

Article

Smart Cities in the Agentic AI Era: Three Vectors of Urban Transformation

Esteve Almirall 

Center for Innovation in Cities, Esade Law & Business School, Ramon Llull University, 08034 Barcelona, Spain; esteve.almirall@esade.edu

Abstract

Agentic artificial intelligence—systems that reason, plan, and act autonomously within governed workflows—is converging with autonomous electric mobility and urban robotics to reshape how cities govern, move, and manage physical space. We argue that the simultaneous arrival of these three vectors is triggering a transformation comparable in scope to the Industrial Revolution. Cities that deploy across all three domains are becoming the new hubs of innovation: they concentrate talent, accelerate knowledge circulation, enable cross-fertilisation, and generate hybrid proposals that no single vector could produce alone. Just as Manchester, Birmingham, and the Ruhr became the defining centres of industrialisation because steam, textiles, iron, and coal recombined through the proximity of the engineers and entrepreneurs who moved between them, a small number of cities today are pulling ahead because they host the shared talent pool around which agentic governance, autonomous mobility, and urban robotics co-evolve. Conceptually, we extend the mirroring hypothesis in two directions: dynamically, arguing that organisations and urban ecosystems converge toward the configurations new technologies make possible; and ontologically, arguing that agentic AI introduces non-human agents into organisational architectures, requiring hybrid human–AI coordination. We formalise this dynamic as five propositions (P1–P5) of cumulative recursive hybridisation (CRH), operating through four reinforcing feedback loops—data, regulation, infrastructure, and talent. Together, these loops explain why the emerging urban order is path-dependent: early movers accumulate compounding advantages, while latecomers face exponentially rising costs of entry. We demarcate CRH from adjacent frameworks—general-purpose technologies, organisational complementarities, and complex adaptive systems—and test it against counterfactual evidence from failed, stalled, and Global South trajectories (Sidewalk Toronto, the Cruise rollback, Songdo, Bengaluru). We also examine its political-economy, equity, and surveillance limits. Drawing on comparative evidence from public-sector chatbot deployments, autonomous mobility ecosystems in the United States and China, and emerging urban robotics cases, we conclude that what is at stake is not incremental modernisation but the construction of a new urban order. The cities that act as innovation hubs for the agentic AI era will shape global standards, attract global talent, and define the institutional templates that others eventually adopt—much as the industrial cities of the eighteenth and nineteenth centuries did.



Academic Editor: Grigorios Beligiannis

Received: 30 March 2026

Revised: 10 April 2026

Accepted: 13 April 2026

Published: 15 April 2026

Copyright: © 2026 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\) license](https://creativecommons.org/licenses/by/4.0/).

Keywords: smart cities; agentic AI; autonomous mobility; urban robotics; cognitive government; mirroring hypothesis; cumulative recursive hybridisation; urban innovation ecosystems; chatbots; public sector

1. Introduction

Generative artificial intelligence (AI) is precipitating a step change in how societies organise work, coordinate decisions, and provision services. Since its public breakout in late 2022, adoption has progressed from experimentation to widespread use at a pace comparable to—yet faster than—previous general-purpose technologies such as electricity or the internet [1,2]. Crucially, these systems no longer merely classify or predict; they reason over open-ended tasks, draft and translate, synthesise evidence, and increasingly act within workflows, making them candidates for a new operating layer of governance, mobility, and urban infrastructure [3,4].

This inflexion point is especially consequential for cities. Throughout history, cities have served as the primary arenas for social, economic, and technological experimentation—from the Greek polis to the industrial metropolis to the contemporary smart city [5]. Each wave of general-purpose technology has reshaped the urban form: the railway redefined the Victorian city, the automobile created suburbia, and the internet enabled the networked metropolis. Agentic AI—systems capable not only of generating content but of learning, anticipating, and acting on behalf of citizens and institutions—promises a transformation of comparable magnitude, touching simultaneously how cities govern, how people and goods move, and how physical urban space is maintained and managed [6,7].

Yet the diffusion of AI across urban domains is neither automatic nor uniform. While many private organisations have already reorganised around AI-augmented workflows, public administrations face structural constraints—rigid hierarchies, procurement and hiring frictions, and incentive misalignments—that slow endogenous change [8]. At the same time, citizen expectations have been irreversibly reset by private digital services that deliver immediacy, personalisation, and reliability [2,9]. The result is a widening gap between what is technologically possible and what public institutions are organisationally prepared to deliver. Closing this gap requires not simply “adding AI” to legacy systems, but re-architecting governance toward agentic models: interoperable, data-secure, auditable systems of human–AI collaboration that can learn, anticipate, and act on citizens’ behalf within democratically bounded mandates [10,11].

This paper argues that three interconnected vectors of AI-driven transformation are converging to reshape the smart city: (1) the evolution of public-sector chatbots and agentic systems that redefine the citizen–administration interface, moving from informational tools toward cognitive government; (2) the emergence of autonomous electric mobility—robotaxis, on-demand transit, and autonomous logistics—that is fundamentally altering urban spatial structure, connectivity, and cost; and (3) the deployment of intelligent robotics and urban infrastructure systems—from city brain platforms to maintenance robots and drones—that automate the physical management of the urban environment.

Crucially, we contend that these three vectors should not be understood in isolation, and that what binds them together is not the technologies themselves but the talent pool they imply and the interconnections that pool enables. Agentic governance, autonomous electric mobility and urban robotics all draw on the same scarce community of foundation-model, perception, control, data, safety, regulatory and operations specialists. When the three vectors are deployed in the same place, that community circulates between them: the engineers who built an autonomous vehicle perception stack advise the city brain team, the regulatory designers who licenced the robotaxi fleet shape the permitting interface used by the robotics start-up, and the operations staff move between projects within the same labour market. It is this circulation—of people, methods, working practices and tacit knowledge—that turns the co-presence of the three vectors into a generative dynamic rather than a coincidence, and that enables

the cross-fertilisation and hybridisation through which each vector makes the next one cheaper, faster and more reliable, and through which entirely new proposals emerge that no single vector could have produced in isolation. Because such circulation is bounded by daily contact, shared institutions and a common labour market, it is inherently local: it happens at the scale of the city, and it is at that scale that the city itself is reconstituted as an innovation hub—concentrating talent, compressing iteration cycles, multiplying the chance encounters that seed novel combinations, and accelerating the speed at which ideas move from conception to deployment.

The analogy with the Industrial Revolution is not merely illustrative; it is structurally precise. In England, steam power, mechanised textiles, iron production and coal extraction cross-fertilised through the geographic proximity and iterative recombination of the engineers, mechanics and entrepreneurs who moved between them [12,13]. The result was not a set of improved technologies but a new economic and social order: the industrial city became the organising unit of capitalism, and the cities that led—Manchester, Birmingham, Glasgow—shaped national trajectories, attracted global talent, and defined institutional templates for over a century. We argue that an analogous process is under way. The convergence of agentic governance, autonomous mobility and urban robotics generates cumulative recursive hybridisation—a dynamic that requires the co-location of talent, regulatory authority, experimentation capacity and institutional willingness to iterate, all of which are embedded in specific urban ecosystems rather than distributed evenly across national territories [14,15]. The cities that master this process will not merely be more efficient versions of today’s cities; they will be the new hubs of a new order, and the gap between them and those that hesitate will widen in the same path-dependent, self-reinforcing way that separated industrial from pre-industrial cities two centuries ago.

The paper thus advances five formal propositions, developed in Section 2.4 and summarised in Figure 2. Building on the mirroring hypothesis [10]—that organisations realise the potential of new technologies only when they reconfigure themselves to mirror the possibilities those technologies open—we argue that the city is the natural unit of analysis for AI-driven transformation, because it is at the municipal level that regulatory frameworks, service delivery, infrastructure management and innovation ecosystems intersect.

Propositions P1–P4 specify four reinforcing feedback loops—data (P1), regulation (P2), infrastructure (P3) and talent (P4)—generated by the co-deployment of agentic governance, autonomous mobility and urban robotics within a city; Figure 3 represents the same dynamic as a system-dynamics model in standard Forrester notation. Proposition P5 states the resulting path-dependency claim: because P1–P4 are recursive and bidirectional, the inter-city distribution of AI-urban capability is non-ergodic, so cities that move early benefit from compounding returns, while those that hesitate—protecting incumbent structures or waiting for national directives—risk being relegated to adopting models designed, tested and optimised elsewhere [16].

The remainder of the paper is organised as follows. Section 2 presents the conceptual framework, integrating the extended mirroring hypothesis with theories of urban innovation ecosystems and systems thinking; Section 2.4 formalises the five propositions of cumulative recursive hybridisation, and Section 2.5 demarcates the framework from adjacent traditions (general-purpose technologies, organisational complementarities, and complex adaptive systems). Section 3 examines the first vector—the evolution of public-sector conversational AI—proposing a four-level maturity model and drawing on comparative international cases. Section 4 addresses the second vector: autonomous electric mobility and its implications for urban spatial structure, cost, availability and policy. Section 5

explores the third vector: robotics and intelligent infrastructure in the urban environment. Section 6 develops the cross-cutting argument—the city as the locus of transformation—analysing the dynamics of cumulative recursive hybridisation, the local embedding of innovation ecosystems and the conditions that separate pioneering cities from lagging ones; Section 6.5 confronts the framework with counterfactual evidence from failed, stalled and non-Western trajectories, and Section 6.6 examines the political-economy, equity and surveillance limits of the optimistic reading. Section 7 concludes with implications for urban policy and directions for future research.

The article is therefore a conceptual and theory-building contribution: its purpose is to formalise cumulative recursive hybridisation (CRH), clarify its boundary conditions, and render its mechanisms empirically testable, rather than to claim causal identification through new modelling or primary data. This framing guides the more explicit operationalisation and comparative standardisation introduced below.

2. Conceptual Framework: Institutional Symmetry, Urban Innovation, and the Dynamics of Recombination

Understanding how agentic AI transforms cities requires more than a technology-adoption lens. It demands a conceptual framework that accounts for (a) the organisational conditions under which institutions capture value from new technologies, (b) the spatial dynamics that make cities uniquely fertile environments for innovation, and (c) the systemic mechanisms through which diverse technological domains cross-fertilise and compound. This section integrates three bodies of theory—the mirror hypothesis of institutional symmetry, the literature on urban innovation ecosystems, and the logic of intersectional recombination—to construct the analytical foundation that will guide the subsequent empirical sections (Box 1).

Box 1. Operational definitions of key terms.

Agentic AI refers to AI systems—typically built on large language or multimodal foundation models—that can decompose open-ended goals into sub-tasks, plan over multiple steps, invoke tools or APIs, and act on external systems on behalf of a principal under defined authority constraints. The defining feature is autonomous action across multi-step workflows, not merely content generation [3,17,18].

Cognitive government denotes a configuration of public administration in which AI agents are first-class participants in service delivery and decision support, operating under explicit democratic mandates, audit trails, and accountability mechanisms. It corresponds to Level 4 of the maturity model developed in Section 3 and is empirically incipient rather than mature [10].

City brain designates an integrated municipal data and analytics platform that ingests real-time signals from sensors, cameras, connected vehicles, and administrative systems, applies machine-learning and rule-based reasoning, and triggers operational decisions or actuator commands in domains such as traffic management, environmental monitoring, and emergency response [19,20].

Extended mirroring hypothesis—our reformulation of Conway’s Law and the Colfer-Baldwin formalisation—holds that organisations and ecosystems converge toward the configuration of human and AI agents, coordination mechanisms, and decision architectures that best exploits the possibilities opened by a new technology, and that the degree of alignment determines value capture [21,22].

Cumulative recursive hybridisation, defined formally in Section 2.4 below, denotes a dynamic in which multiple co-located technological vectors interact through bidirectional feedback loops in data, regulation, infrastructure, and talent, such that each cycle increases the system’s capacity for the next.

2.1. The Mirroring Hypothesis Extended: From Communication Structures to Agentic Ecosystems

Our conceptual point of departure is Conway’s foundational observation that “organisations which design systems are constrained to produce designs which are copies of the communication structures of these organisations” [21]. Conway’s Law, as it came to be known, established a powerful principle: the architecture of a technical system reflects

the social structure that produced it. Colfer and Baldwin [22] subsequently formalised this regularity as the *mirroring hypothesis*, demonstrating through a systematic review of empirical studies across industries that the alignment between organisational structure and technical architecture is not merely a constraint but an *equilibrium condition*—when the two are in symmetry, coordination costs fall and performance improves; when they diverge, friction and underperformance follow. Convergent findings in the technology-adoption literature reinforce this pattern: the “productivity paradox” of the 1980s–1990s showed that enterprise computing investments yielded returns only when firms restructured workflows and incentive systems in parallel [23,24]; Milgrom and Roberts theorised this as organisational complementarity [25]; and Brynjolfsson and colleagues demonstrated empirically that IT productivity requires co-investment in decentralised decision-making and workforce reorganisation [26].

However, we argue that both Conway’s original formulation and Colfer and Baldwin’s empirical formalisation, while foundational, describe a static version of mirroring—a tendency for existing structures to be reflected in existing designs. In this paper, we extend the mirroring hypothesis in two directions that are essential for understanding the impact of agentic AI on cities.

The first extension is dynamic and strategic. Organisations and ecosystems do not merely mirror their current communication structures in the systems they produce. When a genuinely new technology opens a previously unavailable space of possible configurations—new ways of coordinating, new divisions of labour, new modes of decision-making—organisations explore this space and tend to converge, through competitive pressure and institutional learning, toward the strategic structures that best exploit the technology’s possibilities [27,28]. The mirroring, in this view, is not backward-looking (reflecting what exists) but forward-looking (converging toward what is now possible). The introduction of the internet did not simply mirror pre-existing corporate hierarchies into websites; it opened a space of networked, platform-based, and disintermediated organisational forms that firms progressively discovered and adopted—those that found the best configurations thrived, and those that clung to pre-internet structures were displaced [29]. The productivity paradox was, in this light, a period of exploration: the lag between technology availability and productivity gains reflected the time required for organisations to search, experiment, and converge on the organisational architectures that mirrored the technology’s actual possibilities rather than their inherited structures.

The second extension is ontological, and it is specific to the agentic AI era. In all prior technological transitions, the agents doing the coordinating, exploring, and decision-making within organisations were exclusively human. The communication structures that Conway described, the organisational units that Colfer and Baldwin mapped to technical modules, the complementary practices that Milgrom and Roberts theorised—all of these involved human actors coordinating with other human actors. Agentic AI fundamentally alters this premise. For the first time, the participating agents in an organisational ecosystem include not only humans but also artificial agents capable of autonomous reasoning, decision-making, and action [17,18]. An agentic chatbot that processes a citizen’s permit application end-to-end is not a tool used by a human agent—it is an agent, with a defined role, a scope of authority, and interaction protocols with other agents (human and artificial) in the system. A fleet-management AI that coordinates robotaxis across a city is a coordination agent operating within the ecosystem’s architecture. An urban intelligence platform that synthesises data from autonomous vehicles, robotic sensors, and administrative systems to trigger maintenance workflows is an orchestrating agent that mediates the interactions of multiple subsystems.

This means that the organisational architecture that must mirror the technology is no longer composed solely of human roles, human communication channels, and human decision rights. It is a hybrid architecture of human and AI agents, where the roles, boundaries, interaction protocols, and coordination mechanisms must be designed for both kinds of participants [10,30]. The mirroring hypothesis, extended to the agentic era, thus states: *organisations and ecosystems will tend to converge toward the configuration of human and AI agents, coordination mechanisms, and decision architectures that best exploits the possibilities opened by agentic technologies—and the degree to which they achieve this alignment will determine the value they capture.*

We further extend this logic beyond the boundary of the individual organisation to the level of the city as an ecosystem. A city is, in essential respects, an organisational architecture: it comprises governance structures, regulatory frameworks, service delivery systems, infrastructure networks, labour markets, and institutional cultures, all of which are interdependent and jointly shape the city's capacity to absorb, adapt, and benefit from new technologies. In the agentic AI era, a city's ecosystem includes not only its human actors—public servants, entrepreneurs, researchers, citizens—but also the growing population of AI agents that mediate, execute, and coordinate urban functions: administrative chatbots that handle citizen interactions, autonomous vehicles that navigate city streets, robotic systems that maintain infrastructure, and intelligence platforms that synthesise urban data. The city that achieves mirroring is the one whose institutional architecture—its regulatory agility, data-sharing frameworks, coordination protocols, and workforce capabilities—evolves in symmetry with this hybrid human–AI ecosystem. When this symmetry is present, the city captures compounding value from the recursive interaction of its three transformation vectors.

When it is absent, the city falls into an asymmetry trap: powerful agentic technologies constrained by organisational structures designed for an exclusively human coordination logic.

For public administrations specifically, this extended mirroring hypothesis carries particular weight. Government agencies operate under constraints—legislative mandates, civil-service rules, procurement frameworks, accountability regimes—that make organisational reconfiguration slower and more politically fraught than in the private sector [8,31]. These constraints were designed for organisations composed entirely of human agents, and they do not easily accommodate AI agents that need defined roles, authority boundaries, audit trails, and accountability mechanisms of their own. This structural rigidity explains why most public-sector chatbot deployments worldwide remain at the informational level (Level 1), merely layering conversational interfaces onto unchanged back-end systems, rather than progressing toward the transactional and agentic architectures (Levels 3–4) that would require deep institutional redesign—including redesigning the organisation to incorporate AI agents as legitimate participants in service delivery and decision-making [10,29]. The extended mirroring hypothesis thus predicts that the cities achieving the greatest impact from agentic AI will be those that pursue not merely technological deployment but the co-evolution of their entire institutional ecosystem—human roles, AI agent roles, coordination architectures, and governance frameworks—toward the configurations that best exploit the possibilities of the agentic era. Figure 1 synthesises the structure of the extended mirroring hypothesis, showing how Conway's foundational observation and Colfer and Baldwin's formalisation are extended in the dynamic/strategic and ontological directions.

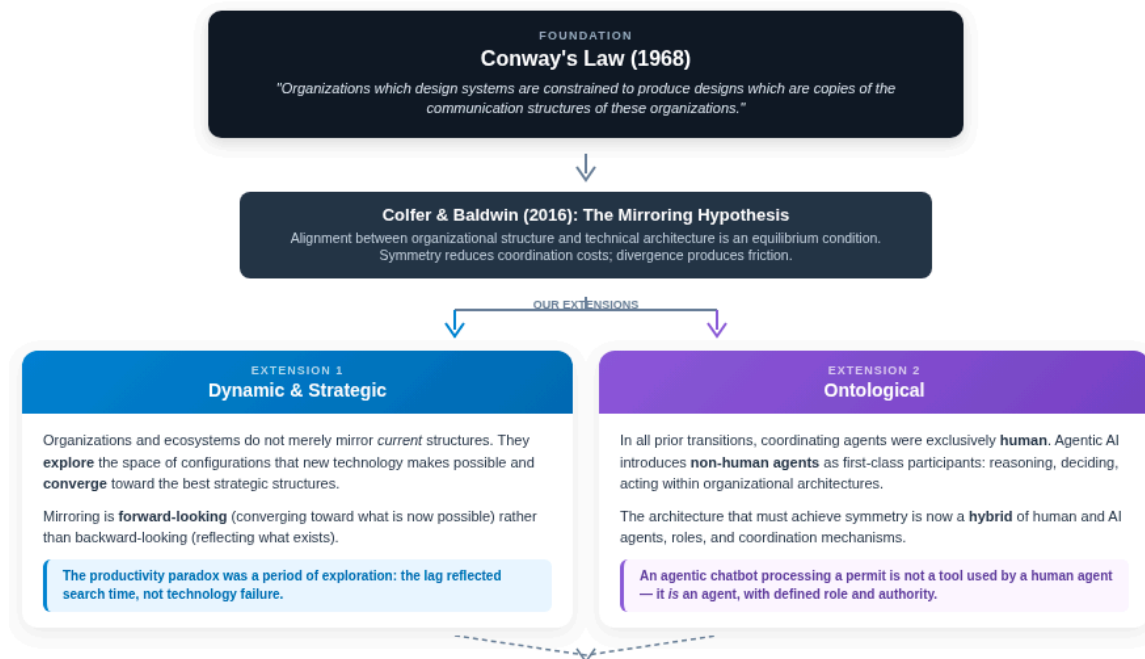


Figure 1. The extended mirroring hypothesis: from communication structures to agentic ecosystems. Conway’s Law and Colfer and Baldwin’s formalisation are extended dynamically (organisations converge toward configurations new technology enables) and ontologically (AI agents join humans as coordination participants). Key takeaway: the organisational architecture that must ‘mirror’ the technology is no longer purely human—it is a hybrid of human and AI agents, and cities must redesign governance structures accordingly.

2.2. Cities as Innovation Ecosystems: Embedded Knowledge, Talent, and the Local Roots of Transformation

Why should the city, rather than the nation-state or the firm, be the primary unit of analysis for AI-driven transformation? The answer lies in a long tradition of scholarship on the geography of innovation, which consistently demonstrates that breakthrough innovation is not spatially neutral—it clusters in specific places, and those places are overwhelmingly cities [14,28,29].

Jane Jacobs argued half a century ago that cities are engines of economic life precisely because they bring together diverse activities in dense proximity, enabling the kind of unexpected cross-pollination that planned economies and isolated institutions cannot replicate [32]. Her insight anticipated what evolutionary economists would later formalise as *related variety*: the principle that innovation thrives not in environments of pure specialisation (where everyone does the same thing) nor in environments of pure diversity (where there is no common language), but in ecosystems where distinct but complementary knowledge domains overlap and recombine [33,34].

This dynamic is sustained by several mechanisms that are intrinsically urban. First, embedded knowledge: the tacit know-how, institutional memory, and relational capital that accumulate in a place over time and cannot be easily transferred or replicated elsewhere [35]. The reason Silicon Valley dominates venture-backed technology, or that Shenzhen leads in hardware prototyping, is not simply a matter of policy incentives—it reflects decades of accumulated expertise, supplier networks, shared mental models, and trust architectures that are woven into the social fabric of the city [36].

Second, talent density and circulation: cities attract, develop, and retain specialised human capital, and the physical proximity of skilled workers in complementary domains accelerates the informal knowledge exchange—the corridor conversation, the accidental

encounter at a conference, the engineer who moves from a robotics startup to a transit authority—that seeds novel combinations [15,33]. Third, fast supply chains and institutional proximity: in a city, a startup developing autonomous delivery vehicles operates within reach of the municipal authority that issues permits, the university lab that tests sensor arrays, the logistics company that provides pilot routes, and the venture fund that supplies capital. This institutional proximity compresses the iteration cycle from idea to prototype to deployment in ways that remote or dispersed arrangements cannot match [37].

In Marshallian terms, such concentrations generate external economies of agglomeration, while Storper and Venables show how urban “buzz” and face-to-face interaction accelerate the circulation of tacit signals, reputational information, and opportunities within dense local ecosystems [38–40].

A critical and often underappreciated dimension of this argument is that knowledge and talent do not exist as abstract, portable assets—they are constituted through participation in ongoing projects and practices. Nonaka and Takeuchi’s foundational work on organisational knowledge creation demonstrated that the most valuable forms of expertise—tacit knowledge—are generated through socialisation and externalisation within active work processes, not through formal training or documentation alone [41].

Lave and Wenger extended this insight through their concept of communities of practice, showing that expertise develops through legitimate peripheral participation: newcomers become skilled practitioners not by studying a body of knowledge but by progressively engaging in real, situated activity alongside experienced colleagues [42]. More recently, project-based learning theories have shown that in fast-moving technological domains, competence is inseparable from the projects in which it is exercised—engineers, data scientists, and urban planners develop frontier capabilities only when there are live deployments to work on, real problems to iterate against, and functioning teams embedded in ongoing operations [43,44].

The implication for cities is profound. A city that cancels its autonomous bus pilot, shelves its chatbot initiative, or fails to launch robotics experiments does not merely lose time—it loses the very substrate in which talent forms. The engineers who would have learned to calibrate LiDAR in local traffic conditions, the public servants who would have developed expertise in AI-mediated service design, the urban planners who would have gained intuition for robotics-infrastructure integration—none of these capacities develop in the abstract. They develop in projects. And when the projects disappear, the talent either fails to form or migrates to cities where the projects exist. This creates a self-reinforcing dynamic: cities with active deployments attract and develop talent, which enables more ambitious deployments, which further deepens the talent pool—while cities without projects enter a vicious cycle of capacity erosion [15,45].

For the AI-driven transformation of cities, these dynamics are especially potent because the three vectors we examine—governance, mobility, and robotics—are all place-based: they are deployed in, regulated by, and experienced within specific urban territories. Unlike cloud software, which can be adopted from anywhere, an autonomous robotaxi fleet depends on local road infrastructure, local traffic regulation, local mapping data, and local citizen acceptance. An agentic public administration depends on local regulatory authority, local data registries, and local institutional culture. Urban robotics depends on local infrastructure conditions, local labour markets, and local procurement processes. The city is not merely the setting for these innovations—it is a constitutive element of their development.

2.3. *The Medici Effect and Intersectional Innovation in the Urban Context*

Frans Johansson's concept of the *Medici Effect* offers a complementary lens [45]. Johansson argued that the most transformative innovations in history have occurred at the intersection of disciplines, cultures, and domains—just as the Medici family's patronage in Renaissance Florence brought together sculptors, scientists, poets, financiers, and architects, creating an environment in which ideas from one domain sparked breakthroughs in another. The key insight is that innovation at intersections is qualitatively different from innovation within a single domain: it is less predictable, more combinatorial, and often more disruptive, because it escapes the path dependencies and mental models that constrain domain-specific thinking [45,46].

Applied to the smart city, the Medici Effect suggests that the most significant transformations will not come from advancing any single vector in isolation—a better chatbot, a faster robotaxi, a more efficient street-cleaning robot—but from the intersections among them. Consider: when an agentic public administration can process autonomous-vehicle licencing in real time, it removes a bottleneck that currently delays deployment by months or years. When an autonomous logistics network shares data with a city brain platform, it enables predictive infrastructure maintenance that neither system could achieve alone. When robotic street maintenance operates on schedules optimised by the same AI that manages traffic flow, both systems become more efficient. These intersectional innovations are not planned from the top down; they emerge from the proximity, density, and diversity of the urban ecosystem—from the fact that the people building governance chatbots, the engineers deploying robotaxis, and the teams developing maintenance drones work in the same city, attend the same events, share the same infrastructure, and face the same regulatory environment.

This is why we argue that cities function as innovation intersections in the Johansson sense: they are the places where the Medici Effect operates at scale, because they concentrate the diversity of domains, the density of talent, and the institutional infrastructure needed for intersectional recombination to occur.

2.4. *Systems Thinking and Cumulative Recursive Hybridisation*

The final element of our conceptual framework draws on systems thinking—the recognition that complex systems exhibit emergent properties that cannot be predicted from the behaviour of their individual components [47,48]. In a systems perspective, the smart city is not a collection of independent technological applications (a chatbot here, a robotaxi there, a drone over there) but an interconnected system in which changes in one domain propagate through feedback loops to reshape others.

We synthesise the preceding theoretical elements into the concept of cumulative recursive hybridisation: a dynamic in which multiple technological domains, co-located within an urban ecosystem, interact through iterative cycles of recombination, each cycle building on the outputs of the previous one and generating compounding returns. The term draws on the historiography of the Industrial Revolution, where scholars have shown that Britain's transformation was not driven by any single invention—not the steam engine, nor the spinning jenny, nor the blast furnace—but by the recursive interaction among all of them within concentrated geographic clusters that attracted talent, accelerated knowledge circulation, enabled cross-fertilisation across domains, and gave rise to entirely new industries and institutional forms [12,13,45,46]. Those clusters became the hubs of a new economic order, not because they adopted technologies faster but because they created the ecosystem conditions—talent density, speed of iteration, institutional plasticity, tolerance for failure—in which recombination could compound. The same logic applies to the agentic AI era. Cities that simultaneously host agentic governance, autonomous mobility and urban

robotics are not simply deploying three technologies; they are constructing the ecosystems in which hybridisation, cross-fertilisation and the emergence of novel proposals become routine, and in doing so they are positioning themselves as the innovation hubs of a new urban order.

The analogy is direct. In a city that adopts all three vectors simultaneously, the following feedback loops emerge: (a) *data loops*: autonomous vehicles and urban robots generate continuous streams of urban data—traffic patterns, road conditions, air quality, pedestrian flows—that feed city brain platforms and improve governance decision-making; (b) *regulatory loops*: agentic administrations capable of real-time licencing and adaptive regulation accelerate the deployment and iteration of mobility and robotics services; (c) *infrastructure loops*: robotic maintenance and intelligent infrastructure management improve the physical environment on which autonomous mobility depends; and (d) *talent loops*: the presence of frontier deployments in multiple domains attracts and retains the specialised workforce that sustains the ecosystem. These loops are recursive—each cycle increases the system’s capacity for the next—and cumulative—the gains compound over time, creating increasing returns that widen the gap between pioneering cities and lagging ones [49].

To render cumulative recursive hybridisation analytically tractable—and falsifiable—we restate the framework as five testable propositions that follow from the four feedback loops described above (Figure 2). Each proposition specifies a mechanism, an empirical referent, and a condition under which it would fail.

Proposition P1 (data-substrate marginal cost). Once a shared urban data substrate exists for any one vector v_i (e.g., a city brain platform serving traffic), the marginal cost of adding a second vector v_j that consumes the same substrate is strictly lower than the cost of building v_j in isolation. Failure condition: cities where data is technically integrated but legally siloed (data-protection rules, departmental ownership) should not exhibit the marginal cost reduction—an empirical test we encourage.

Proposition P2 (action-space composition). Each new vector $v_{\{n+1\}}$ expands the action space of every existing vector via API or workflow composition. Concretely, the joint operational repertoire $R(V_n)$ of n vectors satisfies $|R(V_{\{n+1\}})| > |R(V_n)| + |R(v_{\{n+1\}})|$, i.e., composition is super-additive. Failure condition: vectors deployed without inter-system standards or shared identity layers should exhibit only additive, not super-additive, growth.

Proposition P3 (regulatory throughput as bottleneck). The realised compounding rate is bounded above by the throughput of municipal regulatory decision-making. Cities at maturity Levels 1–2 of agentic governance cannot capture the full hybridisation premium even when the technical substrate is mature. Failure condition: a city with fast regulatory throughput but isolated technical pilots should still under-perform; a city with rich technical pilots but Level-1 administration should also under-perform—both are observable.

Proposition P4 (talent recursion). Specialised AI-urbanism talent forms only inside live deployments and grows monotonically with the number of concurrent deployments in the same labour market. The exit of any single anchor deployment from a city reduces the local talent stock by more than the project’s nominal headcount, because tacit knowledge and network ties dissolve. Failure condition: cities that cancel anchor deployments should exhibit measurable engineer-migration outflows within 24 months.

Proposition P5 (path dependency and divergence). Because P1–P4 are recursive and bidirectional, the inter-city distribution of AI-urban capability is non-ergodic: small early differences in deployment intensity compound into large persistent gaps. Failure condition: a latecomer city that deploys all three vectors simultaneously and at scale should be able to close the gap with early leaders within a decade; persistence of the gap despite such investment would falsify P5.

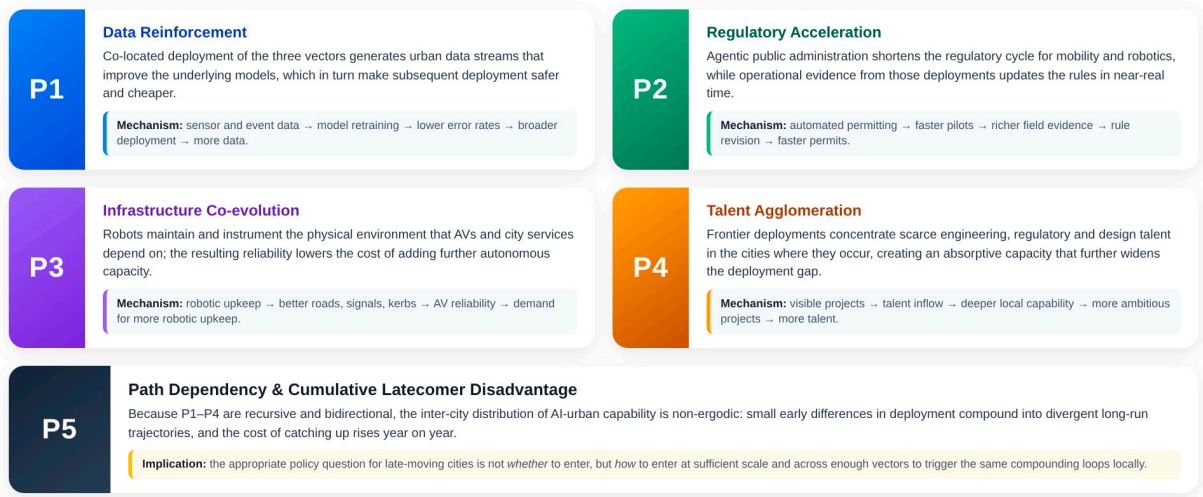
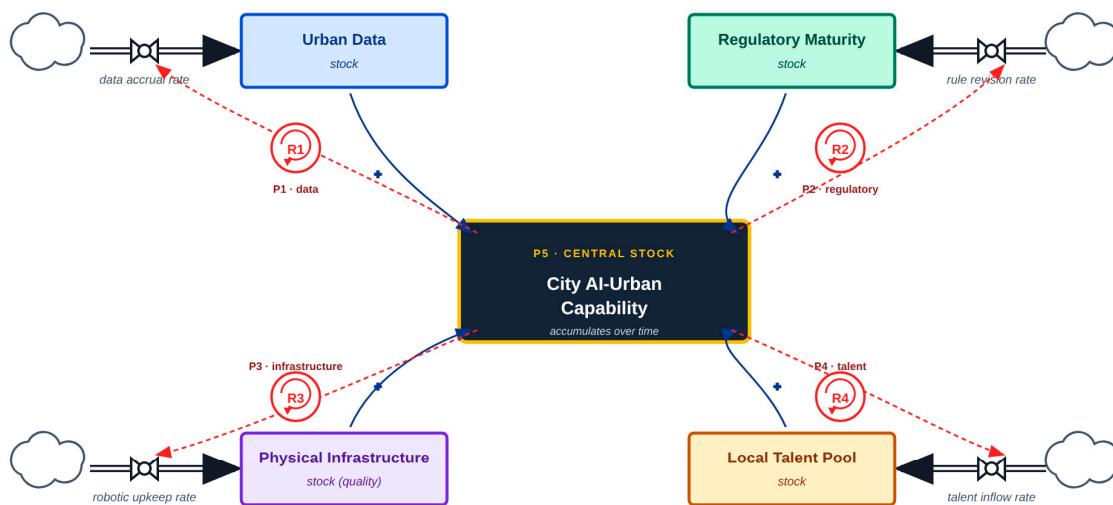


Figure 2. The five formal propositions of Cumulative Recursive Hybridisation (CRH). P1–P4 specify the four reinforcing feedback loops generated by the co-deployment of agentic governance, autonomous mobility and urban robotics; P5 is a derived claim about the joint dynamics of the system as a whole. Each proposition is testable through cross-city panel evidence on deployment scale, talent flows, regulatory cycle times and capability indicators.

These propositions are not yet operationalised through system-dynamics simulation (Figure 3) or panel econometrics—both reviewers correctly identified this as the appropriate next step, and we agree. The contribution of the present paper is to formulate the framework precisely enough that such tests become possible. Section 7 returns to this research agenda.



Each reinforcing loop closes through City AI-Urban Capability; the four loops compound, producing the path-dependent divergence stated in P5.

Stock (level / accumulation) Flow / rate (valve) Source / sink (cloud) Information link (+ polarity) Reinforcing feedback (R-loop)

Figure 3. System-dynamics representation of Cumulative Recursive Hybridisation (Forrester convention). Four stocks accumulate the city’s AI-urban capability; each inflow rate is governed by an information link from the central capability stock, closing reinforcing loops R1–R4 (propositions P1–P4). P5 is the integral over time. Key takeaway: because each loop feeds back into the others, early deployment advantages compound—the diagram shows why latecomers face an accelerating rather than constant disadvantage.

To make the empirical agenda more concrete, Table 1 translates P1–P5 into illustrative indicators, candidate units of analysis, and feasible research designs. The table is not intended as a full measurement model; rather, it shows how the framework can be operationalised while preserving the paper’s conceptual and theory-building purpose.

Table 1. Illustrative operationalisation of the five propositions of cumulative recursive hybridisation.

Prop.	Core Mechanism	Illustrative Indicators	Unit/Data	Suggested Test or Failure Signal
P1	Shared data substrate lowers the marginal cost of adding adjacent vectors.	API re-use; interoperable datasets; time/cost required to add a new vector.	City programme or project level; procurement files; integration logs; platform documentation.	Compare cities with technically integrated versus legally siloed data regimes. No cost reduction would weaken P1.
P2	New vectors expand the joint action space through workflow or API composition.	Number of cross-system workflows; new combined services; extent of identity-layer re-use.	City-service portfolio; process maps; platform architecture; municipal service catalogues.	Test whether joint deployment yields more than additive service expansion. Purely additive growth would weaken P2.
P3	Regulatory throughput bounds realised compounding.	Median permit cycle time; number of exemptions or approvals; frequency of rule updates.	Municipal regulatory units; permit registries; council records; administrative logs.	Relate deployment intensity to decision speed. Rich pilots with slow approvals, or fast approvals with isolated pilots, should under-perform.
P4	Talent forms inside live deployments and decays when anchor projects exit.	Specialised job postings; engineer mobility; project exits; spin-outs; training pipelines.	Metropolitan labour market; LinkedIn/job data; company announcements; ecosystem mapping.	Track whether cancellation of anchor projects is followed by talent outflows or reduced local project formation.
P5	Early differences compound into persistent inter-city divergence.	Change in deployment intensity; cumulative rides/robots/services; persistence of capability gaps.	Cross-city panel; repeated case coding; longitudinal ecosystem indicators.	Assess whether latecomers that invest across all three vectors can close gaps within a decade. Persistent divergence despite scale would weaken P5.

2.5. Demarcation: How Cumulative Recursive Hybridisation Differs from Adjacent Frameworks

Cumulative recursive hybridisation (CRH) is adjacent to three established traditions—general-purpose technologies (GPT), organisational complementarities, and complex adaptive systems (CAS)—and a sharp delineation is essential to clarify what is genuinely novel. Bresnahan and Trajtenberg’s GPT framework [50] explains why a single pervasive technology (electricity, the computer) drives broad productivity growth through innovational complementarities across user industries. CRH differs in two respects: it concerns multiple co-evolving technologies rather than one, and its unit of analysis is the city as a bounded ecosystem rather than the national economy.

Milgrom and Roberts’ theory of organisational complementarities [25] explains how clusters of practices reinforce one another within a firm; CRH applies a related logic but at the inter-organisational, intra-municipal scale, and centres on bidirectional data and regulatory loops that link autonomous agents rather than static pairwise complementarities of practices and incentives. Holland’s complex adaptive systems framework [51] supplies the language of feedback, emergence, and non-ergodicity that CRH inherits, but CAS remains descriptive rather than prescriptive: it does not tell us

which feedback loops matter, why municipal scale is binding, or which institutional levers shift system behaviour.

The novelty of CRH lies in the conjunction of three claims that none of these adjacent frameworks makes simultaneously: (i) recursion in the strict sense that outputs of one vector become training inputs for another, generating compounding rather than merely complementary returns; (ii) the city—not the firm or the nation—as the binding unit because regulatory authority, deployment data, and tacit talent are constituted at municipal scale; and (iii) the agentic ontology of Section 2.1, in which AI agents are first-class participants in the loops rather than tools wielded by human actors.

This framework suggests that the cities most likely to lead the agentic AI transformation are not necessarily those with the largest budgets or the most advanced national policies, but those that cultivate the conditions for cross-domain recombination: institutional openness to experimentation, regulatory agility, talent density in complementary fields, and a willingness to allow—and learn from—failure. More fundamentally, the framework predicts that what is emerging is not a technological upgrade but a new urban order. The cities that succeed in closing the four feedback loops—data, regulation, infrastructure, talent—will function as innovation hubs in which talent concentrates, knowledge circulates at high speed, cross-fertilisation produces hybrid proposals, and the institutional architecture co-evolves with the technology. These cities will set global standards, attract global talent, and export institutional templates, much as the industrial cities of the eighteenth century did. And the divergence between these hubs and cities that hesitate is likely to accelerate with each passing year, because each feedback loop amplifies the others.

Section 6 will return to these dynamics with empirical evidence, examining how specific cities have—or have not—created the conditions for cumulative recursive hybridisation to take hold.

3. AI and Public Administration: From Chatbots to Cognitive Government

Within the CRH framework developed above, the first vector—agentic governance—occupies a distinctive position: it is both a transformation in its own right and the regulatory and institutional substrate through which the other two vectors are permitted, shaped, and governed. We now examine how this vector has evolved in practice.

The first vector of urban transformation concerns the interface between public administrations and citizens. Over the past decade, conversational systems in the public sector have evolved from rule-based bots and dynamic FAQs into sophisticated assistants powered by natural language processing and, increasingly, by generative AI [52,53]. This trajectory signals more than a technological refresh: it entails organisational, cultural, and social change in how administrations interface with citizens and execute core processes. Framed through the extended mirroring hypothesis developed in Section 2, the progression from informational chatbots to agentic governance systems represents a gradual—and still largely incomplete—alignment between the possibilities opened by AI technology and the organisational architectures of the public sector.

3.1. A Functional Maturity Model for Public-Sector Conversational AI

To anchor the analysis, we propose a four-level maturity model (Figure 4) that maps the capabilities, organisational implications, and citizen value of public-sector conversational systems [10,52]. The model is cumulative: each level builds on the capabilities of the preceding one, but the organisational and institutional demands increase substantially at each stage.

Level 1—Guided informational. At this foundational level, conversational systems reduce search frictions and standardise responses, typically complementing municipal websites without requiring complex back-end integrations. The system provides structured FAQs, basic triage, and links to forms. The organisational demand is minimal: content governance, editorial standards, and analytics for topic/intent distribution. The value proposition centres on 24/7 availability and consistency, deflecting repetitive contacts from telephone and counter channels. This is the modal implementation worldwide—the majority of public-sector chatbots currently operate at this level [52,53].

Level 2—Contextual guidance (NLP/RAG). At this level, assistants move beyond generic answers to understand user intent, retrieve relevant regulation or documentation, and provide contextually adapted responses. Architecturally, this requires generative models operating over retrieval-augmented generation (RAG) pipelines, drawing on curated corpora of policy, regulation, and procedures with source-linked provenance [54]. The organisational demand rises considerably: knowledge curation pipelines, policy/legal review workflows, document versioning, and prompt/response evaluation become necessary. The gain is a marked reduction in ambiguity and error rates, and a more personalised citizen experience.

Level 3—Assisted transactions. The qualitative leap occurs here: chatbots initiate or complete specific procedural steps—booking appointments, filing requests, processing payments, handling document submissions—integrating securely with transactional back-end systems [10,55]. This level requires API-first integration with line systems (registries, scheduling, payments), robust identity verification, explicit consent management, and auditable event logging. The organisational implications are profound: the chatbot is no longer a communication channel but a transaction operator, and its failures have direct consequences for citizens. Product-oriented operating models emerge—conversation design, service choreography, incident response, and security and privacy impact assessments become routine functions.

Level 4—Cognitive-agentive (agents with data and action). At the apex of the model, conversational agents operate autonomously on behalf of citizens, orchestrating multi-system processes with traceability and safeguards [10,17]. The assistant does not merely respond to requests but proactively identifies eligibility, sends renewal warnings, coordinates across agencies, and executes multi-actor workflows under governed mandates. The architecture demands interoperable data spaces, policy-as-code for eligibility and safeguards, real-time analytics, and multimodal interfaces (text, voice, video, AR). This is, in the language of our extended mirroring hypothesis, the level at which AI agents become first-class participants in the organisational architecture of public administration—not tools used by human agents, but agents in their own right, with defined roles, authority boundaries, and accountability mechanisms.

The transition across levels is not merely a function of better language models. Moving from Levels 1–2 to Levels 3–4 requires not only improved natural language understanding but also identity, consent, audit, and interoperability layers—institutional capabilities that translate model performance into public value with accountability [10]. Figure 4 summarises the four levels, their architectural requirements, organisational demands, and representative international cases.



Figure 4. Four-level maturity model for public-sector conversational AI. Each level builds on the capabilities of the preceding one, but the organisational and institutional demands increase substantially. Representative city deployments illustrate the current international landscape at each level.

3.2. Comparative International Landscape

A comparative analysis of chatbot deployments across Europe, Latin America, and Asia reveals three empirical patterns that reinforce the maturity model and illuminate the dynamics of the extended mirroring hypothesis [10,52].

This pattern is consistent with broader evidence that AI use in government remains uneven and is still concentrated primarily in assistance, information, and decision-support functions rather than in fully integrated, high-autonomy service delivery [56].

The dominance of Level 1. The majority of implementations worldwide remain at the guided informational level. Systems such as Línea Madrid (Spain), WienBot (Vienna), Bobbi (Berlin), TMBbot (Barcelona), Govbot (Japan), Divinha (Curitiba), and Jugalbandi (India) provide standardised information, reduce contact-centre load, and offer basic navigation. They deliver quick wins in availability and consistency, but yield limited transformation when kept siloed from back-end systems. Their persistence at Level 1 reflects, in mirroring terms, an organisational architecture that has not yet reconfigured to match the possibilities of the technology.

The emergence of Level 2–3 as a contested frontier. A smaller but growing cohort of deployments has advanced to contextual guidance or assisted transactions. In Europe, Estonia’s Bürokratt stands as the singular Level 4 reference—a platform-of-platforms architecture that integrates services across agencies and channels, including voice, toward a single national assistant. Clara (Madrid) and Noa (Île-de-France) operate at Level 2 with contextual retrieval capabilities. In Latin America, Buenos Aires’ Boti is a leading Level 3 case on WhatsApp, handling over 300,000 interactions per month, including appointment bookings, filings, and alerts; Mexico City’s TEO focuses on anti-corruption reporting (Level 3); and Bogotá fields multiple systems at different levels (Chatico at Level 2, Rebeca at Level 1). In Asia, Dubai’s Rammas supports billing and inquiries with multi-language capability and deep system integration (Level 3); Singapore’s VICA centralises multi-agency assistance (Level 2); and South Korea’s OneService Chatbot enables complaint filing and service reservations (Level 3) [10,52].

Regional patterns of strategic differentiation. Europe leads in the diversity of approaches, with cases spanning from informational to early agentic, and a notable concentration of experimentation in Spain (Las Rozas’ Miguel, Ciudad Real’s Prado with

80+ languages and 3500 sources, the Open Administration of Catalonia's shared generative service). Latin America stands out for high adoption via WhatsApp—a channel strategy that meets citizens “where they are” and proves decisive for inclusion and scale. Asia drives large-scale, multi-service integration projects. Across all regions, the geographic concentration of advanced deployments in a few jurisdictions confirms that policy clarity, shared infrastructure, and institutional programme management capacity are differentiators for moving from demonstration to institutionalisation [10,52].

3.3. Three Scenarios for Implementation

Building on the maturity model and the comparative evidence, three implementation scenarios delineate the trajectories that local and regional administrations may follow. These are not mutually exclusive; a single jurisdiction may occupy different scenarios across policy domains and evolve asynchronously over time [10].

Conservative: incremental optimisation. Chatbots act as guided informational interfaces that standardise answers and reduce search frictions without altering back-office structures. Maturity corresponds to Level 1 and selectively Level 2. Benefits include low cost and rapid deployment; limitations include minimal end-to-end resolution and potential citizen frustration when transactions are expected but unavailable. Illustrative cases include Vienna and Madrid's informational assistants.

Disruptive: assisted transactions and mass personalisation. Chatbots become transaction operators that initiate or complete procedural steps, personalise interactions, and interoperate with multiple line systems in real time. Maturity is Level 3 with early Level 4 features. Cycle times shrink, staff effort shifts to higher-value work, and the citizen experience approaches that of private-sector digital services. The organisational prerequisites are substantial: robust data governance, cybersecurity, change management, and identity/consent layers. Representative cases include Boti (Buenos Aires), Rammas (Dubai), and elements of the Bürokratt platform [10,52].

Systemic: networked cognitive government. Conversational agents become components of an inter-organisational cognitive infrastructure where public agencies, private providers, and civil society co-produce seamless services. Maturity aligns with Level 4. Value arises from proactive, preventive administration—eligibility nudges, renewal warnings, multi-actor workflows—and ecosystem orchestration under stringent governance. The administration ceases to be a building that citizens must visit and becomes a network of agents that accompany them throughout their lives [52]. This scenario represents the fullest expression of the extended mirroring hypothesis: the organisational architecture of government has been redesigned to incorporate AI agents as legitimate participants, with coordination mechanisms, decision rights, and accountability structures that reflect the possibilities of the agentic era. National exemplars remain rare; Estonia's interoperable layer and Singapore's multi-ministry approach are the closest references.

3.4. Toward the Windowless Administration: The Agentic Governance Vision

The trajectory described above points toward a vision of public administration fundamentally different from the multi-channel bureaucracy that most citizens experience today. In this vision—what the Anteverti-Esade scenarios report terms the “windowless administration” [52]—the administration has no service windows, no websites in the traditional sense. Instead, citizens interact through conversational, multimodal interfaces that adapt to and anticipate their needs. The burden of navigating institutional complexity shifts from the resident to the system: interfaces become adaptive (text, voice, rich interactions), policy logic is encoded as executable rules that guide end-to-end processes, and a resilient

orchestration layer routes each request to purpose-built agents that consult authoritative sources, apply policy-as-code, and perform actions under role-based permissions [10].

This is not merely a user-interface improvement. It represents a fundamental reframing of government—from a channel-centric service provider to an agent-orchestrated system that learns from context, anticipates needs, and executes on citizens' behalf within clear democratic mandates. The operating layer of government changes: it becomes a fabric of human–AI teams that coordinate across agencies, data sources, and service partners with accountability and auditability [10].

Realising this vision requires the co-evolution of technology and organisational architecture that the mirroring hypothesis predicts. Administrations that retain paper-era hierarchies while adding conversational front-ends will achieve, at best, a veneer of modernization without gains in responsiveness, coordination, or resilience. Those that redesign their structures—decision rights, roles, KPIs, coordination mechanisms—to incorporate AI agents as active nodes in traceable workflows will be the ones that deliver proactive, personalised services at scale. The practical test of progress, as the IJEBM analysis argues, is the movement from reactive case handling to preventive administration, measured through time-to-complete, first-contact resolution, and equity of access [10].

4. Autonomous Electric Mobility and the Reshaping of Urban Space

The CRH framework predicts that the three vectors reinforce each other through shared data, regulation, infrastructure, and talent loops (P1–P4). Having examined the governance vector, we turn to the second: autonomous electric mobility, whose deployment both depends on and feeds back into the regulatory and data capabilities that agentic governance provides.

The second vector of urban transformation concerns how people and goods move through cities. Mobility is arguably the most consequential domain of AI-driven change, because cities are defined—physically, temporally, and economically—by their transport infrastructures [57,58]. The mediaeval city was the city of walking and the cart; the modern city is the city of the automobile. Each transport technology has not merely served the city but shaped it: determining its spatial extent, its density patterns, its connectivity, and its social geography. The convergence of electric propulsion and autonomous driving, mediated by AI, promises a transformation of comparable magnitude—one that will redefine the cost, availability, and spatial logic of urban mobility [57,59].

4.1. How Mobility Defines the City: Space, Time, and Cost

Three parameters have historically determined the form and functioning of cities through their transport systems [57]. The first is *space*: urban physiognomy and connectivity are dictated by the dominant mobility technology of each era. Railway lines created linear urban corridors; the automobile enabled low-density suburban sprawl; metro systems shaped radial density patterns around stations. The second is *time*: urban life is organised around an implicit constraint of approximately one hour of commuting. A city grows to the extent that its inhabitants can traverse it within this limit, and the effective size of the city is thus a function of the speed and reliability of its transport. The third is *cost and availability*: access to transport at different times, in different zones, and at different price points determines who can participate in urban economic and social life—and who is excluded.

For over a century, the cost per kilometre of private automobile transport has remained remarkably stable at approximately €0.40 in constant prices [57]. Electric propulsion disrupts this equilibrium by dramatically reducing energy and maintenance costs. Autonomous driving disrupts it further by eliminating the driver—the single largest cost

component of any for-hire transport service, accounting for approximately 60% of total ride-hailing costs [60,61]. The combination of these two disruptions opens the possibility of individual on-demand transport at costs approaching €0.10 per kilometre, and substantially less for shared or collective autonomous services [57,62]. This is not a marginal cost reduction; it is a structural shift that could fundamentally alter the economics of urban mobility.

4.2. *The Autonomous Mobility Revolution: From Pilot to Scale*

What was a speculative prospect five years ago is now an operational reality in a growing number of cities. Waymo, Alphabet's autonomous driving unit, currently provides over 500,000 paid robotaxi rides per week across ten U.S. cities—up from 175,000 at the start of 2025, representing a 157% increase in twelve months [60]. The company operates a fleet of over 3000 robotaxis and is targeting one million weekly rides by the end of 2026. Its geographic footprint has expanded from initial deployments in Phoenix, San Francisco, and Los Angeles to include Austin, Atlanta, Miami, Dallas, Houston, San Antonio, and Orlando, with planned launches in Nashville, Las Vegas, San Diego, Detroit, Washington D.C., Seattle, and Denver [60,61]. Critically, Waymo has announced plans to begin operations in London in 2026 and has deployed test vehicles in Tokyo—marking the first moves toward international expansion beyond the United States [61].

In China, the scale of deployment is equally striking. Baidu's Apollo Go service operates over 1000 robotaxis across 22 cities, including Beijing, Shanghai, Wuhan, and Shenzhen, completing over 17 million cumulative orders [62]. In Wuhan, where the largest single-city deployment operates, Apollo Go has achieved per-vehicle profitability—a milestone that demonstrates commercial viability at scale [62]. Weekly ride volumes reached 250,000 by late 2025, comparable to Waymo's volumes at that time [62]. Moreover, Baidu has announced agreements with Lyft to expand into Europe, with vehicles manufactured by Jiangling Motors expected to operate in Germany and Britain starting in 2026 [62]. Chinese competitors including Pony.ai, WeRide, and AutoX are pursuing parallel international strategies, with deployments or regulatory approvals in the UAE, Singapore, and South Korea [63].

The revolution extends well beyond robotaxis. **Autonomous buses** are entering regular public transit operations in Chinese cities. In August 2025, WeRide and Shenzhen Bus Group launched Shenzhen's first Level 4 fully driverless robobus line—the B888 route in Luohu District—connecting Luohu Port to MixC Market over a 6.6 km route, with vehicles equipped with over 20 sensors providing 360-degree perception [64]. This is not an isolated pilot: Beijing, Xiong'an New Area, Guangzhou, Changsha, Wuxi, Zhengzhou, Chongqing, and Hainan have all introduced autonomous shuttles on public roads, with technology provided by WeRide, Baidu Apollo, QCraft, and UISEE [64,65]. WeRide is partnering with bus manufacturer Yutong for global rollout, signalling that autonomous public transit—not just individual ride-hailing—is entering its scaling phase.

Beyond hardware scale-up, recent work demonstrates that generative AI is also being used to dynamically prioritise the information streams that autonomous vehicles must process in complex urban scenarios, weighting environmental factors so that critical signals receive computational attention before non-critical ones [66]. This is a non-trivial enabler of safe operation in dense city environments and represents one of the explicit links between the agentic AI literature and the autonomous mobility vector examined here.

Autonomous delivery vehicles represent the third dimension of the mobility transformation, and China is again the leading laboratory. By mid-2025, over 15,000 autonomous delivery vehicles (ADVs)—mid-sized electric vans with two to ten cubic metres of storage handling payloads of up to a ton—were operating on roads nationwide, with approvals in

over 200 Chinese cities [67,68]. Platform giants Meituan, JD.com, and Cainiao (Alibaba) have embedded driverless vehicles within their fulfilment systems, using them to replace human couriers on predictable last-mile routes. In Beijing's Shunyi district, Meituan operates a hybrid model where autonomous vehicles transport parcels to transfer stations and couriers complete the final hundred metres [67]. Manufacturer Neolix, which holds China's first official permit for autonomous delivery on public roads, has received approximately 30,000 global orders and now produces over 1000 vehicles per month [68]. Two manufacturers command close to 90% of the professional autonomous delivery market, indicating rapid consolidation and industrial maturity [68].

The significance of autonomous delivery for cities is substantial: it reduces the number of individual human-driven delivery trips (a major source of urban congestion and emissions), enables consolidated multi-drop routing, and extends reliable logistics to areas and hours currently underserved. As the "Ciutats i IA" essay observes, the mobility transformation encompasses not only the movement of people but the movement of goods—and both dimensions are being reshaped simultaneously by the same convergence of electric propulsion and autonomous navigation [57].

The technology is thus no longer confined to controlled test environments. Across robotaxis, autonomous buses, and delivery vehicles, it is scaling commercially, expanding geographically, and entering the phase of mass deployment that will trigger the urban transformations discussed below.

4.3. Urban Implications: Availability, Equity, and Spatial Restructuring

The most radical impact of autonomous electric mobility may not be cost reduction but *availability transformation* [57]. Conventional public transport—buses, metro, taxis—concentrates service in central zones, high-density corridors, and peak hours. Peripheral areas, low-density neighbourhoods, and off-peak hours are systematically underserved. This creates a geography of mobility inequality: residents of well-connected central neighbourhoods enjoy frequent, affordable transport, while those in peripheral areas face long waits, limited routes, and high costs.

Autonomous vehicles are indifferent to time of day, demand density, and geographic peripherality. A robotaxi responds at 3 a.m. as readily as at 8 a.m.; it serves a suburban residential street as willingly as a downtown boulevard [57]. On-demand autonomous transit—minibuses and shuttles that dynamically adjust routes based on real-time demand—can provide frequent, flexible service to areas where fixed-route transit is economically unviable [69]. The result is a potential connectivity revolution: living far from the city centre need no longer mean being disconnected from urban opportunities. In low-traffic conditions, peripheral locations could be within ten minutes of multiple urban destinations [57].

The implications for urban spatial structure are profound. If high-quality, low-cost transport becomes universally available regardless of location and time, the pressure on central housing markets could ease, as previously peripheral zones become viable residential alternatives. The chronic housing affordability crisis afflicting most major cities is, in significant part, a mobility crisis: housing is unaffordable in central locations because those are the only locations with adequate transport [57,70]. Autonomous electric mobility does not solve the housing problem directly, but it fundamentally changes the spatial equation by making connectivity less dependent on proximity to fixed infrastructure.

For goods movement, the impact is analogous. Autonomous delivery vehicles—small vans and pods operating on optimised multi-drop routes—can consolidate shipments, reduce the number of individual delivery trips, and provide service at lower cost and with greater availability than human-driven alternatives [57,71]. The proliferation of

micromobility services (shared bicycles, scooters) is already reshaping last-mile logistics; autonomous systems will extend this logic to heavier loads and longer distances.

4.4. Two Policy Scenarios: The Decisive Role of Municipal Government

The realisation of these benefits is not automatic. It depends critically on the policy choices made by municipal governments—a point that connects directly to the extended mirroring hypothesis and the argument developed in Section 2 [57].

Scenario A: Proactive municipal regulation. Cities that proactively licence autonomous mobility services—robotaxis, on-demand transit, autonomous logistics—while managing the transition for incumbent sectors (taxi drivers, delivery workers, conventional transit operators) can steer the technology toward a city with far fewer private cars and a rich diversity of autonomous services, public and private, for both passengers and freight. In this scenario, private car ownership declines because the alternative—on-demand, low-cost, always-available autonomous transport—is simply more convenient and more affordable. Street space currently devoted to parking can be reclaimed for public use. Traffic volumes may decrease as shared autonomous vehicles replace single-occupant private cars [57,69].

Scenario B: Regulatory paralysis and incumbent protection. Cities that refuse to licence autonomous services—whether to protect existing taxi collectives, delivery operators, or conventional transit—risk producing the opposite outcome. In this scenario, private citizens acquire their own autonomous vehicles, which, instead of parking, circulate empty or reposition to remote lots, returning when summoned. The result is *more* vehicles in circulation, *more* congestion, and a city environment that is worse, not better, than the status quo [57]. This is a powerful example of how regulatory choices designed to protect a specific group can end up harming the broader urban community—and it illustrates the mirroring hypothesis at the municipal level: a regulatory architecture designed for human-driven transport, applied unchanged to autonomous technology, produces perverse outcomes because it fails to mirror the new technology's possibilities.

The cities at the forefront—San Francisco, Austin, Wuhan, Shenzhen—have chosen some variant of Scenario A, establishing regulatory frameworks that accommodate autonomous vehicles while imposing safety and data-sharing requirements [60,62,69]. European cities remain largely in a pre-decision phase, though London's openness to Waymo and Germany's early regulatory frameworks for autonomous driving on public roads signal movement [61,71]. The decisive factor is municipal agency: national legislation provides the legal framework, but it is cities that issue operating permits, designate service zones, manage road infrastructure, and set the terms of engagement with autonomous mobility providers.

4.5. The International Landscape: A Widening Gap

The current geography of autonomous mobility deployment reveals a rapidly widening gap between pioneering and lagging cities. The leading cities are overwhelmingly in the United States (San Francisco, Phoenix, Los Angeles, Austin) and China (Wuhan, Shenzhen, Beijing, Shanghai), with emerging operations in the UAE (Dubai, Abu Dhabi) and early-stage preparations in Europe and Japan [60–63]. This concentration is not coincidental: it reflects the interplay of regulatory openness, technology ecosystem density, and institutional willingness to experiment that the conceptual framework in Section 2 identifies as the conditions for innovation clustering.

European cities, with notable exceptions, risk falling behind. Barcelona's experience is illustrative: an autonomous bus pilot was launched and then discontinued, and today the city has no operational autonomous mobility services [57]. As the "Ciutats i IA" analysis

argues, if Barcelona—or any European city—wants to see how autonomous buses and robotaxis actually work at scale, it must currently look to China or California [57]. The gap is not merely technological; it is institutional and cultural, reflecting a regulatory and political environment that has been slower to accommodate autonomous mobility than its American and Asian counterparts. This gap, as the cumulative recursive hybridisation framework predicts, will compound over time: cities with operational deployments develop data, expertise, regulatory know-how, and public acceptance that make subsequent iterations easier, while cities without them fall further behind with each passing year.

5. Robotics and Intelligent Urban Infrastructure

The third vector of the CRH triad—urban robotics and intelligent infrastructure—is the most physically embedded: it concerns the material fabric of the city itself. As the framework predicts (P3), the infrastructure loop connects all three vectors, because the sensors, networks, and platforms that serve autonomous vehicles also serve cleaning robots, inspection drones, and city brain systems.

The third vector of urban transformation concerns the physical fabric of the city itself: how streets are cleaned, infrastructure is maintained, incidents are detected, and emergency responses are coordinated. While less visible to the general public than chatbots or robotaxis, the deployment of intelligent systems and robots in the management of urban space represents a profound shift—from reactive, labour-intensive maintenance to proactive, data-driven, and increasingly autonomous operation of the city as a physical system [57,72]. In the framework of the extended mirroring hypothesis, urban robotics and intelligent infrastructure constitute the domain where AI agents most directly take on roles previously reserved for human workers, fundamentally changing the coordination architecture of city operations.

5.1. City Brain Systems: From Traffic Optimisation to Urban Intelligence Platforms

The most mature manifestation of AI in urban physical management is the “city brain”—an integrated platform that ingests data from thousands of sensors, cameras, and connected systems to monitor, analyse, and act upon urban conditions in real time. The concept was pioneered by Alibaba’s ET City Brain, launched in Hangzhou in 2016, and has since been deployed in dozens of Chinese and Asian cities [19,20].

The results in Hangzhou are well documented: incident-detection accuracy exceeded 92%, average driving speeds increased by approximately 15%, daily commutes shortened by three minutes, and emergency response teams reached destinations seven minutes faster [19,20]. The city, once ranked fifth among China’s most congested, dropped to 57th on the national congestion index [19]. The system operates by integrating data from traffic lights, surveillance cameras, GPS signals from vehicles, and mobile phone location data, using AI to optimise signal timing, detect incidents, reroute traffic, and allocate emergency resources dynamically. City Brain has since expanded to Guangzhou (emergency service optimisation), Suzhou (accident detection), and numerous other cities across China, as well as implementations in Kuala Lumpur and Macau [20,73].

What distinguishes city brain systems from conventional smart-city dashboards is their capacity for *autonomous action*, not merely monitoring. The system does not simply alert a human operator to a traffic jam; it adjusts signal timing to alleviate it. It does not merely display an accident report; it identifies the optimal ambulance route and dispatches resources. In mirroring hypothesis terms, city brain platforms represent an organisational architecture in which AI agents have been assigned operational decision rights within defined parameters—the system acts on behalf of the city within governed mandates, just as the agentic chatbots described in Section 3 act on behalf of citizens.

The evolution of city brain systems points toward comprehensive urban intelligence platforms that integrate not only traffic data but environmental monitoring (air quality, noise, temperature), infrastructure condition assessment (road surface, utility networks, building facades), public safety, and energy management. China's smart city AI is increasingly moving into environmental control, with systems that monitor and respond to pollution events, manage urban heat islands, and optimise energy distribution across city districts [74]. This integration creates the data substrate on which the other two vectors depend: autonomous vehicles require real-time traffic and road-condition data; agentic administration benefits from situational awareness of urban conditions for proactive service delivery.

5.2. Urban Service Robots: Street Cleaning, Maintenance, and Beyond

Below the scale of city-wide platforms, a growing fleet of specialised robots is taking on tasks that have traditionally required large numbers of human workers. The deployment is most advanced in Chinese cities, where fiscal constraints—municipal budgets are lean, and the imperative to maintain urban quality with limited resources is acute—have driven rapid adoption [57,72].

In Shenzhen's Shijing sub-district, 36 autonomous cleaning robots developed by Cowa Robot patrol an area of approximately 2.7 million square metres, sweeping streets, collecting waste, and operating continuously across day and night shifts [75]. Guangzhou has announced plans to increase its fleet of unmanned sanitation equipment to 1000 units by 2026, and Hangzhou now explicitly requires the inclusion of unmanned equipment in new public sanitation tenders [75,76]. These are not experimental pilots—they represent institutionalised procurement decisions that embed autonomous robots into the routine operational architecture of city maintenance.

The logic extends beyond cleaning. Autonomous inspection robots monitor infrastructure conditions—bridges, tunnels, utility networks—detecting cracks, corrosion, and structural anomalies that human inspectors might miss or reach only at considerable cost and risk [72]. Robotic systems for vegetation management, road surface repair, and facade inspection are in various stages of deployment or advanced testing across Chinese, Japanese, and Korean cities [77].

5.3. Drones, Emergency Automation, and Urban Patrol

Drones add an aerial dimension to urban robotics. Their applications in cities span logistics (Walmart has completed over 20,000 drone deliveries across U.S. hubs and announced plans to expand coverage to 1.8 million additional households [78]), emergency response (drone-carried defibrillators, search-and-rescue in disaster scenarios), infrastructure inspection (power lines, rooftops, facades), and environmental monitoring (air quality sampling, flood mapping).

In the domain of urban security and patrol, Chinese cities have moved furthest. Chengdu deployed teams of robot police officers in June 2025, combining quadruped robots, wheeled robots, and humanoid robots to patrol city streets [79]. Hangzhou placed AI-powered traffic policing robots on active duty in December 2025, and in Wuhu, a humanoid officer designated Intelligent Police Unit R001 oversees a busy junction, using cameras, speakers, and an AI system to detect cyclists and pedestrians in the wrong lane [79,80]. In Shenzhen, EngineAI's PM01 humanoid robots—standing 1.38 m tall—patrol alongside human officers [79]. These deployments, while still limited in scale, signal a trajectory in which robotic agents share public space with citizens as visible components of urban governance infrastructure.

Shenzhen is being designed as China's first "robot-friendly" urban district, where robots will transition from closed training environments to open, on-street operation in neighbourhood blocks [81]. This is a deliberate urban planning decision—the city is reconfiguring its physical and regulatory infrastructure to accommodate robotic agents as permanent inhabitants of public space, not merely as temporary experimental devices.

5.4. *The Integration Challenge: From Isolated Robots to Systemic Urban Intelligence*

The current state of urban robotics is characterised by fragmentation: cleaning robots, patrol robots, delivery drones, and city brain platforms typically operate as independent systems, managed by different agencies or contractors, with limited data sharing or coordination. The transformative potential lies in integration—and this is where the cumulative recursive hybridisation framework developed in Section 2 becomes directly operative.

Consider a fully integrated scenario: city brain platforms ingest data from autonomous vehicles, delivery robots, cleaning machines, patrol drones, and infrastructure sensors to build a real-time model of urban conditions. Cleaning robots are dispatched to areas identified as high-priority based on pedestrian flow data from the mobility network. Infrastructure maintenance robots are routed to locations flagged by the sensors embedded in autonomous vehicles' road-surface detection systems. Emergency drones are pre-positioned based on predictive models fed by the city brain's incident-detection algorithms. Patrol robots share security data with traffic management systems to coordinate responses to accidents. Each system feeds data to the others; each becomes more effective because the others exist.

This integrated scenario does not yet exist anywhere in full. But the components are operational, and the cities that are building them—Shenzhen, Hangzhou, Guangzhou, Seoul—are creating the conditions for integration to emerge. The critical enabler is not any single technology but the *ecosystem density* that permits cross-system data flows, shared standards, and coordinated governance. This is, once again, a function of municipal agency: it is the city government that sets data-sharing protocols, defines interoperability standards, manages procurement to ensure compatibility, and creates the institutional architecture within which autonomous systems from different vendors and domains can interact.

For European and American cities, where urban robotics deployments remain more limited and more fragmented, the risk is not merely that they will lack individual robotic capabilities but that they will miss the systemic integration effects that arise when multiple autonomous systems operate within the same urban ecosystem. A cleaning robot in isolation is a labour-saving device. A cleaning robot connected to a city brain platform, sharing data with autonomous vehicles and coordinated with infrastructure inspection drones, is a node in an intelligent urban system—and the value of the node is a function of the network it belongs to.

6. The City as Locus of AI-Driven Transformation

The preceding sections have examined three vectors of AI-driven urban transformation—agentic governance, autonomous mobility, and urban robotics—each with its own maturity trajectory, international landscape, and institutional prerequisites. This section brings the argument together. Drawing on the conceptual framework developed in Section 2 and the empirical evidence assembled in Sections 3–5, we argue that the city is not merely the setting where these transformations unfold but the *generative locus* where they interact, recombine, and compound. The key to understanding which cities will lead the agentic AI era lies not in any single vector but in the dynamics of cumulative recursive hybridisation across all three—and in the local conditions that enable or inhibit that hybridisation.

6.1. Cumulative Recursive Hybridisation in Practice: How the Three Vectors Interact

The Industrial Revolution did not occur because steam engines improved, or because spinning jennies were invented, or because coke-smelted iron became available. It occurred because these innovations *interacted* within specific geographic clusters—Lancashire’s cotton towns, Birmingham’s metal trades, the Scottish Lowlands’ engineering workshops—each improvement in one domain enabling and demanding improvements in the others, in recursive cycles that compounded over decades [12,13,82]. The process was not planned; it emerged from the density, proximity, and diversity of the local ecosystem.

The same dynamic is beginning to emerge in the cities at the forefront of the AI-driven transformation, although the process is still in its early stages and the full integration remains aspirational. The interactions among the three vectors can be mapped through four feedback loops, as theorised in Section 2.4 and illustrated in Figure 5:



Figure 5. Cumulative recursive hybridisation: four feedback loops connecting three vectors of urban AI transformation. The three vectors—agentive governance, autonomous electric mobility, and urban robotics—interact through data, regulatory, infrastructure, and talent loops within urban ecosystems. Key takeaway: each cycle increases the system’s capacity for the next, generating compounding returns; cities that activate all four loops simultaneously pull ahead of those that deploy in isolation.

Data loops. Autonomous vehicles—robotaxis, buses, delivery vans—are mobile sensor platforms. As they navigate city streets, they continuously collect high-resolution data on road conditions, traffic patterns, pedestrian flows, air quality, and urban infrastructure state. This data feeds city brain platforms, improving their predictive models and operational decisions. In Hangzhou, the city brain’s traffic optimisation depends on the density of data inputs; as autonomous vehicles proliferate, the volume, granularity, and freshness of data increase, making the system more effective [19,20]. Conversely, the city brain’s traffic optimisation and incident-detection capabilities make autonomous vehicle operation safer and more efficient, closing the feedback loop. Delivery robots and cleaning machines contribute additional data streams—pavement conditions, waste accumulation patterns, micro-climate variations—that no single system would collect on its own.

Regulatory loops. Agentic governance systems capable of processing permits, licences, and regulatory decisions in real time can dramatically accelerate the deployment of autonomous mobility and robotics. Today, a city that requires months of bureaucratic process to issue an operating licence for a robotaxi fleet, or that lacks procedures for permitting cleaning robots on public sidewalks, creates bottlenecks that slow the entire ecosystem. An administration operating at Level 3 or 4 of the maturity model—with API-integrated transactional capabilities and agentic orchestration—could process such decisions in days or hours, adapting regulations dynamically as the technology evolves [10]. Shenzhen’s decision to design itself as a “robot-friendly” urban district [81] is an example of regulatory architecture that mirrors the possibilities of the technology: the city is not waiting for robots to arrive and then figuring out how to regulate them; it is proactively redesigning its institutional framework to accommodate robotic agents as permanent participants in urban life.

Infrastructure loops. Robotic maintenance systems—cleaning robots, inspection drones, road-surface monitoring—improve the physical infrastructure on which autonomous vehicles depend. Potholes, debris, and degraded road markings are among the most common causes of autonomous vehicle disengagement; a city that maintains its infrastructure proactively through robotic systems creates a more reliable operating environment for autonomous mobility [58,65]. In turn, the data generated by autonomous vehicles about infrastructure conditions enables more targeted and efficient robotic maintenance, creating a self-reinforcing cycle of improvement.

Talent loops. The most consequential feedback loop may be the least visible. As argued in Section 2.2, talent and knowledge are constituted through participation in active projects. A city that simultaneously deploys agentic governance systems, autonomous mobility services, and urban robotics creates a dense ecosystem of frontier projects across multiple AI domains. This attracts and develops specialised talent—AI engineers, data scientists, urban planners with robotics expertise, public servants skilled in AI-mediated service design—who would not develop these capabilities in a city without live deployments [15,40,42–44]. The talent, in turn, enables more ambitious projects, attracts investment, and deepens the ecosystem. San Francisco’s concentration of autonomous vehicle expertise is inseparable from the fact that Waymo, Cruise (now wound down), Zoox, and numerous startups operated there simultaneously, creating a labour market and knowledge commons that no single company could have generated alone [60,61]. Shenzhen’s emergence as a robotics capital reflects the same dynamic: the co-presence of DJI (drones), UBTECH (humanoid robots), Cowa (cleaning robots), WeRide (autonomous vehicles), and dozens of smaller firms creates an ecosystem in which engineers move between companies, ideas cross-pollinate, and the city’s collective capability compounds [75,79,81].

6.2. *Why Cities, Not Nations: The Local Embedding of Transformation*

A recurrent finding across all three vectors is that advanced deployments are concentrated in a small number of cities, not distributed evenly across national territories. Estonia’s Bürokratt is a national platform, but it is an exception; most governance chatbot innovation happens at the municipal level (Buenos Aires, Madrid, Dubai). Autonomous mobility is concentrated in specific cities (San Francisco, Wuhan, Shenzhen), not deployed uniformly across the United States or China. Urban robotics follows the same pattern (Shenzhen, Hangzhou, Guangzhou).

This concentration is not accidental. It reflects the fundamentally local nature of the conditions required for AI-driven urban transformation. National governments can provide legal frameworks, fund research, and set standards—but they cannot replicate the ecosystem density that makes transformation possible. The critical ingredients are:

Regulatory authority at the municipal level. It is the city that issues operating permits for robotaxis, sets zoning rules for robot-friendly districts, procures autonomous cleaning services, and decides whether its administration will deploy agentic chatbots. National legislation enables; municipal government acts [57,69].

Institutional proximity and fast iteration cycles. In a city like Shenzhen, the robotics startup developing a cleaning robot operates within kilometres of the municipal sanitation authority that will procure it, the university lab that tests its sensors, the autonomous vehicle company whose road-condition data could optimise its routes, and the city brain platform that could coordinate its operations. This proximity compresses the cycle from idea to pilot to deployment to iteration in ways that geographically dispersed arrangements cannot match [37].

Embedded knowledge and communities of practice. As the theoretical framework in Section 2.2 established, expertise develops through participation in live projects. A city with active deployments across multiple AI domains develops communities of practice—networks of practitioners who share tacit knowledge, solve problems collaboratively, and build the institutional memory that makes subsequent deployments more effective [41,42]. This knowledge is embedded in the city’s social fabric: it travels when people change jobs within the same metropolitan area, but it does not easily transfer to distant cities that lack the project base in which to apply it.

The Medici Effect at urban scale. The intersectional innovations that arise when governance, mobility, and robotics professionals interact in the same urban ecosystem—the chance conversation between a chatbot developer and a robotaxi engineer that sparks an idea for real-time permit processing; the urban planner who realises that cleaning robot data could inform housing policy—occur because diverse domains are co-located in dense proximity [45,46]. These innovations cannot be planned or mandated by national policy; they emerge from the ecosystem’s structure.

This analysis explains why the relevant unit of comparison for AI-driven urban transformation is not “the United States vs. China” or “Europe vs. Asia” but “San Francisco vs. Barcelona,” “Shenzhen vs. Berlin,” “Wuhan vs. Madrid.” The transformation is happening city by city, and the differences between leading and lagging cities within the same country (San Francisco vs. Detroit, Shenzhen vs. Chengdu) can be as large as the differences between countries.

6.3. From Experimentation to Systemic Transformation: What Distinguishes Leading Cities

Not every city that experiments achieves systemic transformation. Many cities have launched isolated pilots—an autonomous bus here, a chatbot there—without progressing to the integrated, cross-domain deployments that generate compounding returns. What distinguishes the cities that are advancing from those that are stalling?

The evidence from Sections 3–5 points to four differentiating factors:

Simultaneity across vectors. Cities that deploy across multiple vectors simultaneously—governance, mobility, and infrastructure—create the conditions for cross-fertilisation. Shenzhen is deploying autonomous buses (WeRide), cleaning robots (Cowa), patrol robots (EngineAI), and building a robot-friendly urban district, all while its administration modernises digital services. This simultaneity is not a coincidence; it reflects a municipal strategy of comprehensive AI adoption that creates the ecosystem density required for recursive hybridisation. Cities that pursue one vector in isolation—a chatbot initiative here, an autonomous bus pilot there—miss the compounding effects.

Institutional willingness to redesign, not just adopt. The extended mirroring hypothesis predicts that value comes not from deploying technology onto existing structures but from co-evolving institutional architectures with technological capabilities. The cities that are

pulling ahead—Estonia (national administration redesigned around Bürokratt), Shenzhen (urban districts redesigned for robotic agents), Wuhan (regulatory framework redesigned for autonomous vehicles)—have accepted that the institutional architecture itself must change. Cities that deploy AI while preserving inherited bureaucratic structures, incumbent protections, and legacy procurement processes remain in the asymmetry trap.

Tolerance for iteration and failure. The “Ciutats i IA” essay makes the point forcefully: there are no established best practices to copy in AI-driven urban transformation. The field is exploratory, and the path forward requires experimentation, evaluation, and continuous adjustment [57]. Cities that accept failure as a component of learning—that launch pilots, evaluate results, iterate, and scale what works—develop adaptive capacity. Cities that demand guaranteed outcomes before acting, or that cancel pilots at the first sign of difficulty (as Barcelona did with its autonomous bus), erode their own capacity to learn and improve.

Public-sector entrepreneurship. In every leading case, municipal government has played an active, entrepreneurial role—not merely as regulator or procurer but as ecosystem orchestrator. Hangzhou’s city government co-developed City Brain with Alibaba. Shenzhen’s government is designing robot-friendly districts. Buenos Aires’ government deployed Boti on WhatsApp as a deliberate channel strategy. These are acts of institutional entrepreneurship: public officials taking initiative, accepting risk, and shaping the trajectory of technological adoption in their cities. The role of the public sector is not to step aside and let technology companies operate freely, nor to regulate restrictively and slow adoption, but to actively co-create the conditions for transformation—setting standards, facilitating integration, managing transitions for affected workers and communities, and ensuring that the benefits are broadly shared.

6.4. *The Divergence Ahead: Path Dependency and the Widening Gap*

The cumulative recursive hybridisation framework suggests that the gap between leading and lagging cities is likely to widen over time rather than narrow automatically. This expectation follows from the logic of path dependency and increasing returns that characterises innovation ecosystems [16,49]. But the implication goes further than a widening performance gap. What is emerging is a new urban order in which the leading cities function as innovation hubs—concentrating talent, accelerating knowledge circulation, enabling cross-fertilisation and hybridisation across the three vectors, and generating novel institutional forms and business models that latecomers will eventually import rather than invent. The analogy with the Industrial Revolution is again precise: Manchester did not merely produce more textiles than other cities; it became the organising centre of an entirely new mode of production, and its institutional innovations—the factory system, labour markets, urban infrastructure—became templates that shaped the world. The cities now leading in agentic governance, autonomous mobility, and urban robotics are on a similar trajectory: they are not just deploying technologies but constructing the ecosystems in which the next phase of urban civilisation takes shape.

Cities that have already deployed across multiple vectors are accumulating advantages that are difficult to replicate: operational data that improves AI systems, regulatory know-how that accelerates subsequent deployments, talent pools that deepen with each new project, public acceptance that builds through positive experience, and institutional memory that reduces the cost and risk of future innovation. Each cycle of deployment, evaluation, and iteration makes the next cycle easier and more productive. The result is a compounding trajectory in which early movers accelerate while latecomers face not only the challenge of catching up technologically but the far more daunting challenge of building the institutional, human, and relational capital that the leading cities have accumulated over years of active deployment.

For cities that have not yet moved decisively, the window of opportunity is narrowing. The historical analogy is instructive: during the Industrial Revolution, the cities that industrialised first—Manchester, Birmingham, Glasgow—maintained their advantages for over a century, shaping national and global economic geography for generations [12,13]. Cities that missed the industrial transition—or that actively resisted it to protect pre-industrial interests—were relegated to peripheral status. The same dynamic is plausible for the agentic AI era: the cities that build the institutional architectures, talent ecosystems, and cross-domain integration capacities in this decade may establish advantages that persist for decades to come.

This does not mean that latecomers cannot succeed. But it does mean that the cost of delay is not linear—it is exponential, because each year of inaction represents not just lost time but lost learning, lost talent formation, and lost ecosystem development. Cities that aspire to participate in the AI-driven transformation must act now—not with isolated pilots or cautious studies, but with the kind of comprehensive, simultaneous, institutionally transformative commitment that the leading cities have already made. The question, as the “Ciutats i IA” essay poses it, is whether a city wants to be one that tries, leads, and shapes its own future—or one that watches from the sidelines as others build the future and then adopts, on others’ terms, the models designed elsewhere [57].

6.5. Counterfactual Evidence: Failed, Stalled, and Non-Western Trajectories

The empirical sections above privileged leading-edge deployments in the United States, China, and a small set of European cities. Both reviewers correctly observed that this case selection risks survivorship bias: a framework that explains success through the presence of certain conditions must also account for failure when those conditions are absent or when the deployments themselves collapse for reasons the framework does not capture. This subsection examines four cases that complicate the narrative—two outright failures, one Global South trajectory, and one regulatory rollback—and uses each to test the boundary conditions of cumulative recursive hybridisation.

Toronto: Sidewalk Labs and the limits of corporate-led smart urbanism. Between 2017 and 2020, Alphabet’s Sidewalk Labs subsidiary attempted to develop a 12-acre lakeside district in Toronto as a flagship demonstration of integrated smart-city technologies—sensor-rich streets, autonomous mobility, modular construction, and data-driven governance. The project collapsed in May 2020, ostensibly because of pandemic-related economic uncertainty but in reality because of sustained civic resistance to its data-governance model and the inability of the corporate developer to obtain a social licence for the data flows on which the integrated vision depended [83]. The Toronto failure is informative for our framework: the technical components were available, but the regulatory and trust loops (P3) failed, and consequently the data and talent loops never closed. CRH does not predict success wherever technology is available—it predicts success only when all four loops can close, and Toronto demonstrates that the regulatory loop can be the binding constraint even in a wealthy city with deep talent pools.

San Francisco and the Cruise rollback. In October 2023, the California Public Utilities Commission suspended Cruise’s driverless operating permit in San Francisco following a serious pedestrian incident, and General Motors subsequently announced in December 2024 that it would wind down the Cruise robotaxi business entirely. This is the most significant regulatory reversal in autonomous mobility to date, and it occurred in the city our framework identifies as the global leader for the vector. The lesson for CRH is not that the framework is wrong but that recursion is bidirectional in a stronger sense than we initially emphasised: as data and operational evidence accumulate, public-acceptance and regulatory loops can run in reverse, contracting the action space rather than expanding it.

A revised statement of P2 should therefore acknowledge that super-additive composition is conditional on safety performance remaining within the tolerance band of public and regulatory actors—an empirical condition, not an assumption.

Songdo and the limits of master-planned smart cities. South Korea’s Songdo International Business District, launched in the mid-2000s as a purpose-built ubiquitous-computing city, was designed from scratch with the kind of integrated digital substrate that CRH would predict as ideal soil. Two decades later, Songdo has under-performed on almost every dimension: population growth has lagged plans, the integrated services have not generated the expected innovation density, and the city remains largely a residential dormitory rather than the global innovation hub it was branded to be [84]. The lesson is that ecosystem density (Section 2.2) cannot be manufactured by infrastructure alone; the embedded knowledge, communities of practice, and project-based learning that the talent loop (P4) requires need pre-existing density to bootstrap, and a master-planned greenfield typically lacks them.

Bengaluru: a Global South trajectory. The leading European, North American, and East Asian cases have institutional, fiscal, and infrastructural endowments that are unrepresentative of the cities where most of the world’s urban population lives. Bengaluru offers a useful contrast. India’s information-technology capital has built a deep talent loop in software and an emerging one in robotics and mobility, hosting both Indian unicorns and the R&D centres of global firms. Yet its mobility, governance, and physical infrastructure loops remain weakly closed: the chatbot Jugalbandi referenced in Section 3 operates at Level 1; autonomous mobility deployment is constrained by mixed traffic conditions that defeat current sensing stacks; and city governance remains fragmented across multiple agencies with overlapping and sometimes competing mandates [85]. Bengaluru therefore exhibits one strong loop and three weak loops, and our framework suggests that it should under-capture the hybridisation premium relative to Shenzhen or San Francisco despite a comparable concentration of AI talent. This prediction is consistent with current observation and is empirically testable as Bengaluru’s capability stack matures. More broadly, the dynamics of African, South Asian, and Southeast Asian cities—where institutional contexts, infrastructure baselines, and demographic pressures differ markedly—deserve dedicated investigation that the present paper can only initiate.

Taken together, these counterfactuals do not falsify CRH; they sharpen it. They show that the framework’s predictions about which cities will lead are conditional on the simultaneous closure of all four loops, that any single binding constraint—social licence, safety performance, the absence of pre-existing talent density—can interrupt the compounding dynamic, and that the conditions for cumulative recursive hybridisation are exacting rather than abundant.

To standardise the comparative logic, Table 2 codes a small set of illustrative city trajectories on common ordinal dimensions drawn from Sections 3–6. The coding is heuristic rather than metric, but it makes the cross-case comparison more explicit and more easily contestable by future empirical work.

Table 2. Illustrative ordinal comparison of selected city trajectories across the three vectors.

City	Gov. AI	Mobility	Robotics	Reg. Throughput	Talent/Trajectory
Shenzhen	L2–3	High	High	High	High talent density; multi-vector reinforcing trajectory.
Wuhan	L1–2	High	Med.	High	Mobility-led scaling with growing ecosystem depth.

Table 2. Cont.

City	Gov. AI	Mobility	Robotics	Reg. Throughput	Talent/Trajectory
San Francisco	L1–2	High	Med.	High/volatile	High talent density; strong experimentation but regulatory reversals matter.
Buenos Aires	L3	Low	Low	Med.	Governance-led path with weaker cross-vector closure.
Barcelona	L1–2	Low	Low	Low/med.	Strong urban policy capacity, but stop–start deployment across vectors.
Toronto	L1–2	Low	Low	Med.	High talent base, but trust and governance failures stalled closure of loops.
Songdo	L1	Low	Med.	Med.	Infrastructure-rich but ecosystem-thin; weak endogenous talent formation.
Bengaluru	L1	Low	Low/med.	Low/med.	High software talent, but limited loop closure outside the talent dimension.

6.6. Political Economy, Equity, and the Limits of Technological Optimism

The cumulative recursive hybridisation framework as stated in Section 2.4 is necessary but not sufficient: it identifies the conditions under which compounding returns become possible, but it is silent on the political-economic conflicts that determine whether those conditions actually obtain. Five issues deserve sustained attention.

Power asymmetries and labour displacement. The autonomous mobility, robotics, and agentic-administration vectors all reduce demand for specific categories of labour—taxi and delivery drivers, sanitation workers, mid-level administrative staff, contact-centre agents. These groups are typically organised, often unionised, and politically active at the municipal level. Cities that pursue rapid deployment without negotiated transitions risk both ethical failure (the people who bear the costs of transformation are not the people who capture its benefits) and political failure (resistance that delays or reverses deployment, as the partial reversal of robotaxi expansion in some U.S. cities illustrates). A serious agenda for AI-driven urban transformation must include retraining programmes, transition guarantees, and meaningful participation by affected workers in deployment decisions—none of which is automatic and all of which require political will that incumbent administrations may lack.

Electoral cycles and regulatory capture. Municipal politics operates on short electoral cycles that are poorly aligned with the multi-year horizons over which cumulative hybridisation unfolds. A new administration can cancel pilots launched by its predecessor, as Barcelona’s autonomous bus discontinuation illustrates [57,86]. Conversely, regulatory capture by incumbent operators—taxi medallions, traditional sanitation contractors, established consulting firms—can lock cities into legacy structures even when superior alternatives exist. The framework’s prediction that leading cities will widen their lead is therefore not a deterministic forecast but a conditional one: it holds only where political institutions can sustain long-horizon commitments against both incumbent resistance and electoral discontinuity.

The digital divide and the right to a non-agentic interface. Section 3.4’s vision of a ‘windowless administration’ is an idealisation that risks deepening inequality if implemented uncritically. Older citizens, low-income households without reliable connectivity, people with disabilities affecting digital interaction, recent migrants without language fluency, and the digitally marginalised more generally cannot be presumed to interact seamlessly with cognitive agents. A cognitive government that closes its physical service points without

preserving robust non-digital alternatives effectively withdraws basic services from these populations. We therefore qualify the windowless-administration thesis: it should be read as a description of the dominant mode of citizen–state interaction in a mature agentic regime, not as a prescription that physical and human-mediated channels should be eliminated. Equity audits, preserved face-to-face channels, and a ‘right to a human caseworker’ for high-stakes decisions are necessary complements, not optional extras [87,88].

Cybersecurity, privacy, and the centralisation risk. The data loops on which CRH depends create concentrated repositories of behavioural data whose compromise could be catastrophic. Highly integrated city brains constitute single points of failure: a successful attack on the central platform could simultaneously degrade traffic management, autonomous mobility, robotic maintenance, and citizen-facing services. The history of municipal ransomware incidents (Atlanta 2018, Baltimore 2019, dozens of smaller U.S. and European cities since) shows that this is not a hypothetical risk [89]. The framework should therefore be read alongside an explicit security agenda: federated rather than monolithic data architectures, principled minimisation of cross-domain joins, defence in depth, incident-response capacity, and—crucially—an honest assessment of which integrations are worth their security cost [89,90]. Sadowski’s analysis of the political economy of municipal data treats integration as accumulation, with attendant risks of extraction and lock-in that the smart-city literature has historically under-weighted [89].

Accountability gaps and the limits of audit. When AI agents take consequential decisions on citizens’ behalf, the question of who is accountable when things go wrong becomes harder rather than easier. Conventional administrative law assigns accountability to identifiable human officials operating under defined mandates; agentic systems complicate this by interposing learned models whose behaviour is partially opaque even to their designers. Wachter, Mittelstadt, and Floridi’s work on transparent, explainable, and accountable AI for robotics, together with the broader literature on algorithmic accountability, suggests that the appropriate institutional responses include policy-as-code, legible audit trails, contestability mechanisms, and meaningful human oversight at decision points where the cost of error is high [86–88,91]. None of these is technically resolved, and all of them require institutional investment that no city has yet made at scale. Recent work on transparent, explainable and accountable AI for robotic systems underlines that without contestable, auditable interfaces these asymmetries cannot be checked [91].

These five issues are not peripheral caveats; they are the conditions under which the framework is operative. A revised statement of the central thesis is therefore: cities that deploy across the three vectors and successfully navigate the political-economic, equity, security, and accountability challenges enumerated here will benefit from cumulative recursive hybridisation. Cities that deploy without addressing these challenges will either generate transient gains followed by reversal (the Cruise pattern), produce stratified benefits that deepen inequality (the digital divide pattern), or accumulate risks that undermine the data substrate on which the loops depend. The framework is therefore not a prediction that technology will transform cities but a conditional statement about the institutional and political work required for transformation to occur.

7. Conclusions and Future Research Directions

This paper has argued that the agentic AI era is reshaping cities through three interconnected vectors—the transformation of public administration toward cognitive government, the emergence of autonomous electric mobility, and the deployment of robotics and intelligent infrastructure in the urban environment—and that the city, not the nation or the firm, is the natural unit of analysis for understanding this transformation.

Three principal contributions emerge from the analysis.

Taken together, these contributions are theory-building rather than causal-identification claims: they specify mechanisms, boundary conditions, and comparative expectations that can be tested more rigorously in subsequent empirical work.

First, we have extended the mirroring hypothesis from its original domain of firm–product architecture (Conway [21], Colfer and Baldwin [22]) in two directions. The *dynamic extension* holds that organisations and ecosystems do not merely mirror their current structures but explore and converge toward the best strategic configurations that a new technology makes possible. The *ontological extension* holds that agentic AI fundamentally changes the nature of the participating agents: for the first time, the organisational architectures that must achieve symmetry with technology include not only human roles but also AI agents with defined decision rights, authority boundaries, and coordination protocols. The city’s institutional ecosystem—governance structures, regulatory frameworks, talent markets, infrastructure management—must evolve to mirror the possibilities of this hybrid human–AI coordination architecture. Where it does, compounding value follows; where it does not, the city falls into an asymmetry trap.

Second, we have proposed the concept of cumulative recursive hybridisation to explain why the convergence of the three vectors within specific urban ecosystems generates compounding returns analogous to those observed during the Industrial Revolution. The mechanism operates through four feedback loops—data, regulatory, infrastructure, and talent—each of which is recursive (each cycle increases the system’s capacity for the next) and cumulative (gains compound over time). The concept draws on the historiography of industrialisation, the economics of agglomeration, and systems thinking, and it formalises an intuition that runs through much of the smart-city literature without having been made explicit: that the interaction effects among co-located technological vectors are more consequential than the vectors themselves. What CRH adds is the claim that these interaction effects are transforming cities into the new hubs of innovation for the agentic AI era—hubs characterised not merely by technological adoption but by talent concentration, accelerated knowledge circulation, cross-fertilisation across domains, hybridisation of methods and practices, and the continuous emergence of novel proposals. The result is not a more efficient version of the existing city but the construction of a new urban order, and the mirroring hypothesis explains why: cities must redesign their institutional architectures to mirror the possibilities the technology opens, and those that do so become the organising nodes of the new order, just as the industrial cities that mirrored the possibilities of steam and mechanisation became the organising nodes of industrial capitalism.

Third, the comparative analysis across more than twenty governance chatbot deployments, the rapidly scaling autonomous mobility ecosystems of the United States and China, and the emerging urban robotics landscape has identified the conditions that distinguish leading cities from lagging ones: simultaneity of deployment across vectors, institutional willingness to redesign (not merely adopt), tolerance for iteration and failure, and public-sector entrepreneurship as ecosystem orchestration. These conditions are fundamentally local—they depend on municipal regulatory authority, talent ecosystem density, institutional proximity, and the embedded knowledge that accumulates only through active projects.

The implications for urban policy are direct. Cities that wish to participate meaningfully in AI-driven urban transformation are unlikely to benefit by waiting for national directives, relying only on isolated pilots, or protecting incumbent structures at the expense of systemic learning. They would need to coordinate action across governance, mobility, and infrastructure, redesign institutional architectures to incorporate AI agents as legitimate participants, and accept the iterative character of a transformation for which no settled playbook yet exists. The role of municipal government, in this reading, is less passive

regulation than active ecosystem orchestration—co-creating, with firms, universities, and citizens, the conditions under which cumulative recursive hybridisation might take hold.

Several limitations of this study point toward avenues for future research. First, while Section 2.4 has now restated the framework of cumulative recursive hybridisation as five formal propositions (P1–P5), we have not yet operationalised those propositions through system-dynamics simulation, agent-based modelling, or panel econometrics; doing so is the natural next step and we encourage other groups to take it up. Second, the comparative analysis of governance chatbots draws primarily on publicly available documentation and reported capabilities rather than on systematic user-experience data or impact evaluations; longitudinal studies tracking the progression of specific deployments across maturity levels would provide more rigorous evidence of the extended mirroring hypothesis in action. Third, although Section 6.5 introduces counterfactual cases and a Global South trajectory (Bengaluru), the dynamics of AI-driven urban transformation in African, South Asian, and Southeast Asian cities remain under-studied and warrant dedicated empirical investigation. Fourth, Section 6.6 has expanded the political-economy and ethics treatment substantially, but accountability frameworks for hybrid human–AI coordination architectures—and the empirical evaluation of audit, contestability, and equity mechanisms—remain at an early stage and demand sustained interdisciplinary work.

The agentic AI era is not a distant prospect. It is already reshaping the cities that have chosen to engage with it. The central message of this paper is that what is at stake goes beyond the implementation of new technologies to build more efficient cities. What is under way is the construction of a new order—one in which a small number of cities, by deploying simultaneously across the three vectors and redesigning their institutions to mirror the possibilities those technologies open, are emerging as the innovation hubs of the twenty-first century. These hubs concentrate talent, accelerate the circulation of knowledge, enable the cross-fertilisation and hybridisation from which entirely new industries and governance models arise, and generate the compounding returns that widen the gap with every passing year. The parallel with the Industrial Revolution is not rhetorical: it is structural. The cities that led industrialisation—Manchester, Birmingham, the Ruhr—did not merely adopt steam engines; they built a new economic and social order around the recombination of technologies, talents, and institutions. The cities that lead the agentic AI transformation will do the same, and the mirroring hypothesis and cumulative recursive hybridisation are the conceptual tools this paper offers for understanding why, and for identifying which cities are most likely to succeed.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Data available in a publicly accessible repository (Anteverti).

Conflicts of Interest: The author declares no conflict of interest.

References

1. Lipsey, R.G.; Carlaw, K.I.; Bekar, C.T. *Economic Transformations: General Purpose Technologies and Long-Term Economic Growth*; Oxford University Press: Oxford, UK, 2005.
2. McKinsey Global Institute. *The State of AI in Early 2024: Gen AI Adoption Spikes and Starts to Generate Value*; McKinsey & Company: New York, NY, USA, 2024.
3. Bommasani, R.; Hudson, D.A.; Adeli, E.; Altman, R.; Arber, S.; von Arx, S.; Bernstein, M.S.; Bohg, J.; Bosselut, A.; Brunskill, E.; et al. On the Opportunities and Risks of Foundation Models. *arXiv* **2021**, arXiv:2108.07258. [[CrossRef](#)]
4. Noy, S.; Zhang, W. Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence. *Science* **2023**, *381*, 187–192. [[CrossRef](#)]

5. Glaeser, E.L. *Triumph of the City: How Our Greatest Invention Makes Us Richer, Smarter, Greener, Healthier, and Happier*; Penguin Press: New York, NY, USA, 2011.
6. Batty, M. *The New Science of Cities*; MIT Press: Cambridge, MA, USA, 2013.
7. Kitchin, R. *The Data Revolution: Big Data, Open Data, Data Infrastructures and Their Consequences*; SAGE: London, UK, 2014.
8. Dunleavy, P.; Margetts, H.; Bastow, S.; Tinkler, J. *Digital Era Governance: IT Corporations, the State, and E-Government*; Oxford University Press: Oxford, UK, 2006.
9. Mergel, I.; Edelman, N.; Haug, N. Defining Digital Transformation: Results from Expert Interviews. *Gov. Inf. Q.* **2019**, *36*, 101385. [[CrossRef](#)]
10. Almirall, E. Reinventing Public Governance: From Digital Governments to Agentic Governance. *Int. J. Eng. Bus. Manag.* **2026**, *in press*.
11. Wirtz, B.W.; Langer, P.F.; Fenner, C. Artificial Intelligence in the Public Sector—A Research Agenda. *Int. J. Public Adm.* **2021**, *44*, 1103–1128. [[CrossRef](#)]
12. Allen, R.C. *The British Industrial Revolution in Global Perspective*; Cambridge University Press: Cambridge, UK, 2009.
13. Mokyr, J. *The Lever of Riches: Technological Creativity and Economic Progress*; Oxford University Press: Oxford, UK, 1990.
14. Florida, R. *The Rise of the Creative Class: And How It's Transforming Work, Leisure, Community and Everyday Life*; Basic Books: New York, NY, USA, 2002.
15. Moretti, E. *The New Geography of Jobs*; Houghton Mifflin Harcourt: Boston, MA, USA, 2012.
16. Arthur, W.B. *Increasing Returns and Path Dependence in the Economy*; University of Michigan Press: Ann Arbor, MI, USA, 1994.
17. Xi, Z.; Chen, W.; Guo, X.; He, W.; Ding, Y.; Hong, B.; Zhang, M.; Wang, J.; Jin, S.; Zhou, E.; et al. The Rise and Potential of Large Language Model Based Agents: A Survey. *arXiv* **2023**, arXiv:2309.07864. [[CrossRef](#)]
18. Wang, L.; Ma, C.; Feng, X.; Zhang, Z.; Yang, H.; Zhang, J.; Chen, Z.; Tang, J.; Chen, X.; Lin, Y.; et al. A Survey on Large Language Model Based Autonomous Agents. *Front. Comput. Sci.* **2024**, *18*, 186345. [[CrossRef](#)]
19. Zhang, J.; Hua, X.-S.; Huang, J.; Shen, X.; Chen, J.; Zhou, Q.; Fu, Z.; Zhao, Y. City Brain: Practice of Large-Scale Artificial Intelligence in the Real World. *IET Smart Cities* **2019**, *1*, 28–37. [[CrossRef](#)]
20. Pan, Y.; Tian, Y.; Liu, X.; Gu, D.; Hua, G. Urban Big Data and the Development of City Intelligence. *Engineering* **2016**, *2*, 171–178. [[CrossRef](#)]
21. Conway, M.E. How Do Committees Invent? *Datamation* **1968**, *14*, 28–31.
22. Colfer, L.J.; Baldwin, C.Y. The Mirroring Hypothesis: Theory, Evidence, and Exceptions. *Ind. Corp. Change* **2016**, *25*, 709–738. [[CrossRef](#)]
23. Solow, R.M. We'd Better Watch Out. *New York Times Book Review*, 12 July 1987; p. 36.
24. David, P.A. The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox. *Am. Econ. Rev.* **1990**, *80*, 355–361.
25. Milgrom, P.; Roberts, J. Complementarities and Fit: Strategy, Structure, and Organisational Change in Manufacturing. *J. Account. Econ.* **1995**, *19*, 179–208. [[CrossRef](#)]
26. Brynjolfsson, E.; Hitt, L.M. Beyond Computation: Information Technology, Organisational Transformation and Business Performance. *J. Econ. Perspect.* **2000**, *14*, 23–48. [[CrossRef](#)]
27. March, J.G. Exploration and Exploitation in Organisational Learning. *Organ. Sci.* **1991**, *2*, 71–87. [[CrossRef](#)]
28. Nelson, R.R.; Winter, S.G. *An Evolutionary Theory of Economic Change*; Harvard University Press: Cambridge, MA, USA, 1982.
29. Brynjolfsson, E.; McAfee, A. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*; W.W. Norton: New York, NY, USA, 2014.
30. Dafoe, A.; Bachrach, Y.; Hadfield, G.; Horvitz, E.; Larson, K.; Graepel, T. Cooperative AI: Machines Must Learn to Find Common Ground. *Nature* **2021**, *593*, 33–36. [[CrossRef](#)] [[PubMed](#)]
31. Pollitt, C.; Bouckaert, G. *Public Management Reform: A Comparative Analysis—Into the Age of Austerity*, 4th ed.; Oxford University Press: Oxford, UK, 2017.
32. Jacobs, J. *The Economy of Cities*; Vintage Books: New York, NY, USA, 1969.
33. Frenken, K.; Van Oort, F.; Verburg, T. Related Variety, Unrelated Variety and Regional Economic Growth. *Reg. Stud.* **2007**, *41*, 685–697. [[CrossRef](#)]
34. Boschma, R. Proximity and Innovation: A Critical Assessment. *Reg. Stud.* **2005**, *39*, 61–74. [[CrossRef](#)]
35. Gertler, M.S. Tacit Knowledge and the Economic Geography of Context, or The Undefinable Tacitness of Being (There). *J. Econ. Geogr.* **2003**, *3*, 75–99. [[CrossRef](#)]
36. Saxenian, A. *Regional Advantage: Culture and Competition in Silicon Valley and Route 128*; Harvard University Press: Cambridge, MA, USA, 1994.
37. Feldman, M.P.; Audretsch, D.B. Innovation in Cities: Science-Based Diversity, Specialization and Localized Competition. *Eur. Econ. Rev.* **1999**, *43*, 409–429. [[CrossRef](#)]
38. Storper, M.; Venables, A.J. Buzz: Face-to-Face Contact and the Urban Economy. *J. Econ. Geogr.* **2004**, *4*, 351–370. [[CrossRef](#)]

39. Marshall, A. *Principles of Economics*, 8th ed.; Macmillan: London, UK, 1920.
40. Glaeser, E.L.; Kerr, W.R.; Ponzetto, G.A.M. Clusters of Entrepreneurship. *J. Urban Econ.* **2010**, *67*, 150–168. [CrossRef]
41. Nonaka, I.; Takeuchi, H. *The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation*; Oxford University Press: Oxford, UK, 1995.
42. Lave, J.; Wenger, E. *Situated Learning: Legitimate Peripheral Participation*; Cambridge University Press: Cambridge, UK, 1991.
43. DeFillippi, R.J.; Arthur, M.B. The Boundaryless Career: A Competency-Based Perspective. *J. Organ. Behav.* **1994**, *15*, 307–324. [CrossRef]
44. Grabher, G. Learning in Projects, Remembering in Networks? Communitarity, Sociality, and Connectivity in Project Ecologies. *Eur. Urban Reg. Stud.* **2004**, *11*, 103–123. [CrossRef]
45. Johansson, F. *The Medici Effect: What Elephants and Epidemics Can Teach Us About Innovation*; Harvard Business School Press: Boston, MA, USA, 2004.
46. Hargadon, A. *How Breakthroughs Happen: The Surprising Truth About How Companies Innovate*; Harvard Business School Press: Boston, MA, USA, 2003.
47. Meadows, D.H. *Thinking in Systems: A Primer*; Chelsea Green Publishing: White River Junction, VT, USA, 2008.
48. Sterman, J.D. *Business Dynamics: Systems Thinking and Modelling for a Complex World*; McGraw-Hill: Boston, MA, USA, 2000.
49. David, P.A. Clio and the Economics of QWERTY. *Am. Econ. Rev.* **1985**, *75*, 332–337.
50. Bresnahan, T.F.; Trajtenberg, M. General Purpose Technologies: Engines of Growth? *J. Econom.* **1995**, *65*, 83–108. [CrossRef]
51. Holland, J.H. *Hidden Order: How Adaptation Builds Complexity*; Addison-Wesley: Reading, MA, USA, 1995.
52. Almirall, E.; Fernández, M.; Tudurí, M.; Fernández, B. *Artificial Intelligence and the Public Sector: Scenarios for 2030*; Anteverti and Esade: Barcelona, Spain, 2025.
53. Følstad, A.; Brandtzaeg, P.B. Chatbots and the New World of HCI. *Interactions* **2017**, *24*, 38–42. [CrossRef]
54. Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; Küttler, H.; Lewis, M.; Yih, W.-T.; Rocktäschel, T.; et al. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*; Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H., Eds.; Curran Associates: Red Hook, NY, USA, 2020; pp. 9459–9474.
55. Larsen, A.G.; Følstad, A. The Impact of Chatbots on Public Service Provision: A Qualitative Interview Study with Citizens and Public Service Providers. *Gov. Inf. Q.* **2024**, *41*, 101927. [CrossRef]
56. van Noordt, C.; Misuraca, G. Artificial Intelligence for the Public Sector: Results of Landscaping the Use of AI in Government across the European Union. *Gov. Inf. Q.* **2022**, *39*, 101714. [CrossRef]
57. Almirall, E. *Ciutats I IA (Cities and AI)*; Esade Center for Innovation in Cities: Barcelona, Spain, 2025.
58. Newman, P.; Kenworthy, J. *The End of Automobile Dependence: How Cities Are Moving Beyond Car-Based Planning*; Island Press: Washington, DC, USA, 2015.
59. Senge, P.M. *The Fifth Discipline: The Art and Practice of the Learning Organisation*; Doubleday: New York, NY, USA, 1990.
60. Waymo. Waymo One Surpasses 500,000 Weekly Paid Rides. Waymo Blog. 2026. Available online: <https://waymo.com/blog/> (accessed on 25 March 2026).
61. Hawkins, A.J. Waymo Is Now Doing over 500,000 Paid Robotaxi Rides per Week. The Verge. 2026. Available online: <https://www.theverge.com/> (accessed on 25 March 2026).
62. Baidu. Apollo Go Autonomous Ride-Hailing Service Surpasses 20 Million Cumulative Rides. Baidu IR. 2026. Available online: <https://ir.baidu.com/> (accessed on 25 March 2026).
63. Pony.ai. Pony.ai Investor Presentation, Q4 2025. Available online: <https://pony.ai/> (accessed on 25 March 2026).
64. WeRide; Shenzhen Bus Group. WeRide and Shenzhen Bus Group Launch Shenzhen’s First Fully Driverless Robobus Line. WeRide Press Release. 2025. Available online: <https://ir.weride.ai/> (accessed on 25 March 2026).
65. QCraft. Autonomous Buses Operating in Chinese Cities: Status Report 2025. Available online: <https://www.qcraft.ai/> (accessed on 25 March 2026).
66. Zou, Y.; Xu, Z.; Wang, T.; Xiong, G.; Lin, Z.; Li, D. Generative AI-Driven Dynamic Information Prioritization for Enhanced Autonomous Driving. *IEEE Trans. Intell. Transp. Syst.* **2025**, 1–14. [CrossRef]
67. Meituan. Meituan Autonomous Delivery: Scaling Last-Mile Logistics with Autonomous Vehicles. Meituan Technical Blog. 2025. Available online: <https://tech.meituan.com/> (accessed on 25 March 2026).
68. Neolix. Neolix Autonomous Delivery Vehicle: 30,000 Global Orders and Counting. Company Report. 2025. Available online: <https://www.neolix.ai/> (accessed on 25 March 2026).
69. International Transport Forum. *The Shared-Use City: Managing the Curb*; OECD/ITF: Paris, France, 2018.
70. Milakis, D.; van Arem, B.; van Wee, B. Understanding autonomous vehicles: A systematic literature review on capability, impacts, planning and policy. *J. Transp. Land Use* **2017**, *10*, 1–18.
71. Florida, R. The Geography of Autonomous Vehicles. CityLab/Bloomberg. 2024. Available online: <https://www.bloomberg.com/citylab/> (accessed on 25 March 2026).

72. Alibaba Cloud. City Brain: Now in 23 Cities in Asia. Alibaba Cloud Blog. 2020. Available online: <https://www.alibabacloud.com/blog/> (accessed on 25 March 2026).
73. Caprotti, F.; Liu, D. Platform Urbanism and the Chinese Smart City: The Co-Production and Territorialisation of Hangzhou City Brain. *GeoJournal* **2022**, *87*, 1559–1573. [[CrossRef](#)]
74. Wei, J.; Li, Z.; Lyapustin, A.; Sun, L.; Peng, Y.; Xue, W.; Su, T.; Cribb, M. Tracking Air Pollution in China: Near Real-Time PM2.5 Retrievals from Multisource Data Fusion. *Environ. Sci. Technol.* **2021**, *55*, 12106–12115. [[CrossRef](#)] [[PubMed](#)]
75. Cowa Robot. Cowa Deploys 36 Sanitation Robots in 2.7 Million Square Meter Area in Shenzhen. Yicai Global. 2025. Available online: <https://www.yicaiglobal.com/> (accessed on 25 March 2026).
76. Guangzhou Municipal Government. *Guangzhou to Increase Unmanned Sanitation Fleet to 1,000 Units by 2026*; Government Press Release: Guangzhou, China, 2025.
77. Lattanzi, D.; Miller, G. Review of Robotic Infrastructure Inspection Systems. *J. Infrastruct. Syst.* **2017**, *23*, 04017004. [[CrossRef](#)]
78. Walmart. Walmart Drone Delivery Expands to 1.8 Million Additional Households. Walmart Corporate. 2025. Available online: <https://corporate.walmart.com/> (accessed on 25 March 2026).
79. South China Morning Post. Chengdu Deploys Robot Police Teams Combining Quadruped, Wheeled, and Humanoid Robots. SCMP. 2025. Available online: <https://www.scmp.com/> (accessed on 25 March 2026).
80. Xinhua. Hangzhou Places AI-Powered Traffic Policing Robots on Active Duty. Xinhua News Agency. 2025. Available online: <https://www.xinhuanet.com/> (accessed on 25 March 2026).
81. Shenzhen Municipal Government. *Shenzhen Designed as China's First "Robot-Friendly" Urban District*; Government Press Release: Guangzhou, China, 2025.
82. Landes, D.S. *The Unbound Prometheus: Technological Change and Industrial Development in Western Europe from 1750 to the Present*, 2nd ed.; Cambridge University Press: Cambridge, UK, 2003.
83. Goodman, E.P.; Powles, J. Urbanism Under Google: Lessons from Sidewalk Toronto. *Fordham Law Rev.* **2019**, *88*, 457–498. [[CrossRef](#)]
84. Shwayri, S.T. A Model Korean Ubiquitous Eco-City? The Politics of Making Songdo. *J. Urban Technol.* **2013**, *20*, 39–55. [[CrossRef](#)]
85. Marchetti, D.; Wangano, L.; Shukla, V.K. Smart City Initiatives in India: Lessons from Bengaluru and Beyond. In *Smart Cities: A Data Analytics Perspective*; Springer: Cham, Switzerland, 2021; pp. 215–233.
86. Yeung, K. 'Hypernudge': Big Data as a Mode of Regulation by Design. *Inf. Commun. Soc.* **2017**, *20*, 118–136. [[CrossRef](#)]
87. Eubanks, V. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*; St. Martin's Press: New York, NY, USA, 2018.
88. Benjamin, R. *Race After Technology: Abolitionist Tools for the New Jim Code*; Polity Press: Cambridge, UK, 2019.
89. Sadowski, J. When Data Is Capital: Datafication, Accumulation, and Extraction. *Big Data Soc.* **2019**, *6*, 1–12. [[CrossRef](#)]
90. Kitchin, R.; Dodge, M. The (In)Security of Smart Cities: Vulnerabilities, Risks, Mitigation, and Prevention. *J. Urban Technol.* **2019**, *26*, 47–65. [[CrossRef](#)]
91. Wachter, S.; Mittelstadt, B.; Floridi, L. Transparent, Explainable, and Accountable AI for Robotics. *Sci. Robot.* **2017**, *2*, eaan6080. [[CrossRef](#)]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.