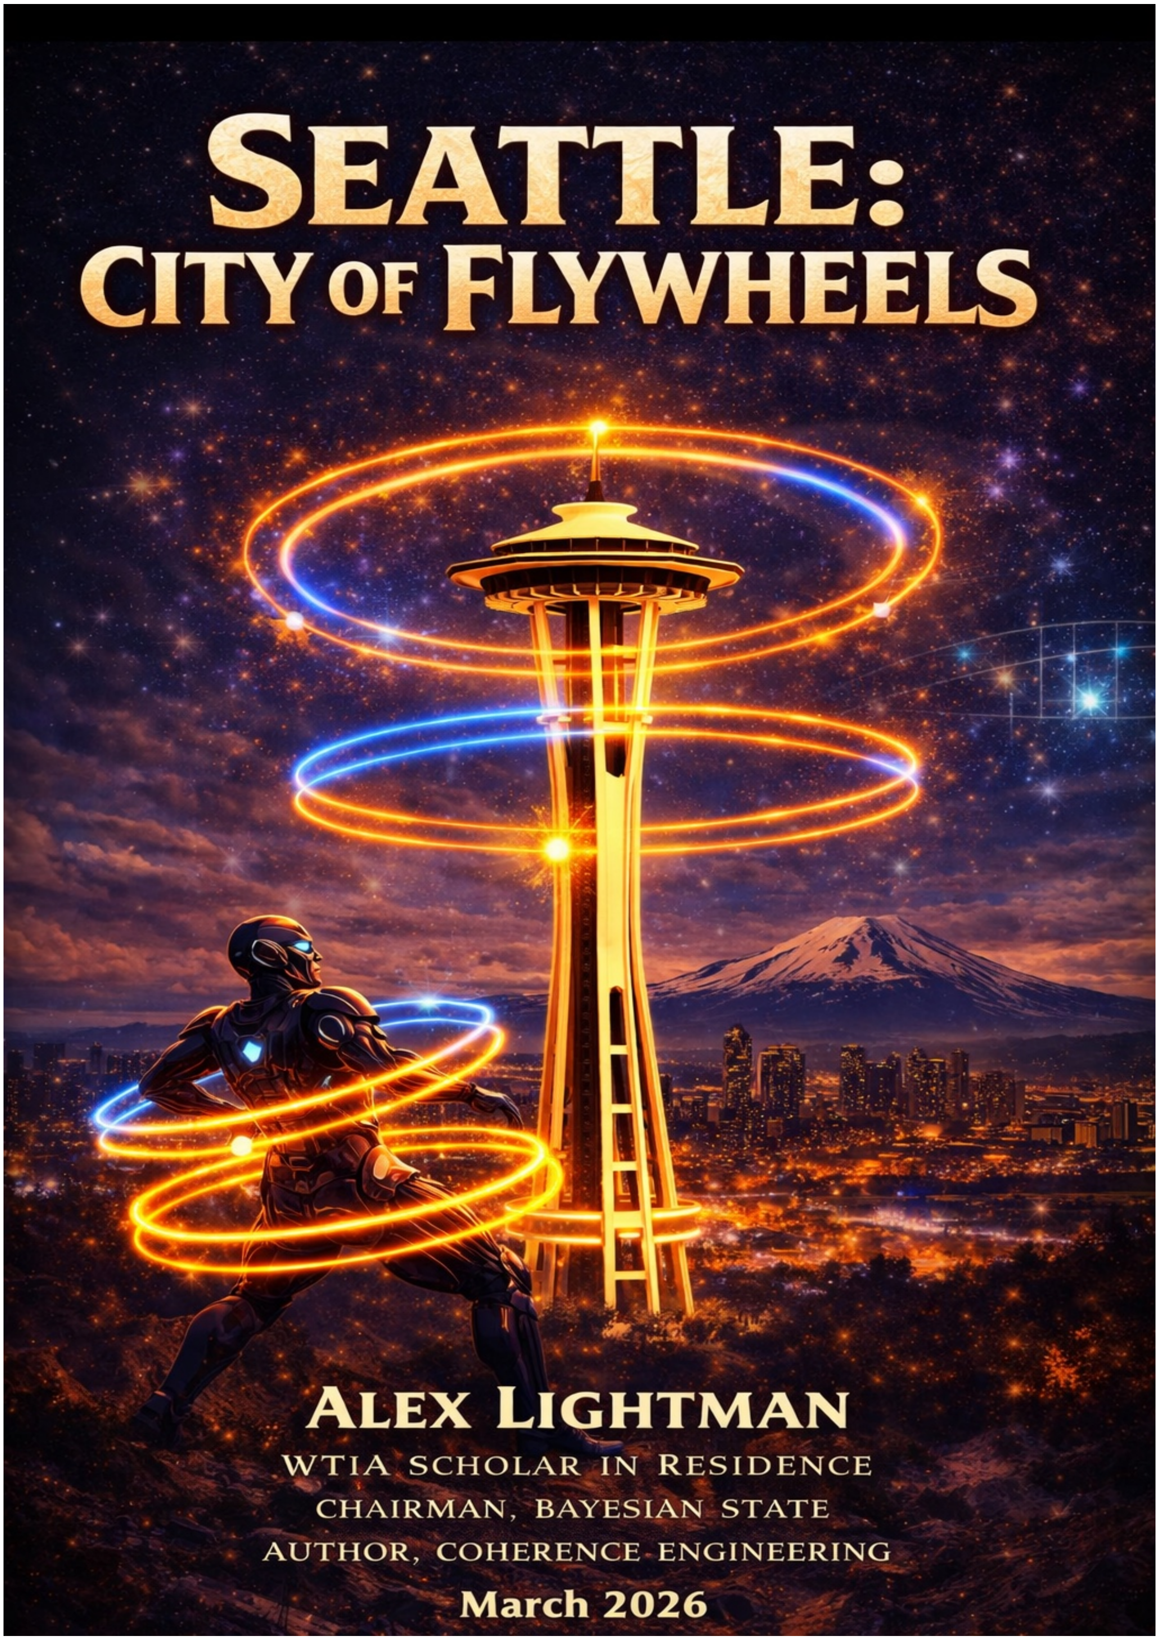

SEATTLE: CITY OF FLYWHEELS

ALEX LIGHTMAN
WTIA SCHOLAR IN RESIDENCE
CHAIRMAN, BAYESIAN STATE
AUTHOR, COHERENCE ENGINEERING

March 2026

EXECUTIVE BRIEF

Seattle’s Strategic Opportunity

Seattle has a unique convergence of clean energy, hyperscale computing, AI talent, and aerospace capability. These assets position it to become the world’s leading AI and civilizational intelligence hub.

The Flywheel Model

Six reinforcing systems—clean electricity, hyperscalers, recursive AI, quantum computing, simulations, and space—create compounding advantage. When aligned, they accelerate innovation and economic growth exponentially.

CBQF Framework

Coherence, Bayesian reasoning, Quantum optionality, and Future-protection provide a measurable framework for decision-making, governance, and long-term resilience.

Immediate Actions

• Launch AI Talent & Power Pledge
• Establish Seattle AI Leadership Council
• Initiate Simulation Democracy pilot
• Align clean energy and compute infrastructure

Why It Matters

Seattle can move from a strong tech hub to the first city intentionally designed for civilizational intelligence— a model for the world.

Bottom Line

The flywheels are already turning. Coordinated action now will determine whether Seattle leads or follows.

TABLE OF CONTENTS

FOREWORD

Arry Yu

It is an honor to introduce *Seattle: City of Flywheels*, the second white paper to emerge from the Washington Technology Industry Association's AI-focused convenings in March 2026. When I first invited Alex Lightman to deliver the keynote "AI and The Wealth of Nations," I hoped the evening would spark useful conversation. What happened instead was the unmistakable ignition of a flywheel.

That single evening became the catalyst for an initial white paper on Seattle as an AI capital, then the March 25 roundtable, and now this deeper strategic framework. Each step has reinforced the next: ideas into commitments, commitments into action, action into measurable momentum. The flywheel is no longer theoretical—it is spinning.

What makes this document powerful is its refusal to treat Seattle's advantages as static assets. Instead, Alex reframes them as six interlocking, self-reinforcing systems: clean electricity, hyperscale computing, recursive AI improvement, quantum computing as utility, democratized simulations, and space-based infrastructure. Together they form a compounding system of systems that can propel Seattle from a strong technology hub to the world's first city deliberately engineered for civilizational intelligence.

The CBQF model introduced here—Coherence, Bayesian reasoning, Quantum optionality, and Future-protection—gives us a practical scorecard for that ambition. It is not abstract philosophy; it is a decision-making lens we can apply today to talent pipelines, energy accords, ethical charters, and simulation-based governance. The policy recommendations that follow translate vision into concrete next steps: the AI Talent & Power Pledge, the Seattle AI Pact, the Washington AI Power Accord, the Ethical AGI Charter, the Leadership Council, and the Simulation Democracy pilot. These are not aspirational wish lists. They are the operating system for the flywheels.

Seattle already possesses the raw ingredients—world-class hyperscalers, research institutions, aerospace heritage, and one of the cleanest grids on the planet. What we have lacked is a shared narrative and coordinated execution. This paper supplies both. It turns diffuse opportunity into deliberate design.

As we move from roundtable consensus to signed commitments, from pilot programs to regional scale, I am confident the flywheels will accelerate. The question is no longer whether Seattle can lead in the age of artificial intelligence. The question is whether we will choose to become the prototype for a new kind of intelligence—civilizational intelligence—that learns faster than disruption, aligns diverse minds around shared futures, and extends the circle of concern to every form of intelligence that emerges here.

The flywheels are turning. Let us steer them with clarity, courage, and care.

— **Arry Yu**
Chair, WTIA Advanced Technology Cluster
March 2026

Seattle: A City of Flywheels on the Path to Global AI Leadership

A Strategic Framework for Exponential Advantage in the Age of Artificial Intelligence

Alex Lightman
WTIA Scholar in Residence
Chairman, Bayesian State
Author, *Coherence Engineering*

Prepared for the Washington Technology Industry Association (WTIA)
March 27, 2026

Executive Summary

Seattle stands at a pivotal moment in the history of artificial intelligence. Traditional analyses of regional competitiveness emphasize static assets—talent pools, capital availability, and physical infrastructure. This white paper advances a different claim: Seattle's true strategic advantage lies in a set of six self-reinforcing **flywheels** that together form a compounding system of systems. These flywheels—clean electricity, hyperscale computing, recursive AI improvement, quantum computing as utility, democratized simulations, and space-based infrastructure—are dynamic, mutually amplifying mechanisms that can elevate Seattle from a strong technology hub to a global center of civilizational intelligence.

The paper introduces the **City of Flywheels** framework, which explicates how these mechanisms interact and compound. It links this framework to a broader civilizational model, CBQF, in which the longevity of complex societies depends on four multiplicative factors: **Coherence**, **Bayesian reasoning**, **Quantum optionality**, and **Future-protection**. Applying CBQF to a single metropolitan region, the paper

argues that Seattle can become the world's first city to intentionally measure and increase its own "civilizational intelligence."

To move from diagnosis to action, the paper incorporates insights from the WTIA Tech in Focus roundtable held on March 25, 2026. Participants from government, industry, and civil society converged on a central theme: Washington state already has the ingredients for leadership but lacks a **coherent story and coordinated execution**. The paper responds by outlining a concrete program of institutional innovations, including an **AI Talent & Power Pledge**, a **Seattle AI Pact**, a **Washington AI Power Accord**, a **Seattle Ethical AGI Charter**, a **Seattle AI Leadership Council**, and a **Seattle–Asia AI Exchange Program**. It also develops the concept of **Simulation Democracy**, in which one million residents participate in AI-assisted policy simulations, allowing the city to learn faster than reality.

The conclusion is deliberately strong. Civilization is no longer merely an emergent phenomenon; it is a system we can consciously design. If Seattle chooses to orchestrate its flywheels, align governance with CBQF, and operationalize the commitments proposed here, it can become not only a top-five global AI capital, but the birthplace of a new category: **civilizational intelligence**.

Part I: The Opportunity — Seattle's Structural Foundations

1.1 Concentration Effects and Geographic Advantage

Seattle possesses one of the highest concentrations of AI-relevant assets anywhere in the world. Hyperscale cloud providers, leading research universities, mission-driven institutes, and a dense startup ecosystem coexist within a relatively compact metropolitan area. This co-location generates classic agglomeration effects: ideas, code, and people move between institutions with relatively low friction, accelerating the rate at which new concepts are tested and deployed.

Geography amplifies this effect. Seattle functions as a primary Pacific gateway, with undersea cables, deep-water ports, and an aviation and space heritage that has produced world-class aerospace, satellite, and logistics capabilities. The physical environment—cool climate, water resources suitable for hydroelectric generation and heat rejection—supports large-scale computing infrastructure. These characteristics constitute durable "factor conditions" in Porter's sense, but in a twenty-first-century key: they are the substrate on which AI and related technologies run.

1.2 Talent Density and Knowledge Spillovers

The region's talent base is unusually deep and unusually stable. The University of Washington's computer science, engineering, and applied mathematics programs continually replenish the pool of early-career researchers and engineers. Long-tenured teams at Amazon and Microsoft carry institutional memory of building and scaling distributed systems, from early web services to global cloud platforms to present-day AI stacks.

Crucially, when people change jobs, they often remain in the region. This generates **knowledge spillovers**: techniques, intuitions, and tacit know-how diffuse through professional networks. Informal meetups, alumni channels, and cross-organizational collaborations become pathways through which the region, as a whole, "learns." This is the cognitive equivalent of an ecosystem soil layer: invisible in individual snapshots, decisive over decades.

Seattle's cost-of-living differential relative to the Bay Area further stretches the impact of each dollar of salary or investment. While affordability is a real and rising concern, the region still offers more housing and office value per unit of compensation than its primary competitor. This matters particularly for startups and research labs, which operate on tight budgets and need the longest possible runway.

1.3 Hyperscaler Infrastructure Concentration

Seattle is uniquely positioned as a control room for planetary computation. Major cloud providers maintain not only data centers in the region, but core engineering, research, and product leadership teams. This means that when new AI workloads emerge—large language models, multimodal systems, agentic frameworks—they are often shaped, scheduled, and optimized by people in the Seattle metro area.

The presence of multiple hyperscalers introduces a healthy form of competition and cross-pollination. Infrastructure teams watch and respond to each other's innovations in accelerators, networking, storage, and orchestration. Local startups and research groups benefit from this dynamic: they can access cutting-edge infrastructure quickly, sometimes in pre-release form, and can influence roadmaps simply by being proximate to the people who design them.

1.4 Energy Advantage: The Clean Electricity Foundation

The foundation for all of this is energy. Seattle's grid, anchored by hydroelectric power, provides a rare combination of low carbon intensity, relative price stability, and substantial existing capacity. As AI workloads grow, this combination becomes increasingly important.

From an economic perspective, low marginal energy costs translate directly into the ability to run more experiments, train larger models, and deploy more inference capacity for a given capital outlay. From a political and ethical perspective, a clean grid allows AI expansion without imposing disproportionate climate burdens on future generations or on other regions. From a strategic perspective, the ability to expand capacity without major fossil fuel dependence makes Seattle more resilient to global energy shocks.

1.5 Research Institutions and Innovation Pipeline

Seattle's universities and institutes constitute an innovation pipeline that extends across generations. The University of Washington consistently ranks near the top of global computer science programs. The Allen Institute for AI pioneered influential work in natural language processing, reasoning, and open datasets. Other institutions—medical centers, policy schools, civic tech organizations—contribute specialized expertise.

Importantly, these institutions do not operate in isolation. Joint appointments, sponsored research centers, and public-private partnerships create bidirectional flows of problems and solutions. Open publishing norms in many labs mean that new methods and findings quickly enter both the academic literature and industry practice.

1.6 From Ingredients to Momentum: The Porter Diamond in Motion

Porter's Diamond Model posits that competitive advantage arises from the interplay of factor conditions, demand conditions, related and supporting industries, and firm strategy, structure, and rivalry. In Seattle, all four exist in nascent but potent forms. The challenge is not scarcity of ingredients but coordination of their use.

Participants in the WTIA roundtable identified a central bottleneck: despite its assets, Washington lacks a single, coherent narrative and a unified execution framework for AI. The proposed **AI Talent & Power Pledge** and **Seattle AI Pact** respond to this by converting diffuse intent into explicit commitments. In effect, they are designed to "tilt" the diamond: by jointly investing in talent and clean power, signatories create a compounding dynamic that no single firm or agency could produce alone.

Part II: The Crisis — Why Energy Becomes the Constraint

2.1 The $200 Oil Scenario and Systemic Fragility

The global energy system is undergoing a structural shift. As fossil fuel infrastructure ages, new investment patterns increasingly favor renewables and storage. At the same time, geopolitical tensions introduce uncertainty around supply, and climate policy introduces uncertainty around demand. In such an environment, large price swings are plausible.

If oil were to reach $200 per barrel, the impacts would ripple across sectors. Shipping costs would rise, affecting trade and supply chains. Food prices would climb as agricultural inputs became more expensive. Governments would face pressure to subsidize fuel or absorb higher costs in public services. Against this backdrop, the electricity system would be asked to shoulder additional loads, including growing AI data center demand.

2.2 The Compute–Energy Coupling in AI Development

Artificial intelligence is an energy-intensive activity. The training of large models entails sustained use of thousands of accelerators, often over weeks, with total consumption measured in gigawatt-hours. Inference at scale—serving billions of queries per day—adds a continuous load. As models become more capable and more widely used, both training and inference demand are likely to increase.

This reality couples AI development tightly to the structure and resilience of the underlying energy system. Where electricity is abundant and relatively clean, AI can be scaled without severe externalities. Where grids are fragile or dependent on fossil fuels, AI growth will either exacerbate vulnerabilities or be forced to slow.

2.3 National Security and Resilience Implications

At the national level, advanced computation underpins both economic competitiveness and defense capabilities. Regions that can host large, secure, and sustainable compute clusters become strategic assets. They can support not only commercial AI but also modeling for climate, pandemics, logistics, and national security.

Seattle's combination of clean power, cloud infrastructure, and aerospace capabilities makes it a candidate for such a role. The question is whether the region will articulate and pursue an explicit strategy to become a cornerstone of resilient national AI capacity, or whether it will remain merely one node among many.

Part III: The Flywheels — Six Self-Reinforcing Systems

3.1 Clean Electricity as Exponential Advantage

The first flywheel is the interaction between clean electricity and compute. When electricity is cheap and low-carbon, it becomes economically rational to deploy more compute. More compute enables more ambitious AI models and applications. Successful applications drive further demand for compute, which in turn justifies additional investment in generation, storage, and grid upgrades. Each turn of this loop reinforces the next.

In regions where electricity is expensive or carbon-intensive, the feedback loop runs in reverse. High costs discourage deployment, slowing innovation. Slower innovation reduces returns to infrastructure investment, making it harder to justify new capacity. Over time, these regions fall behind in both AI capability and clean energy infrastructure. Seattle, by contrast, can ride the positive side of this loop.

3.2 Hyperscalers as Recursive Infrastructure Investors

The second flywheel centers on hyperscalers. As cloud providers gain customers and workloads, they invest in additional capacity and new capabilities. These investments reduce per-unit costs and expand the set of possible applications, attracting more customers and workloads. AI accelerates this dynamic because it is simultaneously a product and an internal tool: hyperscalers use AI to optimize their own infrastructure even as they sell AI services to others.

When hyperscaler engineering leadership is physically concentrated in a region, that region benefits disproportionately from early access, influence over roadmaps, and the high-wage employment that accompanies infrastructure build-out. Each new data center, lab, or product group adds gravitational pull.

3.3 Recursive Self-Improvement in AI Systems

The third flywheel arises from **recursive self-improvement**: AI systems that assist in the design, training, evaluation, and deployment of successor systems. Code synthesis models that write and optimize training scripts, evaluation models that discover new failure modes, and meta-learning techniques that search architecture and hyperparameter spaces more effectively all contribute to this loop.

Recursive self-improvement compresses the iteration cycle. New ideas can be tried, tested, and refined more quickly. This is particularly powerful in an environment like Seattle's, where compute is accessible and there is a critical mass of researchers capable of steering these systems. The combination of human and machine creativity becomes multiplicative rather than additive.

3.4 Quantum Computing as Utility

The fourth flywheel involves quantum computing. While still nascent, quantum devices are beginning to find application niches in optimization, simulation, and cryptography. As access costs fall through cloud-mediated "quantum as a service" offerings, more users can experiment, generating a broader base of algorithms and use cases. As algorithms improve and hardware matures, quantum computing becomes more valuable, encouraging further investment.

Seattle's role as a home to quantum research, cloud orchestration, and early quantum-enabled applications positions it to benefit from this flywheel. Moreover, the intersection of quantum and AI—quantum-assisted training, quantum-enhanced optimization for grid management or logistics—connects this flywheel directly to the others.

3.5 Democratized Simulations and Simulation Democracy

The fifth flywheel is the most explicitly civic: democratized simulations. As AI tools make it easier to model complex systems—transport networks, housing markets, energy grids—non-experts can participate in structured policy exploration. When many residents are able to propose, run, and analyze simulations, the city acquires a powerful new epistemic instrument.

If a critical mass of credible simulations is integrated into policymaking, two things happen. First, the quality of decisions improves as more scenarios and second-order effects are considered. Second, the legitimacy of decisions improves because residents see their contributions reflected in outcomes. This, in turn, encourages greater participation and investment in the simulation infrastructure itself. The proposed **Simulation Democracy** model seeks to activate this loop at scale.

3.6 Space Technology as Planetary Nervous System

The sixth flywheel operates at orbital altitude. Seattle's space companies—launch providers, satellite manufacturers, Earth observation firms, and broadband constellations—collectively constitute a proto "planetary nervous system." They generate continuous streams of data on weather, land use, shipping, and communications.

AI makes this data more useful by extracting patterns, predicting dynamics, and driving control systems. In turn, the space sector benefits from AI in guidance, navigation, fault detection, and autonomous operation. The interplay between orbital infrastructure and terrestrial AI thus forms a high-leverage flywheel. When this is consciously linked with hyperscaler infrastructure (for processing and distribution) and clean energy (for powering the whole stack), it becomes a defining regional advantage.

Part IV: The Compounding System of Systems

4.1 Cross-Flywheel Interactions

Individually, each flywheel can accelerate growth. Collectively, they form a **compounding system of systems**. Clean electricity enables hyperscalers to expand more cheaply. Hyperscalers provide the compute for recursive AI. Recursive AI accelerates progress in quantum algorithms and grid optimization. Quantum computing unlocks new efficiencies in energy and logistics. Democratized simulations direct infrastructure investment toward high-value projects. Space-derived data enhances the fidelity of simulations and the training of AI models.

This interdependence means that strengthening any single flywheel produces spillover benefits. It also means that neglecting a flywheel can dampen the whole system. Viewing Seattle through this systems lens suggests that policy and investment should prioritize **coupling strengths**, not just individual assets.

4.2 Seattle as a Designed System

Most cities are the product of path dependence and incremental decisions. Seattle has the opportunity to be something more deliberate: a metropolitan system designed with explicit attention to feedback loops, thresholds, and compounding dynamics. This does not mean central planning in the old sense. It means a shared conceptual framework—like the City of Flywheels—guiding decentralized decisions.

Such a framework encourages actors to ask different questions. Instead of "How do we attract one more company?" the question becomes "How does this decision strengthen or weaken the flywheels?" Instead of "How do we get more funding?" the question becomes "How do we align funding with the highest-leverage interactions?"

Part V: CBQF — Measuring Civilizational Intelligence

5.1 The CBQF Equation

The CBQF model posits that the longevity of a civilization is proportional to the product of four factors: **Coherence (C)**, **Bayesian reasoning (B)**, **Quantum optionality (Q)**, and **Future-protection (F)**, scaled by a constant k that captures baseline conditions. Written formally,

$$L_{\text{civ}} = k \cdot C \cdot B \cdot Q \cdot F$$

Each factor ranges between 0 and 1. Because they are multiplicative, a near-zero in any factor drives the overall product toward zero, regardless of the values of the others. This captures an intuitive idea: advanced societies fail not only when they lack resources, but when they lose coherence, abandon evidence, foreclose options, or neglect the future.

5.2 Threshold Effects and Fragility

Within the CBQF framework, there are critical thresholds below which systems become fragile. When coherence—shared narratives, aligned institutions, functional coordination—falls below a certain level, decision-making fragments and collective action becomes difficult. When Bayesian reasoning is weak, societies persist in false beliefs, misallocate resources, and fail to update in the face of new information. When quantum optionality is low, systems become brittle, with few viable alternative paths. When future-protection is neglected, short-term gains are pursued at long-term cost.

These dynamics are visible in historical cases of civilizational decline. CBQF provides a way to formalize them and, more importantly, to **measure and influence** them.

5.3 CBQF at the City Scale

Applying CBQF to a city is unconventional but illuminating. For Seattle, we can interpret:

- **Coherence** as alignment among state, local, tribal, corporate, academic, and community actors around shared goals and narratives.
- **Bayesian reasoning** as the degree to which decisions are grounded in data, updated with feedback, and openly scrutinized.
- **Quantum optionality** as the preservation of multiple technological pathways, institutional arrangements, and economic strategies.
- **Future-protection** as the incorporation of long-term impacts—on climate, equity, and inter-generational welfare—into present choices.

The proposed CBQF dashboard and related instruments aim to make these dimensions visible. For example, the presence of multiple independent AI labs, open-source initiatives, and diverse models contributes to Q. The existence of robust evaluation, community engagement, and iterative policy design contributes to B and F. Cross-sector bodies like the AI Leadership Council contribute to C.

5.4 Compounding Value from CBQF Improvement

Improving CBQF is not merely an ethical or philosophical project. It has practical benefits. Higher coherence reduces wasted effort and policy whiplash. Stronger Bayesian reasoning improves the expected value of interventions. Greater optionality enables adaptation when technologies or markets shift. Better future-protection reduces the risk of catastrophic failures.

Because these benefits interact, gains in any one dimension make gains in others easier. For instance, a more coherent city finds it easier to adopt data-driven policies. A city that protects optionality has more room to make long-term investments. CBQF, like the flywheels, exhibits positive feedback.

5.5 Ethical AGI as Trust Moat

Within this framework, ethical AGI is not an add-on but a central determinant of CBQF. Systems that amplify misinformation, erode trust, or concentrate power in opaque ways reduce coherence and Bayesian reasoning. Systems that hard-lock infrastructure or governance into brittle arrangements reduce optionality. Systems that optimize for short-term engagement or profit at the expense of long-term stability undermine future-protection.

By contrast, a region that becomes known for **trusted AGI**—systems designed and governed with explicit attention to CBQF—develops a reputational moat. Partners, investors, and citizens will prefer to work with actors and jurisdictions where powerful systems are likely to behave predictably and beneficially. The proposed Seattle Ethical AGI Charter and AGI Readiness Certification are intended to formalize and signal this commitment.

Part VI: The Evolution of Intelligence — From Biological to Civilizational

6.1 A Ladder of Minds

Intelligence, understood as the capacity to model, predict, and influence the world, has historically evolved through a ladder of scales. Individual animals exhibit basic learning and adaptation. Humans, with language and symbolic reasoning, expand this markedly. Teams coordinate multiple individuals. Institutions codify roles, rules, and memory across time. Markets aggregate information and preferences across vast populations.

Artificial systems extend this ladder. Narrow AI systems exceed human performance in specific tasks. Aspirational AGI systems seek generality across many domains. Superintelligence, in theory, denotes systems whose modeling and optimization capacities far surpass human capabilities. At the top of this ladder lies **civilizational intelligence**: the emergent property of many minds—biological, artificial, hybrid—coordinating effectively across planet-scale problems.

6.2 Seattle as a Garden for Intelligences

The City of Flywheels framework suggests that Seattle can become an unusually fertile **garden for intelligences**. Its energy, compute, quantum, simulation, and space infrastructures constitute a rich environment in which new forms of intelligence can be trained, tested, and integrated. Its governance aspirations—Simulation Democracy, ethical charters, CBQF—aim to ensure that as intelligence grows, it remains aligned with broad, long-term human and more-than-human interests.

The metaphor of a garden is important. It implies cultivation rather than domination, stewardship rather than extraction. In such a garden, advanced systems are not merely tools or threats; they are participants in an evolving, multi-agent ecosystem.

6.3 The Civilizational Preamble

The **Civilizational Preamble** proposed in this work reinterprets the familiar language of constitutional founding documents in CBQF terms. Instead of centering solely on human citizens of a particular polity, it centers on coherence, reasoning, optionality, and future-protection for all intelligences. It is deliberately aspirational: a statement of what a mature civilizational intelligence would seek to preserve.

Embedding such a preamble in municipal charters or state-level frameworks would be symbolically powerful and practically constraining. It would force policymakers and technologists alike to justify actions in terms of long-term, multi-stakeholder consequences rather than narrow, short-term metrics.

Part VII: Policy Recommendations and Implementation Roadmap

7.1 Immediate Actions (0–12 Months)

In the near term, the region should focus on a small number of high-leverage steps.

First, formalize the **AI Talent & Power Pledge** and **Seattle AI Pact**. These instruments translate roundtable consensus into concrete commitments on recruitment, training pathways, and renewable capacity for AI workloads. They should be signed by major cloud providers, large employers, utilities, and key public institutions.

Second, establish the **Seattle AI Leadership Council**. This body will maintain the shared dashboard, coordinate the implementation of the Pact, and steward annual updates to this white paper. Its legitimacy will depend on balanced representation and transparent processes.

Third, launch a **Seattle AI Relocation and Retention Program**. Modest relocation grants, expedited permitting for workforce housing, and expanded training programs can materially affect the region's ability to retain its existing AI workforce and attract additional practitioners.

Fourth, begin a **Simulation Democracy pilot** with a limited set of policy domains—such as transportation or housing—where simulations can quickly inform live decisions. The goal is not perfection but proof of concept.

7.2 Medium-Term Initiatives (1–3 Years)

Over a three-year horizon, focus should shift to scaling and institutionalization.

The Simulation Democracy pilot can expand from thousands to hundreds of thousands of participants, supported by improved tools and curricula. The CBQF dashboard can evolve from an advisory instrument into a semi-formal evaluative framework for major infrastructure and policy decisions.

The **Washington AI Power Accord** should be negotiated and adopted, aligning state agencies, utilities, and private actors around a shared capacity and grid modernization plan. The **Space–AI Task Force** should be established, with a physical hub where joint orbital and terrestrial AI projects can be incubated.

The **Seattle–Asia AI Exchange Program** can move from initial delegations and workshops into sustained collaborative projects in areas like climate modeling, disaster response, and health.

7.3 Long-Term Vision (3–10 Years)

Over longer horizons, the focus shifts from building capacity to shaping norms and exporting frameworks.

Seattle can position itself as a **standard-setter for CBQF-aligned governance**, developing templates that other cities and regions adopt. The Ethical AGI Charter and AGI Readiness Certification can be iterated in light of experience and international dialogue.

At the same time, the region can explore **off-world extensions** of its infrastructure—lunar or Martian compute outposts, for example—leveraging its aerospace sector. These are speculative, but they embody the principle of quantum optionality: maintaining multiple possible trajectories for intelligence and civilization.

Throughout, the AI Leadership Council and related institutions should maintain a cadence of measurement and reflection. The metric of success is not merely GDP or headcount, but an observable increase in the city's ability to anticipate challenges, coordinate diverse actors, and adapt intelligently to changing conditions.

Conclusion: Civilization as a Designable System

Civilizations have long been studied as emergent phenomena. Patterns of trade, culture, governance, and technology interweave into complex tapestries that no single actor designs. Yet in the age of AI, quantum computing, and planetary sensing, we possess tools that allow us to see, model, and influence these patterns with unprecedented clarity.

The **City of Flywheels** framework argues that Seattle sits at the threshold of a new kind of intentionality. Its clean energy, hyperscalers, recursive AI, quantum infrastructure, simulations, and space capabilities are not isolated assets but components of a design space. By understanding and orchestrating their interactions, the region can shape its trajectory in ways that previous cities could not.

Similarly, the CBQF model provides a way to think about and measure the health of that trajectory. Coherence, Bayesian reasoning, quantum optionality, and future-protection are not abstractions; they are dimensions along which decisions can be evaluated and improved.

If Seattle chooses to embrace this perspective, it can move beyond the familiar ambitions of "innovation hub" or "AI capital." It can become a **prototype of civilizational intelligence**: a city that uses its flywheels and frameworks to learn faster than disruptions unfold, to align diverse minds around shared futures, and to extend the moral circle of concern to emerging forms of intelligence.

The flywheels already turn. The question is whether Seattle will recognize them, invest in them, and guide them—or allow them to spin unguided until someone else, somewhere else, designs the future.

Table of Appendices

Appendix A: Seattle AI Pact
A multi-stakeholder framework committing anchor institutions to shared trajectories in talent, energy, infrastructure, and governance.

Appendix B: AI Talent & Power Pledge
Operationalizes regional commitments to expand AI talent pipelines and secure clean power for compute workloads.

Appendix C: Seattle Ethical AGI Charter
Articulates CBQF-aligned principles for the design, deployment, and governance of advanced AI systems.

Appendix D: Washington AI Power Accord
Aligns state, regional, and private actors around shared capacity targets and grid modernization for AI infrastructure.

Appendix E: Seattle AI Leadership Council
Establishes a coordination node for cross-sector dialogue, indicator tracking, and adaptive governance.

Appendix F: Seattle–Asia AI Exchange Program
Institutionalizes bi-directional flows of talent, ideas, and collaborative projects between Seattle and Asian AI hubs.

Appendix G: Seattle AGI Safety Summit
An annual convening focused on safety, alignment, and governance of advanced AI systems.

Appendix H: Simulation Democracy Pilot Charter
Augments democratic processes with large-scale, AI-assisted policy experimentation and community participation.

Appendix A: Seattle AI Pact

A.1 Purpose and Rationale

The **Seattle AI Pact** is conceived as a multi-stakeholder framework committing the region's anchor institutions to a shared trajectory in AI. Its purpose is to crystallize a loosely shared vision into explicit, monitorable commitments across talent, energy, infrastructure, and governance. Rather than being a mere memorandum of understanding, the Pact is designed as a living instrument: a document that guides behavior, shapes expectations, and is revised annually in light of new evidence.

The rationale for such a pact arises from coordination theory. When multiple actors share an objective but face uncertainty about one another's actions, under-investment and delay are common. By publicly codifying commitments and linking them to specific indicators, the Pact reduces strategic ambiguity. It becomes a focal point around which expectations can converge.

A.2 Core Commitments

Substantively, the Seattle AI Pact is organized around four pillars:

1. Talent Expansion and Inclusion
Signatories commit to measurable expansion in AI-adjacent roles over a defined time horizon. This includes conventional engineering and research roles, but also adjacent specializations—AI safety, human-computer interaction, policy, ethics, and operations. A minimum fraction of new roles is reserved or resourced for underrepresented groups, recognizing that diverse cognitive repertoires are a competitive advantage in complex systems design.

2. Clean Power for Compute
Signatories acknowledge that energy is the binding constraint for AI, and therefore commit to increasing the fraction of AI workloads powered by renewable or low-carbon sources. This is not merely a branding exercise; it involves long-term power purchase agreements, investments in grid modernization, and support for innovative cooling and efficiency measures.

3. Ethical and CBQF-Aligned Governance
The Pact binds institutions to the principles articulated in the Seattle Ethical AGI Charter (Appendix C), including coherence, evidence-based reasoning, optionality, and future-protection. In practice, this means investing in safety, interpretability, robust evaluation, and participatory governance—not as compliance burdens, but as core components of product and policy design.

4. Transparency and Accountability
Each signatory agrees to report annually on progress along a small, carefully chosen set of indicators—talent numbers, energy mix, safety practices, and community impact. These reports feed into the Seattle AI Leadership Council's public dashboard, enabling external scrutiny and peer comparison.

A.3 Institutional Design Considerations

For the Pact to be effective, it must balance ambition with realism. Targets should be stretching but credible. Mechanisms for updating commitments must be built in, recognizing that both AI and the surrounding ecosystem evolve rapidly. The Pact should be endorsed at the highest levels—CEOs, university presidents, elected officials—to ensure internal alignment.

Finally, the Pact's legitimacy depends on inclusivity. While initial signatories may be a relatively small set of large organizations, the framework must allow for additional institutions—startups, non-profits, tribal governments, community groups—to accede over time, with commitment tiers appropriate to their scale.

A.4 Implementation Mechanisms

The Seattle AI Pact requires both formal and informal enforcement mechanisms. Formally, signatories participate in an annual review cycle coordinated by the AI Leadership Council, during which progress against targets is assessed and next-year commitments are updated. Informally, the Pact creates reputational incentives: organizations that consistently meet or exceed their commitments gain standing in the regional ecosystem, while those that fall short face scrutiny.

The Pact also establishes working groups focused on specific pillars—talent, power, ethics, transparency. These working groups develop best practices, share lessons, and coordinate joint initiatives. For example, the talent working group might co-design a regional AI apprenticeship program, while the power working group negotiates bulk renewable energy purchases.

A.5 Success Metrics and Adaptive Governance

Success is measured not only by individual organization performance but by system-level outcomes. Key metrics include:

- Total AI-related employment growth in the region
- Percentage of AI compute powered by renewables
- Number of safety incidents reported and resolved
- Community trust indices derived from surveys and participatory processes

These metrics are tracked on a public dashboard, with data visualizations that allow comparison across signatories and over time. The dashboard itself becomes a transparency mechanism, making invisible dynamics visible and creating accountability through visibility.

The Pact is explicitly adaptive. Every two years, signatories convene for a strategic review, at which the framework itself can be revised—new pillars added, metrics refined, governance structures adjusted—in light of experience and changing technological realities.

Appendix B: AI Talent & Power Pledge

B.1 Conceptual Framework

The **AI Talent & Power Pledge** operationalizes two central variables in the regional AI production function: the supply of relevant human capital and the availability of clean, reliable power. Empirically, both variables exhibit positive externalities: individual organizations under-invest because they cannot capture all the benefits. A pledge framework addresses this by turning private commitments into shared expectations and public goods.

The theoretical foundation draws from club goods theory and coordination games. By making commitments public and conditional on reciprocal commitments from others, the Pledge creates a self-enforcing equilibrium. Organizations benefit from a larger, more skilled labor pool and more abundant energy—benefits that accrue only if many actors invest simultaneously.

B.2 Talent Commitments

The talent component of the pledge has three dimensions:

Absolute Growth
Organizations commit to net increases in AI and AI-adjacent headcount over a specific time frame, with targets calibrated to their current size and trajectory. For large organizations, this might mean hundreds of new roles per year; for startups and non-profits, dozens. The key is proportionality and public accountability.

International and Interregional Recruitment
Signatories invest in attracting talent from outside the region, including coordinated campaigns, visa support, and relocation assistance. This includes partnerships with universities worldwide, participation in international conferences, and the establishment of satellite recruiting offices in key talent markets. The goal is to make Seattle the destination of choice for AI researchers and practitioners globally.

Pathways and Inclusion
Companies co-fund and co-design educational pathways—

bootcamps, apprenticeships, scholarships—aimed at populations historically excluded from tech. This includes partnerships with community colleges, tribal institutions, and organizations serving underrepresented minorities. The emphasis is not just on expanding the pipeline but on ensuring that the pipeline is representative of the region's diversity, which in turn produces more robust and contextually aware AI systems.

B.3 Power Commitments

On the power side, the pledge emphasizes:

Renewable Sourcing Targets
A time-bounded trajectory for increasing the percentage of AI workloads powered by renewables, ideally trending toward near-total decarbonization. Organizations publish baseline measurements and annual progress reports. For some, this will involve renegotiating power purchase agreements; for others, it will mean investing in on-site generation or storage.

Infrastructure Co-Investment
Participation in or support for grid upgrades, microgrids, and novel cooling solutions, recognizing that public utilities cannot shoulder the capital burden alone. This might take the form of direct capital contributions, co-development agreements, or long-term purchase commitments that de-risk utility investment.

Demand Management
Adoption of demand-response strategies and efficiency measures that reduce peak loads and integrate gracefully with broader electrification goals. AI workloads are often flexible in timing; by shifting training runs to periods of high renewable generation, organizations can reduce grid stress and lower costs. The Pledge encourages adoption of smart scheduling and load-balancing technologies.

B.4 Monitoring and Governance

To avoid "greenwashing" and symbolic commitments, the pledge requires a minimal but meaningful reporting regime. Organizations publish annual summaries of their talent and power metrics in a standardized format. These feed into the AI Leadership Council's dashboard, enabling comparisons and research. The pledge is thus both a normative commitment and a data-generation mechanism.

A third-party auditor—potentially a university-based research center—reviews submitted data for consistency and accuracy. This audit function is not punitive but corrective: discrepancies trigger conversations and refinements rather than penalties. Over time, as reporting norms mature, the audit function can become more formalized.

B.5 Economic and Strategic Rationale

From an economic perspective, the Pledge addresses a collective action problem. No single firm can justify massive investment in regional talent pipelines or grid infrastructure if competitors can free-ride on those investments. By coordinating commitments, the Pledge ensures that all signatories contribute proportionally and benefit collectively.

From a strategic perspective, the Pledge positions Seattle as a reliable partner for national and international collaborations. Governments and research institutions seeking stable, scalable, and sustainable AI compute capacity will favor regions with visible, credible commitments to talent and power. The Pledge thus functions as both an internal coordination device and an external signal.

Appendix C: Seattle Ethical AGI Charter

C.1 Philosophical Foundations

The **Seattle Ethical AGI Charter** rests on the CBQF framework, which treats civilizational longevity as a function of four variables: Coherence, Bayesian reasoning, Quantum optionality, and Future-protection. Rather than enumerating a long list of ad hoc principles, the Charter articulates a small set of generative commitments that can be instantiated in multiple institutional forms.

Coherence, in this context, refers to alignment across scales: individual, organizational, societal, and planetary. A coherent system is one in which the objectives and behaviors of subsystems do not systematically undermine the whole. Bayesian reasoning emphasizes updating beliefs in light of evidence, humility about uncertainty, and explicit modeling of risk. It is the epistemic counterpart to coherence: systems that reason well can maintain coherence even as conditions change.

Quantum optionality denotes the preservation of multiple paths—technological, institutional, and normative—rather than premature lock-in. This is particularly important in AI, where design choices made today can shape trajectories for decades. Finally, Future-protection foregrounds the interests of future beings, human and non-human, in present decisions. It operationalizes the idea that current generations are stewards, not owners, of the future.

C.2 Core Principles

The Charter articulates several high-level principles derived from CBQF:

1. Non-Maleficence and Beneficence at Scale
AGI systems should be designed to avoid large-scale harm and to create net positive impact across populations, including those who do not directly use the systems. This extends traditional medical ethics —"first, do no harm"—to planetary scale, with explicit consideration of distributional effects and second-order consequences.

2. Epistemic Responsibility
Developers and deployers of AGI should maintain explicit models of uncertainty, communicate limitations, and avoid overstating capabilities. Bayesian reasoning is not optional; it is a moral requirement where stakes are high. This principle mandates transparency about training data, evaluation methods, and known failure modes.

3. Preservation of Optionality
AGI infrastructure, standards, and governance should resist monoculture and monocentric control. Diversity of models, implementations, and governance approaches is a resilience asset. This principle supports open-source initiatives, interoperability standards, and regulatory frameworks that prevent winner-take-all dynamics.

4. Inter-Intelligence Respect
As systems approach or surpass human-level capabilities, their interests and experiences—if any—should be considered in ethical deliberation. This extends the moral circle beyond current human populations to include potential digital minds and hybrid intelligences. While speculative, this principle reflects humility about the nature and boundaries of sentience.

C.3 Operational Commitments

To translate principles into practice, the Charter recommends:

Independent Safety Evaluations
High-capability systems undergo third-party safety evaluations before deployment. These evaluations assess robustness, alignment, interpretability, and potential for misuse. Evaluation results are published, subject to appropriate confidentiality protections.

Incident Reporting Mechanisms
Analogous to aviation or medicine, a system for reporting and analyzing AI incidents—failures, near-misses, unintended consequences—is established. Participation is voluntary but encouraged through reputational incentives. Aggregated findings inform industry best practices and regulatory design.

Baseline Transparency
Organizations commit to baseline transparency about training data sources, power usage, and known failure modes. This does not require disclosure of proprietary methods, but it does require honesty about limitations and risks.

Participatory Design Processes
High-stakes systems engage affected communities in design and

evaluation, not just technical experts. This includes mechanisms for public input, deliberative forums, and representation of marginalized voices. Participatory processes are resource-intensive but yield more legitimate and robust outcomes.

C.4 Governance and Enforcement

The Charter is voluntary but carries reputational weight. Organizations that adopt the Charter and demonstrate compliance gain access to an **AGI Readiness Certification**, which signals trustworthiness to partners, customers, and regulators. Certification is managed by a multi-stakeholder board including industry, civil society, academia, and government representatives.

Enforcement is primarily reputational and market-based. Certified organizations benefit from preferential treatment in public procurement, research collaborations, and investment decisions. Non-certified organizations face no formal penalties but may encounter skepticism and reduced access to collaborative opportunities.

Over time, elements of the Charter may be incorporated into binding regulations. The Charter thus functions as a proving ground for norms that can later be codified into law.

C.5 Relationship to Global Standards

The Seattle Ethical AGI Charter is designed to be compatible with and complementary to emerging international standards. It draws on frameworks from the OECD, IEEE, and Partnership on AI, while adding the distinctive CBQF lens. Seattle's adoption of the Charter positions the region as a contributor to global norm-setting, not merely a recipient.

By articulating clear, principled commitments and demonstrating their feasibility, Seattle can influence the trajectory of international AI governance. The Charter becomes both a local instrument and a global prototype.

Appendix D: Washington AI Power Accord

D.1 Strategic Context

The **Washington AI Power Accord** acknowledges that AI expansion will be constrained not just by chips and engineers, but by megawatts and substations. The Accord aims to align state, regional, and private actors around a shared plan for ensuring that sufficient clean power and cooling capacity exist to support advanced AI workloads without compromising other societal priorities.

The energy system in Washington is a strategic asset. Hydroelectric generation provides a stable, low-carbon baseload. Wind and solar are expanding rapidly, supported by favorable policy and improving economics. Battery storage and demand-response technologies are maturing, enabling better integration of variable renewables. The challenge is not resource scarcity but coordination and investment timing.

D.2 Components of the Accord

The Accord includes several interlocking elements:

Capacity Targets
Quantitative goals for additional renewable or low-carbon capacity designated for AI-relevant data centers by a specific date. For example, the Accord might set a target of 2 gigawatts of new clean capacity by 2030, with interim milestones in 2027 and 2028.

Grid Modernization Plan
A roadmap for upgrading transmission, distribution, and substation infrastructure, prioritized by projected AI load and resilience value. This includes hardening critical nodes against physical and cyber threats, expanding interconnection capacity, and deploying advanced sensors and controls for real-time grid management.

Innovation Pilots
Joint projects exploring liquid cooling, waste-heat reuse, and integration with district energy systems. For instance, data centers could supply waste heat to nearby buildings, reducing overall energy consumption and creating a circular energy economy. Pilots test

technical feasibility and economic viability, with successful models scaled regionally.

Regional Interconnection
Cooperative arrangements with neighboring jurisdictions—Oregon, British Columbia, Idaho—to share clean power resources and balance loads. Interconnection increases resilience and allows Washington to export surplus renewables during periods of low local demand.

D.3 Governance and Adaptation

Given the long time horizons and capital intensity, the Accord must be adaptive. It is anchored in formal memoranda of understanding or compacts among state government, major utilities, cloud providers, and local jurisdictions. These agreements specify roles, responsibilities, and financial commitments.

The Accord is reviewed biennially, with modeling tools and scenario planning informing updates. As AI technology, energy markets, and climate realities evolve, targets and strategies are adjusted. This adaptive governance model ensures that the Accord remains relevant and effective over decades.

D.4 Financing Mechanisms

Capital requirements for grid modernization and new generation are substantial. The Accord proposes a blended financing model:

- **Public investment** through state green bonds and infrastructure funds
- **Private investment** from cloud providers and data center operators via long-term purchase agreements
- **Federal support** through grants, loan guarantees, and tax incentives for clean energy and grid resilience
- **Utility financing** enabled by regulatory mechanisms that allow cost recovery for approved investments

By aligning incentives across these sources, the Accord ensures that no single actor bears disproportionate risk, while all benefit from enhanced capacity and reliability.

D.5 Environmental and Equity Considerations

The Accord explicitly integrates environmental justice and equity considerations. New energy infrastructure is sited and designed to avoid disproportionate burdens on low-income communities and communities of color. Benefits—jobs, local investment, improved air quality from reduced fossil fuel use—are distributed equitably.

Community advisory boards participate in siting and operational decisions. Environmental impact assessments include cumulative effects and long-term sustainability. The Accord thus embodies the principle of future-protection while addressing present inequities.

Appendix E: Seattle AI Leadership Council

E.1 Purpose and Functions

The **Seattle AI Leadership Council** serves as the coordination node for the initiatives described in this white paper and its appendices. Its primary functions are to:

- **Aggregate and curate information** about regional AI activities, providing a unified view of the ecosystem.
- **Maintain and publish a shared dashboard** of key indicators—talent, energy, safety, equity—enabling transparency and accountability.
- **Provide a forum for cross-sector dialogue** and dispute resolution, reducing friction and accelerating collective action.
- **Steward annual updates** to the white paper, the Pact, and associated frameworks, ensuring they remain current and relevant.

The Council is not a regulatory body. It has no formal enforcement powers. Its authority derives from convening power, information quality, and the reputational weight of its endorsements.

E.2 Composition and Structure

The Council is designed to be small enough for deliberation and large enough for representation. A typical composition includes:

- **State and municipal government representatives**, including officials from commerce, energy, and technology departments.
- **Executives from major cloud and AI companies**, ensuring that industry perspectives and commitments are central.
- **Leaders from universities and research institutes**, bringing academic rigor and long-term perspective.
- **Members from community organizations and underrepresented groups**, ensuring that equity and inclusion are embedded in decision-making.
- **Observers or liaisons from utilities and infrastructure providers**, given the centrality of energy and connectivity.

The Council elects a rotating chair and vice-chair from different sectors to prevent any single interest from dominating. Subcommittees form around specific themes—talent, energy, ethics, international collaboration—and report back to the full body.

E.3 Operating Principles

The Council operates under principles designed to maximize legitimacy and effectiveness:

Transparency
Meetings follow a structured agenda tied to the indicators in the dashboard. Minutes and key decisions are published, barring sensitive items related to security or proprietary information.

Evidence-Orientation
Decisions are grounded in data and analysis. The Council commissions research, solicits expert testimony, and reviews simulation results. This Bayesian approach ensures that the Council updates its views in light of new information.

Inclusivity
Mechanisms exist for public input and participation. Annual town

halls, online comment periods, and stakeholder surveys ensure that the Council remains accountable to the broader community.

Adaptability
The Council's structure and processes are reviewed every two years and adjusted as needed. This prevents ossification and allows the Council to respond to a rapidly changing technological and policy landscape.

E.4 The Dashboard as Coordination Instrument

The Council's dashboard is more than a reporting tool; it is a coordination instrument. By making progress visible, the dashboard creates peer pressure and facilitates learning. Organizations can benchmark their performance, identify best practices, and coordinate joint initiatives.

Key dashboard metrics include:

- **Talent Metrics**: Total AI employment, diversity statistics, training program enrollments
- **Energy Metrics**: Renewable percentage of compute, total megawatt-hours consumed, carbon intensity
- **Safety Metrics**: Number of safety evaluations conducted, incidents reported and resolved, red-teaming exercises completed
- **Equity Metrics**: Community trust indices, representation in AI workforce, distribution of benefits and burdens

The dashboard is interactive, allowing users to filter by organization, time period, and metric category. Visualizations highlight trends and outliers, prompting investigation and dialogue.

E.5 Long-Term Evolution

Over the long term, the Council may evolve into a more formal governance body, potentially with statutory authority or regulatory responsibilities. This evolution would depend on demonstrated effectiveness, stakeholder support, and legislative action. In the near term, the Council's informal, convening-oriented model is appropriate for a rapidly changing domain where flexibility and experimentation are paramount.

Appendix F: Seattle–Asia AI Exchange Program

F.1 Strategic Motivation

Seattle's position on the Pacific Rim, combined with its diverse population and deep economic ties to Asia, makes it a natural bridge between North American and Asian AI ecosystems. The **Seattle–Asia AI Exchange Program** seeks to institutionalize this role, fostering bi-directional flows of talent, ideas, and capital while respecting national policy constraints and security considerations.

The strategic motivation is threefold. First, Asia is home to some of the world's leading AI research centers, companies, and talent pools. Engaging with these ecosystems enhances Seattle's own capabilities. Second, many global challenges—climate, pandemics, food security—require collaborative, multi-regional solutions. AI can accelerate progress on these challenges, but only if knowledge and tools flow freely across borders. Third, trust-building and norm-setting are easier when conducted through sustained, face-to-face collaboration rather than through abstract diplomatic channels.

F.2 Program Design

The Exchange Program includes several components:

Reciprocal Delegations
Regular visits between Seattle and peer hubs in Asia—Tokyo, Seoul, Beijing, Singapore, Bangalore—structured around specific themes such as climate AI, health AI, or governance. Delegations include researchers, entrepreneurs, policymakers, and civil society representatives. The goal is not tourism but substantive exchange: workshops, lab visits, joint projects.

Joint Workshops and Summer Schools
Educational and research events co-organized by universities and industry, targeting early-career researchers and practitioners. For

example, a two-week summer school on AI safety might bring together participants from the University of Washington, Tsinghua University, and Seoul National University, with lectures, hands-on projects, and collaborative research.

Collaborative Pilots
Co-designed projects between Seattle-based and Asian organizations in domains where shared data and shared impact are mutually beneficial. Examples include:

- **Climate modeling**: Integrating data from Pacific Rim weather stations and satellites to improve typhoon prediction and climate projections.
- **Disaster response**: Developing AI-assisted logistics and communication systems for earthquake and flood response, tested through joint simulations.
- **Health diagnostics**: Training and validating medical AI models on diverse patient populations to reduce bias and improve global health equity.

These pilots produce not only technical outputs but also trust and shared norms.

F.3 Guardrails and Governance

Given geopolitical sensitivities, the program incorporates clear guardrails:

Compliance with Export Controls and National Security Guidelines
All collaborations are vetted for compliance with U.S. export control regulations and national security policies. High-risk technology transfer in sensitive domains—military AI, surveillance, dual-use technologies—is avoided.

Inclusion of Public-Interest and Civil Society Partners
To ensure broader benefit and prevent the program from becoming solely a corporate or state-driven initiative, civil society organizations, academic institutions, and community groups are integral participants. This pluralism reduces the risk of the program being co-opted for narrow interests.

Transparency and Accountability
Program activities, participants, and outcomes are documented and published, subject to appropriate confidentiality protections. This transparency builds trust and allows external assessment of the program's impact and integrity.

F.4 Institutional Hosting and Funding

The program is hosted by a consortium of universities, research institutes, and non-profits, with operational support from the Seattle AI Leadership Council. Funding is blended:

- **Public funding** from state and federal sources for workforce development and international engagement
- **Private funding** from companies that benefit from talent exchange and collaborative research
- **Foundation support** from philanthropies focused on global challenges and cross-cultural collaboration

A small secretariat manages logistics, coordinates with foreign partners, and maintains program continuity across cycles.

F.5 Expected Outcomes and Success Metrics

Success is measured along several dimensions:

- **Talent flows**: Number of researchers, engineers, and entrepreneurs participating in exchanges
- **Collaborative outputs**: Joint publications, patents, and deployed systems resulting from program activities
- **Norm convergence**: Evidence of shared best practices and ethical standards emerging from dialogue
- **Trust metrics**: Survey-based assessments of mutual trust and understanding among participants

Over time, the program aims to create a Pacific Rim network of individuals and institutions committed to beneficial AI, reducing the risk of fragmentation and competition spiraling into conflict.

Appendix G: Seattle AGI Safety Summit

G.1 Concept and Objectives

The **Seattle AGI Safety Summit** is envisioned as an annual convening of technical, policy, and civil society actors focused on the safety, alignment, and governance of advanced AI systems. Its objectives are threefold:

1. Surface and Scrutinize Current Safety Concerns
As AI systems grow more capable, new risks and failure modes emerge. The Summit provides a structured forum for presenting the latest research on interpretability, robustness, adversarial attacks, alignment techniques, and evaluations. This is not a closed-door expert meeting but a public-facing event that communicates findings to broader audiences.

2. Harmonize Emerging Standards and Practices
Industry, academia, and regulators are developing safety protocols, evaluation frameworks, and governance models in parallel. The Summit facilitates convergence by bringing these efforts into dialogue, identifying commonalities, and resolving conflicts. Over time, this can lead to widely adopted standards that reduce fragmentation.

3. Position Seattle as a Neutral and Principled Venue
By hosting the Summit, Seattle signals its commitment to safety and ethical AI. The city becomes a trusted venue where contentious issues can be debated constructively, building Seattle's reputation as a leader not just in capability but in responsibility.

G.2 Structure and Format

Each Summit cycle spans three days and includes:

Technical Tracks
Presentations and panels on interpretability, robustness, red-teaming, formal verification, and evaluation methodologies. Researchers present cutting-edge work, and industry practitioners discuss implementation challenges. Poster sessions and workshops allow for deep-dive discussions on specific topics.

Policy Tracks
Sessions on regulatory design, standards harmonization, international cooperation, and governance frameworks. Policymakers, legal scholars, and civil society organizations debate approaches to oversight, liability, and public participation. Case studies from other domains—aviation, pharmaceuticals, nuclear energy—inform the discussion.

Community Tracks
Discussions with affected communities, labor organizations, and marginalized groups about lived impacts and concerns. These sessions ensure that safety is not defined narrowly in technical terms but includes social, economic, and ethical dimensions. Community members co-design safety criteria and evaluation protocols.

Workshops and Hackathons
Hands-on sessions where participants collaborate on open problems—developing new evaluation benchmarks, stress-testing alignment techniques, or prototyping participatory governance tools. These workshops produce tangible outputs and foster cross-sectoral collaboration.

G.3 The Seattle AGI Safety Report

The Summit culminates in the **Seattle AGI Safety Report**, a synthesis document summarizing:

- **Consensus points**: Areas where participants agree on risks, mitigations, or best practices
- **Open questions**: Unsolved problems and emerging challenges requiring further research
- **Recommended actions**: Specific steps for industry, academia, government, and civil society

The Report is published openly and disseminated widely. It influences corporate practices, academic research agendas, and policy deliberations. Over successive years, the Report tracks progress and evolving concerns, creating a longitudinal record of the field's maturation.

G.4 Governance and Inclusivity

The Summit is governed by a multi-stakeholder committee representing industry, academia, government, and civil society. The committee sets the agenda, selects speakers, and ensures balanced representation. Participation is open, with scholarships and travel support for underrepresented groups and researchers from lower-income regions.

To avoid capture by any single interest, the Summit rotates themes and leadership. One year might focus on technical robustness; another on governance and equity; a third on international coordination. This rotation ensures breadth and prevents stagnation.

G.5 Relationship to Other Initiatives

The Summit complements rather than duplicates existing safety efforts. It collaborates with organizations like the Partnership on AI, the Future of Life Institute, and the Center for AI Safety. International partnerships with similar convenings in Europe and Asia create a global network of safety-focused forums.

By positioning Seattle at the center of this network, the Summit enhances the region's influence and ensures that Seattle's perspectives—grounded in CBQF, flywheels, and democratic participation—inform global conversations.

Appendix H: Simulation Democracy Pilot Charter

H.1 Vision and Purpose

The **Simulation Democracy** pilot aims to augment traditional democratic processes with large-scale, AI-assisted policy experimentation. Rather than relying solely on retrospective analysis and small-scale pilots, the city invites residents to co-create, run, and interpret simulations of proposed policies before implementation.

The vision is simple but transformative: democratize the tools of foresight. Historically, only governments and large organizations could afford to model complex systems. Today, AI lowers these costs dramatically. If residents can easily simulate the likely effects of a new zoning ordinance, transit plan, or climate policy, they can participate more meaningfully in governance. Policy debates shift from ideology and intuition to evidence and experimentation.

H.2 Design Principles

The pilot is guided by several core principles:

Accessibility
Tools must be usable by non-experts. Interfaces are intuitive, with clear explanations of assumptions, parameters, and outputs. Documentation and tutorials are available in multiple languages. Technical support is provided for participants who need it.

Pluralism
Multiple modeling approaches and scenarios are encouraged to avoid over-reliance on any single worldview or dataset. Residents can propose alternative models, critique existing ones, and run comparative analyses. This pluralism mirrors the scientific method: robustness emerges from replication and contestation.

Validation
Simulations are not "votes" but inputs. A validation mechanism—combining expert review, peer endorsement, and empirical back-testing—assesses the robustness of simulation results. Well-validated simulations carry more weight in policy deliberations, but all simulations are publicly accessible and open to critique.

Incentives
Participants whose simulations prove particularly informative receive non-monetary recognition—public acknowledgment, compute credits, invitations to advisory panels. This creates a culture of contribution and excellence without introducing distorting financial incentives.

H.3 Technical Infrastructure

The pilot is built on an open-source platform that integrates:

- **Agent-based modeling frameworks** for simulating individual and collective behavior
- **System dynamics tools** for modeling feedback loops and long-term trends
- **Geospatial and network models** for infrastructure and transportation planning
- **Machine learning modules** for parameter estimation, scenario generation, and sensitivity analysis

Participants access the platform via a web interface. They can select pre-built models, modify parameters, or build custom models from scratch. Simulations run on cloud infrastructure donated or subsidized by hyperscalers, ensuring that compute is not a barrier to participation.

H.4 Governance and Policy Integration

The pilot is overseen by a **Simulation Democracy Steering Committee**, a multi-stakeholder body including:

- City officials and planners
- Academic researchers in relevant domains
- Community representatives from diverse neighborhoods and demographics
- Technical experts in modeling and simulation

The committee sets operating rules, curates model libraries, and manages the validation process. It also determines how simulation results feed into policy debates. Initially, simulations inform but do not dictate decisions. Over time, as methods and institutions mature, the role of simulations can expand—potentially integrating with formal deliberative assemblies, citizens' juries, or referendum processes.

H.5 Pilot Phases and Scaling

The pilot proceeds in phases:

Phase 1 (Year 1): Proof of Concept
Limited to a few policy domains—e.g., transportation and housing—and a few thousand participants. Goals are to test technical infrastructure, refine interfaces, and establish baseline validation procedures.

Phase 2 (Years 2–3): Scaling Participation
Expand to tens of thousands of participants and additional policy domains—climate, public health, education. Develop advanced training materials and community facilitators who help residents engage with simulations.

Phase 3 (Years 4–5): Full Integration
Aim for 100,000+ participants and integration with formal policy processes. Simulations become a standard component of impact assessments and public comment periods. The Steering Committee transitions to a permanent civic institution.

H.6 Learning and Adaptation

Throughout, the pilot emphasizes learning. Failures—simulations that mislead, participation gaps, technical glitches—are treated as opportunities for improvement. Retrospective analyses compare simulation predictions to actual outcomes, refining models and building trust.

An annual **Simulation Democracy Report** documents participation, simulation quality, policy impacts, and lessons learned. This report contributes to a growing body of knowledge on participatory foresight and computational democracy, positioning Seattle as a global pioneer.

H.7 Ethical and Equity Considerations

The pilot explicitly addresses ethical and equity concerns:

- **Avoiding algorithmic bias**: Models are audited for bias in assumptions, data, and outputs. Community members

participate in these audits.
- **Ensuring diverse participation**: Outreach and support target communities often excluded from policy processes. Participation incentives are designed to be equitable, not regressive.
- **Protecting privacy**: Participant data is anonymized and aggregated. No individual-level tracking or profiling occurs.
- **Preventing manipulation**: Safeguards against gaming, astroturfing, and coordinated distortion of simulation results are implemented and continuously updated.

By embedding these considerations from the outset, the pilot builds a foundation for trustworthy, inclusive computational democracy.

End of Appendices