a lot of researchers who are not great about that. We need to find one that can come in and not just do the statistics, but also sit in a room with a manager and find out all the information they need to know."³³

One of Chicago's data science projects relied on precisely this kind of on-the-ground research and relationship building. Several years ago, Schenk began working with the Chicago Department of Public Health (CDPH) to proactively identify the local food establishments that posed the greatest threats to public safety. If Schenk could predict which restaurants were most likely to violate public health regulations—precisely the type of task that machine learning excels at—he could direct food safety inspectors (known as “sanitarians”) to those locations and help CDPH make the best use of its limited resources.

On a technical level, the project sounded easy: develop a machine learning model that references historical food inspections to identify indicators of unsafe establishments. But Schenk knew the project would not be so simple: he needed to study a large city agency with complex operations about which he knew little, and then develop and deploy an algorithm that could be embedded in its daily operations. So instead of focusing solely on how to create the most sophisticated algorithm, Schenk prepared himself for intense research.

When Schenk approached CDPH's food inspections manager about the project, he emphasized that a successful collaboration required him to gain a deep understanding of CDPH's goals and operations. “We're going to ask you a lot of questions that are going to seem very rudimentary,” Schenk told her at the time. Such research is vital to successfully using data science in government, he explains. “It's super easy to miss what departments don't think is important because it's very banal to their process, but is key for our statistical modeling. The stats aren't hard for us. We spend most of our time talking to the client, trying to understand everything so we can apply statistics.”

Even once the machine learning model appeared to be operational, however, it still had to undergo another critical step: experimental evaluation. With a background in policy analysis and medical research, Schenk knew it was important to test every model before deploying it. “We know there can be a disconnect between the logic we think is correct, and what happens in

the real world,” he says. “We need to introduce experiments to make sure things actually work.”

Schenk designed a double-blind experiment to evaluate whether his algorithm could actually help sanitarians catch a greater number of “critical violations” such as a failure to heat or refrigerate food at the proper temperature. On the basis of internal tests and simulations, he expected that his machine learning approach would drastically increase the department’s efficiency at finding critical violations. But the experiment indicated that the algorithm produced only a negligible improvement. “It took us a long time for us to dig into it and find out what was going on,” Schenk recalls.

With such a large gap between expectations and reality, Schenk realized that despite his best efforts, he must have overlooked a key aspect of the food inspections process when designing the algorithm. He went back to CDPH’s food inspections manager to determine what he had missed. During their conversation, she mentioned in passing that all of the sanitarians had just been reassigned to new neighborhoods for the first time in several years. This was the clue Schenk was seeking: he realized that what had appeared to be an important factor in predicting violations—zip codes—was actually just a reflection of differences between sanitarians, as each assigned food violations according to slightly varying standards. Schenk had not been aware that sanitarians were assigned to specific zip codes, and so did not account for this in his modeling.

Although disappointed by the algorithm’s failure during this trial, Schenk saw the experiment as a success: it had uncovered a gap between the assumptions embedded in the model and CDPH’s actual practices. Schenk updated the algorithm and several months later ran another experiment. This time, the improvement was notable. A simulation found that using the predictive model, CDPH could have improved its early detection of critical food safety violations by 26 percent. If CDPH had been following the recommendations of the predictive model, it would have discovered each critical violation a week earlier on average.³⁴ With these encouraging results in hand, Schenk and CDPH were finally confident that the model was ready for deployment. It has been in action, guiding where sanitarians conduct inspections, since 2014.

The Mayor’s Office of New Urban Mechanics (MONUM) in Boston is developing an even sharper focus on science and research. “For many years now, we’ve been talking about the need to become data-driven, and that

is clearly one important direction that we need to explore further,” says MONUM co-founder and co-chair Nigel Jacob. “But there’s a step beyond that. We need to make the transition to being science-driven in how we think about the policies that we’re deploying and the way that we’re developing strategic visions. It’s not enough to be data mining to look for patterns—we need to understand root causes of issues and develop policies to address these issues.”³⁵

In April 2018, MONUM released a “Civic Research Agenda” comprising 254 questions, the answers to which will inform the city’s efforts to improve life for all Bostonians. These questions range from big (“How can we gain a holistic understanding of the kind of future people want for Boston?”) to small (“What can be done to lower the cost of construction?”), from technological (“How does technology play a role in perpetuating or addressing longstanding inequities across our city?”) to nontechnological (“What is at the root of community opposition to new housing?”).³⁶

All of this is necessary, says Kim Lucas, MONUM’s civic research director, to ensure that municipal projects are based on evidence and demonstrated civic needs. “You can’t solve a real-life issue if you don’t understand it, if you don’t ask the right questions, and if you don’t understand how to get the right information,” she explains. “That’s all research is: asking a question and then finding out the right information. And the next step when you have a finding is to do something with it.”³⁷

Staying grounded in research helps Boston avoid the perils of tech goggles. “Technology is a great tool, but it is not the answer,” says Lucas. “Technology is a tool toward getting the answer more efficiently.” Lucas relies on research, in other words, to “find the right tool to answer the right questions. If you’re not asking the right question in the first place, how do you know that technology is the right approach? It may or may not be.”

Which brings us back to the core message of this book: cities are not technology problems, and technology cannot solve many of today’s most pressing urban challenges. Cities don’t need fancy new technology—they need to ask the right questions, understand the issues that residents face, and think creatively about how to address those problems. Sometimes technology can aid these efforts, but technology cannot provide solutions on its own.

While that observation may appear obvious by now, confidence in the possibility of understanding and optimizing society through technical measures has been remarkably persistent not just over the past several

years but over the past several centuries. The next chapter—the book’s conclusion—will discuss the evolution of those beliefs. In exploring the similarities between past and present, along with how historical attempts to rationalize society have gone awry, it will demonstrate why smart cities are bound to fail. The book will conclude by highlighting how we can avoid such misguided and narrow thinking, synthesizing the lessons we have learned to provide a clear framework that can guide the development of Smart Enough Cities.