

Conclusion

Brett M. Frischmann, Michael J. Madison, and Madelyn Rose Sanfilippo

Smart cities require much more than smart tech. Cities need trusted governance and engaged citizens. Integrating surveillance, AI, automation, and smart tech within basic infrastructure, as well as public and private services and spaces, raises a complex set of ethical, economic, political, social, and technological questions that requires systematic study and careful deliberation. Using the GKC framework to structure case studies that examine smart tech deployment and commons governance in different cities has served two fundamental purposes:

First, it provides important, even if incomplete, guidance for communities deploying smart tech. The book has deepened our understanding of community governance institutions, the social dilemmas communities face, and the dynamic relationships between data, technology, and human lives. We have sharpened attention on key areas that practitioners and researchers need to focus on. Much more work is needed, however, to develop and improve guidance in this politically and culturally contentious space. At this stage, we emphasize that design principles for knowledge and data governance institutions (Ostrom 1990) are not available, at least not based on the social science. There are no universal answers, just as there are no panaceas, technological or otherwise, to the many social dilemmas communities face.

Context matters in more ways than one. Cities are incredibly varied and complex. Within every city, there are many unique communities. Nested action arenas at macro, meso, and micro levels involve mixed sets of actors, pursuing various goals and objectives, while grappling with different obstacles. Smart tech may serve as useful tools in these arenas while also generating challenges and even additional social dilemmas. Interdependencies complicate matters dramatically. We could go on painting an incredibly complex picture. But that is neither necessary nor helpful.

We remain confident that a principled approach to smart city governance is possible. It begins, as Chapter 10 suggested, with asking questions and asking them in a structured way. The point of acknowledging complexity and embracing the

contextual nature of social reality in urban environments is to encourage researchers and policymakers looking for a path toward trusted governance and meaningful citizen engagement. They should be empowered to ask relevant questions and explore possible solutions. They should be looking for structured ways to deconstruct and decompose complexity and context.

Second, it is necessary to support interdisciplinary social science. We hope that what we have collected in this volume is the beginning of a sustained, systematic, comparative, and longitudinal research effort focused on smart technologies. More work on more cities and, frankly, more work on the cities studied in this volume, is the only way to develop empirically grounded answers to the many questions raised.

Throughout this book, authors have asked contextual research questions and explored compelling but often distinct answers guided by the shared structure of the GKC framework. In this Conclusion, we discuss some of the key themes across chapters in this volume, considering lessons learned and implications for future research.

RECURRING THEMES

Transparency, Superficial and Deep, Necessary but Insufficient Governance

Across the case studies, cities and communities using smart tech face *transparency* issues. In this context, transparency specifically refers to the *availability* of specific knowledge resources to community members. What is the relevant knowledge and to whom and how it is available varies considerably across and even within cases. The GKC framework provides a useful lens for examining these details. Transparency generally describes features of the governance structure, namely relationships between actors and resources; issues of transparency around who makes data availability decisions, and how, also arise.

In some cases, transparency is an explicit “open government” objective for which smart technologies can be useful tools. Notably, in such cases transparency is not really a primary end. Rather, it is typically part of a broader effort to pursue democratic values, such as government accountability and citizen engagement, and economic values, such as economic growth and entrepreneurship, and to identify and address corresponding obstacles to achieving those ends. Raymond and Kouper (Chapter 3) explain the origins and evolution of open data/government initiatives at the federal government level and how that thinking is reflected in smart city initiatives, such as the Bloomington Open Data Portal. Across the case studies in this volume, there was little direct evidence or detailed discussion of the relationship between transparency and economic values. At best, some lip service was paid to the idea that open data would support entrepreneurship or economic development. For the most part, the focus was on democratic values.

Local governments often aim to make information about their practices, decision-making, and governance processes available to citizens. The idea is to use data and smart tech to educate and enable community members to engage more effectively with government, whether to obtain public services, cast a vote, lodge a complaint, reform or remedy harmful government practices, or pursue other opportunities. Such open government initiatives often take the form of making data publicly available online, through websites and portals, sometimes accompanied by software tools designed to enable citizens to access, interpret, and use data. The results described in the case studies were mixed.

Some cities, such as Bloomington and Philadelphia (Chapter 5), demonstrated moderate success in making datasets and tools available to the public through online portals. These cases suggest two potential lessons (worth bearing in mind as policy guidance and for future research).

First, building transparency and making it useful are not easy. Effective transparency takes resources, planning, coordination among different sets of actors, commitment to shared values, and maintenance. The Bloomington case study provides an especially useful illustration of how using the GKC framework to examine different action arenas can isolate different obstacles to overcome in pursuing an open data/government initiative. Collecting, publishing, and accessing/using the data present different governance challenges for different actors, who need to coordinate with each other over time. As we saw in this case study and others, transparency may present new risks to consider, such as the inadvertent exposure of sensitive personally identifiable information (PII) through an online portal.

Second, transparency may be achieved superficially or deeply. Transparency exists on a continuum, and it varies, based on what resources are made available, to whom, and for what purposes. The GKC framework proves useful for investigating different types and degrees of transparency.

Superficial transparency focuses mostly on making data publicly accessible without consideration of whether and how it is used. At the extreme, nominal transparency would be nothing more than window dressing to appease the public, put off critics, comply minimally with procurement or other rules, or generate the appearance of propriety. For example, smart city critics and open data advocates alike have decried efforts, such as Chicago's open data portal, which does not meet minimum standards of accessibility via the use of proprietary and unstructured datasets that are difficult to use and interpret.

None of the case studies were fully at this extreme, although there were instances suggestive of nominal transparency – for example, in Anna Artyushina's observation that despite hundreds of lengthy public documents describing the plans for Sidewalk Toronto (Chapter 8), details about specific smart technologies and financial aspects of the deal remained hidden.

In the middle of the continuum, we find cases where transparency results from a sincere commitment to support open data/government initiatives as means to

engage and empower citizens. For example, in Bloomington, Philadelphia, and Pittsburgh (Chapter 6), the cities invested substantially in making government data publicly accessible with these ends in mind. Yet in each of these cases, transparency may not deliver fully the intended results if the quality of the data or data portal is poor, or if the public does not know about, trust, and have the capacities needed to effectively use the resources made available. These “transparency dilemmas” demand attention in the design of governance institutions. Just making the data accessible is not enough, which raises a more general point: In these intermediate cases, transparency typically falls short of community needs because of the focus on government data as the only (or the primary) relevant resource to be shared in the smart city knowledge commons.

Deep transparency requires more. Many local governments express intentions (commitments) to use smart tech to become more transparent not only by making government data publicly available but further in sharing information regarding their decision-making processes, the reasoning behind various initiatives, how data collection will occur, and plans for what to do with data and smart technologies. However, we do not see such comprehensive transparency in many of the case studies. As noted in the Bloomington case study, “[b]oth open data and open government are key for advancing government transparency and entrepreneurship, but the transparency does not seem to extend to decision-making in the creation and sharing of the datasets.” Similarly, in the Toronto case study, public documents had the veneer of a strong commitment to transparency, but openness only extended to certain aspects of the project while many important details, such as the specific technologies the partnership planned to implement and financial aspects of the deal, remained secret. With its formal, detailed privacy principles and governance structures for urban data and surveillance technologies, Seattle seems to be an exception, although there are indications that other cities may be following Seattle’s lead (e.g., AlAwadhi and Scholl 2013; Stübinger and Schneider 2020). Notably, Seattle’s commitment to transparency evolved through a series of initiatives and has deepened over the past decade.

It is worth noting that the polycentric and sometimes decentralized character of urban governance generally makes it difficult to practice transparent governance consistently at the level of the city as a whole. Some agencies and city authorities may be better on this score, some may be worse, and some may have more legitimate reasons for blending transparent and nontransparent governance in different ways. The GKC framework usefully enables analysis of these variations across action arenas at different scales.

Related to transparency, clarity, specificity, and other such considerations arose repeatedly in our case studies as challenges to the quality of communications about smart tech decision-making, planning, and governance-related issues. Hype and tech boosterism distort public perceptions, beliefs, and expectations. Ambiguity clouds judgment and disrupts reasoned public debate and deliberation about

supposedly smart technologies, which can lead to failure, as discussed in Chapters 8 and 9 relative to Toronto. These chapters not only identify the inherent limitations in nontransparency with respect to the public, but also in how it shapes interactions and power asymmetries among decision-makers. We can contrast this with efforts such as Seattle, where the focus is on participatory and responsive governance models, coupled with transparency, to cut through the hype and require specificity of features and functions to meet citizens' needs and improve services. We note that this approach to public documentation and responsive governance is not necessarily always coupled with transparency about datasets, as Chapter 7 shows us that even private sector firms and quasi-public spaces, such as Disney World, benefit with respect to public opinion from transparent and participatory mechanisms around data collection and use.

The bottom line – or at least the thematic lesson learned from these case studies – is that transparency may be necessary but insufficient for the achievement of community goals and objectives. Coordination, trade-offs, and participation also play key roles, as explored in the next section.

Interest Alignment, Potential Conflicts of Interest, and Community Participation

Conflicting values, interest (mis)alignment, community participation, and political and economic power cohere as another key theme. It is often difficult to identify and evaluate the relevant set of interests and values at stake in social contexts. It is hard to know what matters, how much, to whom, and how conflicts among different values can or should be assessed. The GKC framework is descriptive and deliberately pluralistic, which means researchers primarily aim to identify different community goals and objectives and to evaluate patterns and outcomes in terms of benefits, costs, risks, and so on based on the community's stated goals and objectives. We are mostly interested in understanding whether and how different governance institutions enable communities who share various knowledge resources to achieve their own goals and objectives.

In this volume, this inquiry is complicated for a few interrelated reasons. First, cities involve many different communities in background contexts with rich political, economic, and cultural histories. Power dynamics are complicated. Communities often have different constituencies and priorities. Values can be highly contested. Second, knowledge commons and corresponding action arenas focused on supposedly smart technologies (including surveillance technologies) often are a layer added onto already existing action arenas – as in the Philadelphia case study's analysis of the meso-level action arena of vacant land management. Smart tech and data may serve as new means to address already existing problems, but they also bring their own set of considerations (resources, community members, goals/objectives, power dynamics, social dilemmas, governance institutions, etc.).

This can mean that new actors become participants in long-standing action arenas. These new actors – whether IT staff working within a government office or a private vendor offering smart tech solutions – bring their own interests to the community. New participants may challenge or reinforce existing governance structures, for better or worse. Third, it can be difficult to figure out how to account for different community members, their roles, and their interests. It is not always clear who to include as a community member. For example, are technology vendors members of smart city communities? Do their interests in making a profit count as community interests – as one of the relevant community goals and objectives? Are all citizens members of smart city communities? Would that mean that the public at large is effectively the community, in which case are we really talking about a commons or commons governance? These and other related questions push toward concepts of nested commons and polycentricity.

These complications surfaced quite starkly in case studies where concerns about potential conflicts of interest arose. A question that emerges in many chapters is whether data collection and usage is more aligned with the interests of the smart tech provider or the interests of the public. The provider could be a government actor (department) or a private company vendor, supplier, or operator. For example, in the Sidewalk Toronto study, Sidewalk Labs proposed new governance entities that superficially seemed inclusive but substantively seemed to “advance the [private] company’s economic and political interests. The heated controversy over the project highlights the underlying tension between the company’s vision of municipal governance as ineffective and the public call for stronger government oversight over the private sector in Canada.” Similarly, from the Toronto case study (Chapter 8):

- “Subsidized by the government, the trusts would seek to maximize profits from the city spaces and infrastructure, often at the expense of the comfort and health of its residents.”
- “My analysis . . . shows that the trusts benefitted some members of the community more than others. Specifically, Sidewalk Labs openly declared its goal to support developers in the project, as well as the businesses coming to operate in the smart city.”
- “Before any technology is implemented, the citizens may want to decide whether they benefit from having an algorithm decide when and where they relocate, or if eliminating municipal jobs harms the community. When designing a trust for the smart city, it is important to know who gets to decide on behalf of the community, and what the community needs are.”

In the Philadelphia case study, “there was a disconnect between the *smart tech planners and users* (mostly government actors but also vendors and consultants, like IBM . . .) and *smart tech beneficiaries*, including residents, businesses, and visitors.”

Recognizing this disconnect seemed to be an important part of the shift in macro-level strategy that led to more deliberate focus on community inclusion in the SmartCityPHL roadmap. Yet, as the authors noted, whether that commitment is realized or constitutes lip service remains to be seen. At the meso-level action arena involving vacant property management, each of the three complications noted above surfaced. While there has been some success in making smart technology tools and datasets publicly available, “the promise of these technologies is not yet fully realized in Philadelphia.” The authors concluded:

“No smart technology can independently overcome the political and organizational issues and the complex economic trends outlined earlier. . . . the implementation of smart technologies must be founded upon effective local efforts to break down city service barriers caused by entrenched political and administrative structures. In addition, equitable, comprehensive, and successful vacant property disposition requires enhanced engagement with residents and community groups. Smart technology can support Philadelphia’s ongoing efforts to address these constraints and put vacant properties back into productive use for its citizens.”

The Bloomington case study similarly observed that

“in many instances of creating an OGD portal, . . . the design process seems to omit citizen cooperation and participation. Instead, legislators and government officials constitute the prime co-creators of such portals through incentivizing or enforcing portal creation efforts. Third parties, whether nonprofit or commercial, also become participants in OGD portal creation, as these organizations provide governments with the technological infrastructure to support these portals.”

In his analysis of modern digital communities, Richard Whitt draws lessons from the Toronto example and suggests: “Our digital communities should embrace the active participation of citizens and visitors alike in the increasingly blended spaces that constitute the self and world, the private and public, and the physical and virtual.”

These patterns echo findings from previous knowledge commons case studies, such as the issues associated with imposed commons governance that are apparent relative to privacy commons arrangements (Sanfilippo, Frischmann, and Strandburg 2021). It is key to understand that those with decision-making power over rules and governance are not always or necessarily the information subjects, thereby limiting fidelity in responsive governance and testing legitimacy in instances where trust is not developed. In the context of this book, the public are the information subjects but have little decision-making power in how the data being collected on them is used and what is being done with it.

Polycentric Coordination Dilemmas

Another prominent theme in the case studies concerns coordination dilemmas. As various chapters explored, cities are incredibly complex, with many centers of decision-making, sometimes competing and sometimes compatible, often intersecting and overlapping. Scholarship of metropolitan governance over the past seventy years has explored the nature of polycentric institutions, building on Polanyi (1951) and Ostrom, Tiebout, and Warren (1961) to identify conceptual and structural distinctions between functional and dysfunctional arrangements, issues of efficiency, dynamics, and complexity as associated with this feature of local government. Beyond an explanation for why garbage collection services are replicated or the nested coordination arrangements to solve public safety and school choice dilemmas, polycentricity describes decision-making and control around modern public interest technology, as data is collected and shared across agencies and levels of government. Cities rely on polycentric governance even before smart tech enters the picture. In terms of public administration, cities are comprised of functional departments with incompletely autonomous decision-making and areas of responsibility. Many cities are parts of counties and regions, which may have their own governance, and cities have to coordinate with other government forms (such as utility districts and schools, and other, adjacent cities). When smart technologies are embedded within an already polycentric system, it becomes more complex and exacerbates existing dynamics. The case studies in this book note that technology adoption does not solve issues of coordination or competition among polycentric public utilities or agencies, but rather perpetuates and may exacerbate these tensions.

The theme is captured in a line from the Pittsburgh case study: “Polycentricity may be a problem to be solved rather than a solution to questions of appropriate, effective, and accountable governance.” Across the studies, we saw many different coordination dilemmas. Lack of coordination and consistency led to inefficiencies, redundancies, confusion, and conflict. The following observation made in the Philadelphia case study in the context of the vacant property action arena captures the issue seen in many other cases: “The fact that several different city agencies are involved in vacant property management and disposition in Philadelphia makes it difficult to organize an effective, coherent approach. In the past, many of these agencies operated as relatively self-contained silos, with little cross-agency data sharing or collaboration.” In accordance with its more pragmatic and collaborative (cross-departmental) approach in the past five years, Philadelphia has tried to reform its organization approach to overcome this problem.

Additionally, lack of coordination and consistency can connect back to the transparency problems because if different agencies are compiling the data differently or have different standards for data, it can be difficult to make the data accessible.

Consolidating efforts of city agencies, or assigning responsibility for city-wide smart tech development and deployment to a single city department, may be one solution to this coordination and consistency problem. As described in Chapter 8, relative to polycentric structure and oversight in Toronto:

“In the Master Innovation Development Plan (MIDP), the company suggested establishing new governance entities that would mediate between the technology vendor and Canadian authorities and help members of the community collectively govern and manage smart infrastructure. These five new governance entities were the Urban Data Trust, the Waterfront Housing Trust, the Open Space Alliance, the Waterfront Transportation Management Association, and the Public Administrator.”

Further, the Open Space Alliance (OSA) would be aimed at coordination and support for functional polycentricity:

“Sidewalk Labs argued that the OSA would fix the problem of intersecting responsibilities, which results in public spaces not being properly cared for. Some municipal services, such as the parks and recreation departments, could be eliminated altogether. Data modeling and residents reporting problems through the app would help Sidewalk Labs plan for when additional help is needed and hire temporary workers.”

The creation of intersecting bureaucratic offices to smooth coordination efforts is not unique to Toronto, as similar efforts were documented in Pittsburgh and Philadelphia. For example, Pittsburgh created a new Department of Innovation and Performance in 2014, under a newly inaugurated mayor, specifically to advance and coordinate the city's uses of up-to-date computer networks and data-focused governance.

The prominence of polycentricity in these cases reflects the pattern observed around privacy-focused GKC cases (Sanfilippo, Frischmann, and Strandburg, 2021), yet offers new insight. Beyond the inherently polycentric nature of knowledge resource governance, reflecting communities that emerge in existing contexts with exogenous or imposed institutions, we recognize that the extent to which polycentricity is functional is critical to smart systems success. As with other contexts, top-down, exogenous rules-in-use are not as often responsive to community needs and norms, as are bottom-up rules-in-use stemming from commons arrangements. Functional coordination between approaches appears to be critical to patterns and outcomes, as evidenced by cases such as Toronto, where coordination efforts were stymied and the project ultimately failed at Quayside. This suggests that a more intentional focus on polycentricity within the framework may be beneficial in future case studies. Inquiries and practical applications should directly assess where decision-making is nested, competing, and/or coordinated among multiple units.

OBSERVATIONS

Smart cities are complex. It is important to put aside buzzwords and marketing hype and consider the ethical, economic, political, social, and technological implications of deploying and integrating supposedly smart technologies throughout urban environments. Smart cities present a host of governance challenges that are too easily obscured by slick marketing and grand promises of technological solutionism. In this volume, we encountered plenty of examples, but we also observed cities edging toward more pragmatic approaches to smart tech, data, and community governance. The penultimate chapter presented a proposal for principled decision-making that reflected this type of pragmatism.

While much of this book has been critical and focused on governance challenges faced by cities, we feel it is important to emphasize this rather mundane observation: *smart technology can be incredibly useful*. The case studies showed many examples, ranging from improving the quality and efficiency of government services to managing public transit, streetlights, and other infrastructure to making government data sources freely available to citizens. In nearly every action arena imaginable in the context of cities, there is a potential case to be made in favor of deploying some form of supposedly smart technology, bearing in mind that these are just tools that leverage more and potentially better data to develop actionable intelligence. The refrain repeated throughout the book about ignoring the hype is an important reminder, given the flood of tech boosterism in the smart city context and its potential to distort public perceptions and conversations. But the reminder is only to clear the deck and push towards more pragmatic public conversations and principled decision-making. There is often a potential case to be made, but it must actually be made taking into account countervailing considerations, including evaluation of alternatives and governance dilemmas.

What are those considerations? Initially, we described them in terms of patterns observed and lessons learned. As we reviewed what we had written, we realized that was too strong, given the small number of case studies in this volume and the variety of different subjects and action arenas among the small group.

First, intelligent governance of smart cities requires comprehensive public knowledge rather than superficial transparency. This means that community members must be informed and capable of action, whether in using data and tools or in voicing concerns about projects. Of course, as the GKC approach makes clear, the relevant community members and their roles and capabilities vary by action arena and context. Comprehensive public knowledge is not easily achieved, however. Some cities have pursued an open government agenda that entails making government data publicly accessible online along with software tools to analyze, map, or otherwise use the data. This is an important first step, yet it presents its own challenges, as seen in various chapters.

Second, comprehensiveness also relates to a theme from prior GKC research concerning the recursive, dynamic nature of knowledge commons governance and the need for a holistic and longitudinal approach to this type of research. Governance of urban data and smart technologies is multifaceted, dynamic, and continuous. It requires governance structures that interoperate as part of a governance system. At a minimum, smart city governance must encompass city planning, procurement, implementation, and management processes. Yet, as seen in various cities, each of these may operate as independent stages or action arenas with independent sets of actors interacting in particular ways, guided by separate rules-in-use. So, for example, procurement policies and practices may govern specific actors within a city department, such as the IT staff, and their interactions with tech vendors, while data governance and privacy policies and practices may govern actions of a different set of departmental actors, such as city officials, and their interactions with citizens. The case studies – the Philadelphia case perhaps most vividly – suggest that such independence may be an unproductive artifact of politics, public administration, and long-standing departmental siloes. From a more holistic perspective, procurement policies, impact assessments, terms of use, privacy policies, and other governance institutions should be seen as interdependent components of a *governance system*, rather than as isolated and independent institutions. This perspective is essential when each individual governance institution is necessary but insufficient for overcoming the social dilemmas or obstacles to achieving community objectives. The Seattle case study illustrates the emergence of a governance system that encompasses and integrates governance institutions across action situations.

Third, smart city conversations point simultaneously in lots of different directions. Current smart city research and advocacy, on all sides, seems to be speaking multiple languages at once. There's the good governance language. There's the privatization language. There's the surveillance and power language. There's the language of play and scripted behaviors. There's techno-solutionism language (on both sides!). There's the language of obscurantism and black boxing. The GKC framework is like a smart city Esperanto, except that it's not an insane language that no one wants to learn (we hope!). It's an accessible language that builds on common-sense intuitions. What are the resources? What are the problems? Who are the communities? How do we draw directional arrows among those things?

This Conclusion has emphasized themes about complexity and asking appropriate questions. We acknowledge that asking questions about complexity is not always a useful strategy for researchers or practitioners, including people in public administration, private sector partners and funders, and community organizers. We believe the GKC framework can help. The framework shows how to break down complexity into manageable chunks. Again, from all sides of the problem, there's a syntax for question-asking and decision-making that speaks in essentially the same terms to

everyone. There's the resource attribute chunk. There's the social dilemmas chunk. There's the community/collective identity and interest chunk. There's the historical contingency chunk. And so on.

LOOKING AHEAD

We have closed previous *Governing Knowledge Commons* books by paying tribute to Elinor Ostrom and other pioneers, thanking our contributors and new community members, and inviting readers to join us in this collaborative, interdisciplinary research and practice community. We do so once more. This volume is just the beginning of a sustained, systematic, comparative, and longitudinal research effort focused on knowledge commons governance in smart cities and other communities.

In a departure from past practice, in this Conclusion, we do not offer amendments to the GKC framework or the list of representative research questions. The framework remains a dynamic work-in-progress, open to future amendment, but we do not have any to propose at this stage.

REFERENCES

- AlAwadhi, S., and H. J. Scholl. 2013. "Aspirations and Realizations: The Smart City of Seattle." In *2013 46th Hawaii International Conference on System Sciences*, 1695–703. New York: IEEE. <https://doi.org/10.1109/HICSS.2013.102>.
- Ostrom, Elinor. 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. New York: Cambridge University Press.
- Ostrom, Vincent, Charles M. Tiebout, and Robert Warren. 1961. The "Organization of Government in Metropolitan Areas: A Theoretical Inquiry." *American Political Science Review* 55 (4): 831–42. <https://doi.org/10.2307/1952530>.
- Polanyi, Michael. 1951. *The Logic of Liberty*. Chicago, IL: University of Chicago Press.
- Sanfilippo, Madelyn R., Brett M. Frischmann, and Kaatherine J. Strandburg, eds. 2021. *Governing Privacy in Knowledge Commons*. Cambridge: Cambridge University Press.
- Stübinger, J., and L. Schneider. 2020. "Understanding Smart City – A Data-Driven Literature Review." *Sustainability* 12 (20): 8460. <https://doi.org/10.3390/su12208460>.