

Article

The Hidden Net Cost of Data Center Construction and Operation for Household Service Pricing

Arezou Shafaghat ^{1,2,*}, Mikhail Klimenko ³, Da Hu ² and Ali Keyvanfar ⁴¹ Sustain X Scale Lab, Scheller College of Business, Georgia Institute of Technology, Atlanta, GA 30308, USA² Department of Civil and Environmental Engineering, Southern Polytechnic College of Engineering and Engineering Technology, Kennesaw State University, Marietta, GA 30060, USA; dhu3@kennesaw.edu³ School of Economics, Ivan Allen College of Liberal Arts, Georgia Institute of Technology, Atlanta, GA 30308, USA; mikhail.klimenko@econ.gatech.edu⁴ Department of Construction Management, College of Architecture and Construction Management, Kennesaw State University, Marietta, GA 30060, USA; akeyvanf@kennesaw.edu

* Correspondence: ashafaghat3@gatech.edu

Abstract

The rapid expansion of artificial intelligence (AI) is accelerating data center construction and creating downstream implications for households. This study examines how AI-era data center costs (comprising construction, energy, water, grid upgrades, cooling, and lifecycle management) move through service supply chains and affect household prices in healthcare, transportation, education, banking, and commerce. It also considers the productivity and welfare benefits that AI may transmit. This study identifies four pass-through channels: utility-rate socialization of energy costs, cloud-platform pricing, sectoral pass-through from AI-adopting industries, and indirect effects through supply chains and labor markets. It introduces the AI-inflated net good basket, defined as transmitted cost minus transmitted benefit, to show how AI reshapes the overall net cost of household consumption rather than simply inflating individual prices. The study develops the AI Infrastructure Net Cost Pass-Through Model (AI-NCPM), a four-layer conceptual framework tracing net cost flows from data center investment to sectoral allocation and household outcomes. The model's parameters are analytically specified but not empirically calibrated; numerical examples are illustrative rather than representing estimated effects. Its main contribution is an integrative framework linking cost pass-through, infrastructure cost socialization, two-sided platform allocation, environmental externalities, and household expenditure incidence within a single net-cost account. Because these effects originate in the design, construction, energy and cooling systems, and lifecycle operation of data centers, the analysis connects AI infrastructure economics to the built environment. The framework suggests that low-income, minority, rural, older adult, and disability-affected households may face disproportionate net burdens, as costs fall heavily on essential services while benefits accrue more readily to affluent and digitally connected households.



Academic Editor: Yasser Mohamed

Received: 15 June 2026

Revised: 10 July 2026

Accepted: 14 July 2026

Published: 15 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: data center; artificial intelligence; Net Cost Pass-Through; two-sided platforms; household affordability; net good basket; digital economy externalities

1. Introduction

As workloads grow rapidly, driven by machine learning, deep learning, and generative models [1–3], the demand for data centers to process and store such workloads is rising sharply [4,5]. Tech giants like Google, Amazon, Microsoft, Apple, and Meta are locked in an

arms race to build hyperscale data centers that can keep pace with the AI boom [6]. These facilities, often spanning hundreds of thousands of square feet and consuming electricity on par with small cities, are the physical backbone of the AI revolution [7].

Yet, there is an overlooked link in the AI value chain: the connection between rapidly growing AI demand, the physical infrastructure needed to support it, and the costs that are ultimately borne by households. As tech firms pour billions into data center construction and operation to gain an edge in the AI market, the costs of this digital arms race are being absorbed, often invisibly, into the pricing of services that households rely on daily. Moreover, this expansion may generate negative externalities—costs imposed on third parties who did not choose to incur them [8]. Data center energy consumption, for example, can strain local power grids and contribute to carbon emissions, while the water demands of cooling systems can stress local water resources [5,7].

This study argues that the economics of AI cannot be fully understood without examining the infrastructure behind it, and that infrastructure transmits both costs and benefits. As firms scale AI capabilities, data center development and operation create cost pressures that extend beyond firms and utilities into household expenditures, even as the same investment lowers the cost of many downstream services. This study examines how the escalating costs of constructing, operating, and managing AI-era data centers are transmitted through service supply chains and reflected in the prices households pay for essential services across primary sectors (healthcare, transportation, education, banking, and commerce), and how a portion of AI's productivity and welfare gains is passed back to those same households; so the household-relevant quantity is net cost, not cost alone. Accordingly, the objectives are six-fold: (1) to quantify the scale of AI-driven data center construction demand; (2) to analyze data centers' operational cost structure; (3) to map the cost and benefit transmission mechanisms and their net incidence on two-sided platforms—to household service pricing; (4) to develop a conceptual 'AI Infrastructure Net Cost Pass-Through Model'; (5) to assess net equity impacts across household segments; and (6) to identify sustainability and policy implications. Framed this way, the AI-era data center is treated as a distinct building typology whose design, construction, energy and cooling systems, siting, and full lifecycle are the physical origin of the economic effects traced in this study. The pass-through cascade developed below does not begin in the abstract space of markets; it begins with concrete decisions in the built environment—how these facilities are planned, engineered, powered, and eventually decommissioned—and it is these building- and infrastructure-scale choices that ultimately propagate into household service prices. Positioning the analysis in this way connects the economics of AI infrastructure directly to the construction sector, infrastructure planning, and building lifecycle concerns that are central to the scope of this journal.

Contribution and positioning. The primary contribution of this study is a new conceptual framework (the AI Infrastructure Net Cost Pass-Through Model (AI-NCPM)) that integrates five previously separate bodies of literature into a single, household-centric account of net cost incidence. Work on cost pass-through has generally analyzed a single price or sector; infrastructure cost socialization has focused on utility ratepayers rather than the downstream household basket; the two-sided platform literature has modeled elasticity-based allocation without connecting it to physical infrastructure costs; the environmental externality literature has treated data-center impacts as a sustainability question rather than a household-pricing one; and household expenditure-incidence research has rarely traced burdens back to an upstream infrastructure origin. The novelty of the AI-NCPM is not any one of these mechanisms in isolation but their unification: it links a physical building-and-infrastructure origin to a multi-channel, multi-sector cascade, applies two-sided allocation at each intermediating platform, nets transmitted benefits against

transmitted costs to define an AI-inflated net good basket, and embeds distributional vulnerability as a structural element rather than an afterthought. Three further contributions follow from, and are subordinate to, this framework: (i) a policy-oriented synthesis that organizes dispersed evidence on data-center costs, rate socialization, cloud pricing, and affordability around a common structure; (ii) an analytical scaffolding; the parameterized index form of Section 7; that specifies the quantities a future formal model of pass-through would need to estimate; and (iii) a structured empirical research agenda identifying the data and identification strategies required to calibrate the framework. The study does not itself estimate pass-through elasticities or solve a structural equilibrium; its aim is to provide the integrative conceptual architecture on which such work can build.

In doing so, it reintroduces the ‘good basket’ concept in net terms (the idea that AI is not just another item in the household consumption basket, but a force reshaping the net cost of the basket itself, where net cost is transmitted cost minus transmitted benefit) tracing AI’s footprint from silicon to sectoral service pricing. It further recognizes that hyperscalers and the platforms that intermediate these flows are two-sided markets, so cost incidence is split between merchants and consumers rather than borne by households alone.

The remainder of the study is structured as follows. Section 2 describes the research design and analytical approach; Section 3 examines AI growth and data center expansion; Section 4 analyzes the cost anatomy of AI-era data center development; Section 5 details the cost and benefit pass-through mechanisms; Section 6 presents the household net good basket perspective; Section 7 introduces the AI-NCPM framework, including its two-sided platform treatment of net cost incidence; Section 8 outlines the limitations of the study; and Section 9 offers conclusions and future research. It should be emphasized at the outset that the AI-NCPM is developed here as a conceptual framework: its pass-through rates, benefit-transmission parameters, and vulnerability weights are specified analytically and are not yet calibrated or validated against empirical data, a limitation discussed and addressed as a priority for future work.

2. Research Design and Analytical Approach

2.1. Research Questions and Study Type

The study is organized around three framework-oriented research questions: (RQ1) how are AI-era data center capital and operational costs transmitted to household prices? (RQ2) how are these costs, and AI’s offsetting benefits, distributed across household segments with different expenditure patterns and benefit access? and (RQ3) what analytical constructs are needed to make this transmission and allocation empirically estimable in future work? Accordingly, the study is conceptual and theory-developing rather than hypothesis-testing, producing a validated framework and testable research agenda rather than parameter estimates.

2.2. Methodological Approach

The framework consists of four steps. First, five categories of literature (cost incidence and tax-shifting, infrastructure cost socialization, two-sided platform pricing, environmental externalities, and household expenditure incidence) were reviewed to identify relevant mechanisms and gaps. Second, these mechanisms were organized into a four-layer cascade: cost origination, channel allocation, household transmission, and distributional weighting. Third, each layer was formalized with defined inputs, outputs, and admissible ranges, producing the index model. Fourth, worked scenarios demonstrated the framework’s logic and identified quantities for future empirical measurement. This sequence follows a design-science logic adapted to economic modeling.

2.3. Analytical Indices and Constructs

The indices are the constructs and are listed in Table 1; they serve as the design instruments. The cost-origination index $C(j)$ measures the cost generated in each data center category j (pre-construction, construction, operations, grid, lifecycle). The allocation index $A(j,k) \in [0,1]$ measures the share of $C(j)$ routed to transmission channel k . The pass-through index $P(k) \in [0,1]$ measures the fraction of allocated cost actually reaching households and is not a free parameter: it is structurally determined by upstream market power $\mu(k)$ and downstream demand elasticity $\varepsilon(k)$ through $P(k) = f(\mu(k), \varepsilon(k))$ with $\partial P / \partial \mu > 0$ and $\partial P / \partial \varepsilon < 0$ (Equation (2)), which supplies the economic microfoundation that makes the index estimable. The vulnerability weight $V(i)$ measures the differential incidence of a given cost increase on segment i , reflecting expenditure share and financial resilience. The gross-burden index $H(i)$ aggregates these as $H(i) = \sum_j \sum_k C(j) \cdot A(j,k) \cdot P(k) \cdot V(i)$ (Equation (1)), and the net-burden index $H_{\text{net}}(i)$ subtracts the parallel benefit terms $Q_b(k) \cdot V(i)$ defined in Section 7.5, so that the framework's central quantity is a net, distributionally weighted household burden rather than a gross cost. Each index carries an explicit unit and a candidate empirical data source in Table 1, which is what allows the qualitative framework to be reconstituted as an estimable model.

Table 1. Nomenclature: symbols, definitions, units, and possible data sources for the AI-NCPM.

Symbol	Definition	Unit	Possible Data Source
$H(i)$	Total household cost burden attributable to AI infrastructure for segment i	USD per household per year (index)	Derived (model output)
$H_{\text{net}}(i)$	Net household burden: transmitted cost minus transmitted benefit for segment i	USD per household per year (index)	Derived (model output)
i	Household segment (low-income, minority, rural, older adult, disability)	Category index	U.S. Census/ACS; BLS Consumer Expenditure Survey
j	Data center cost category (pre-construction, construction, operations, grid, lifecycle)	Category index	Industry cost reports; developer filings
k	Transmission channel (direct energy, digital service, sectoral, indirect)	Category index	Model definition
$C(j)$	Total cost generated in category j (Layer 1)	USD	Turner and Townsend/industry cost benchmarks; developer capex disclosures
$A(j,k)$	Share of cost category j allocated to channel k (Layer 2)	Dimensionless (0–1)	Input-output tables; sector cost-allocation studies
$P(k)$	Pass-through rate of channel k to households (Layer 3)	Dimensionless (0–1)	Utility rate filings; econometric pass-through estimates
$P_b(k)$	Buyer-side (household) pass-through rate on two-sided platforms	Dimensionless (0–1)	Platform pricing data; elasticity estimates
$Q_b(k)$	Buyer-side benefit pass-through rate in channel k	Dimensionless (0–1)	Productivity/price-deflation studies
$G(k)$	Gross efficiency gain (benefit) generated in channel k	USD	Sectoral productivity analyses
$V(i)$	Vulnerability weight (multiplier) for segment i (Layer 4)	Dimensionless (≥ 0)	BLS Consumer Expenditure Survey; affordability indices
$\mu(k)$	Degree of upstream market power in channel k	Dimensionless index	Market concentration data (e.g., HHI)
$\varepsilon(k)$	Price elasticity of household demand for the channel- k service	Dimensionless	Demand-elasticity literature
$\varepsilon_b, \varepsilon_s$	Buyer-side and seller-side price elasticities on a platform	Dimensionless	Platform demand studies
φ	Aggregate pass-through factor (share of Δc recovered from both sides)	Dimensionless (0–1)	Platform margin/competition analysis

Table 1. Cont.

Symbol	Definition	Unit	Possible Data Source
Δc	Incremental platform infrastructure cost per transaction	USD per transaction	Platform cost disclosures
$\Delta p_b, \Delta p_s$	Buyer-side and seller-side price increments set by the platform	USD per transaction	Platform pricing data

2.4. Evidence Base and Its Epistemic Status

The framework uses two distinct evidence types. Type A includes documented, citable magnitudes (data center construction and energy costs, utility rate-recovery filings, hyperscale capacity growth, and household expenditure shares) to show that cost origination and household vulnerability are real and non-trivial. Type B includes illustrative household pass-through figures used in worked applications to show how assumed costs would distribute across segments, rather than measured findings; these are flagged in Sections 5, 6 and 7.7. This distinction prevents the framework from being mistaken for empirical results and clarifies what future studies must estimate.

2.5. Validation Strategy and Falsifiability

Because each index has a defined unit and candidate data source, the framework also specifies its validation path. C(j) can be populated from industry capex benchmarks and developer disclosures; A(j,k) from input–output and sector cost-allocation studies; P(k) from utility rate filings and econometric pass-through estimates, with $\mu(k)$ and $\varepsilon(k)$ derived from market-concentration and demand-elasticity data; and V(i) from consumer-expenditure and access surveys. The framework is falsifiable in two ways: if P(k) in essential-service channels approaches zero, the model’s socialization mechanism would be contradicted; and if $Q_b(k) \cdot V(i)$ is as large for vulnerable segments as for affluent ones, the regressivity thesis would fail. A structured empirical agenda for each index is provided in the Limitations and future-work section.

Figure 1 illustrates the framework of the AI Infrastructure Net Cost Pass-Through Model (AI-NCPM), which is elaborated in the following sections. The authors used Claude (Anthropic) to assist with visualizing Figure 1. All content was reviewed and verified by the authors, who take full responsibility for the publication.




Figure 1. The AI Infrastructure Net Cost Pass-Through Model (AI-NCPM): Four-Layer Cascade Framework.

3. AI Growth and the Strategic Expansion of Data Centers

The past decade has witnessed an exponential surge in AI workloads, fueled by breakthroughs in machine learning algorithms, the availability of vast training datasets, and the proliferation of specialized hardware accelerators [9,10]. This AI boom is driving an unprecedented expansion of data center infrastructure to support the large-scale computational and storage requirements of AI applications [6].

The insatiable appetite for these workloads is evident in the global data center construction pipeline. Hyperscale facilities, operated by tech giants like Google, Microsoft, Amazon, and Meta, are being built at a frenetic pace to keep up with AI demand [5]. In the U.S. alone, over 150 hyperscale campuses are under development, with power requirements often exceeding 300 MW per site—enough to power a midsize city [7]. Google, for example, announced plans to invest \$10 billion in AI-focused data centers across the U.S. in 2024 [11]. Alongside hyperscalers, colocation providers and edge computing operators are also riding the AI wave. Alongside hyperscalers, colocation providers are seeing record demand from AI startups and enterprises [12], while edge data centers are proliferating to support low-latency applications [13]. U.S. data center construction is a complex undertaking [14]. Site selection is driven by power availability, fiber connectivity, tax incentives, and proximity to end-users [15]. Campuses often require dedicated substations and transmission lines; cooling systems are evolving for the high power densities of AI hardware [16]; and timelines are compressed to months [17]. This buildup is transforming data centers into ‘critical infrastructure’ essential to society and the economy [18], making their uninterrupted operation a national imperative [19].

Global data center IP traffic grew rapidly over the second half of the 2010s, with the Cisco Global Cloud Index projecting annual data center IP traffic to rise roughly threefold, from about 6.8 zettabytes per year in 2016 to approximately 20.6 zettabytes [20] per year by 2021 [19]. This sustained expansion, driven largely by the shift in workloads to cloud and hyperscale facilities, established the traffic base that later AI workloads would build upon. The global hyperscale footprint was already highly concentrated by 2020, when the number of large data centers operated by hyperscale providers reached roughly 540 facilities worldwide, having more than doubled over the [21] preceding five years [20]. The United States and China together accounted for over half of these sites, illustrating the geographic concentration of infrastructure that shapes where AI-related costs and externalities are ultimately localized [21]. The hyperscale footprint has expanded sharply since this 2020 snapshot. Synergy Research Group reports that the number of hyperscale data centers reached 1136 at the end of 2024 and 1360 by the end of 2025, roughly doubling from fewer than 600 in 2019, with generative-AI demand cited as the primary driver of both the rising facility count and their increasing average scale [22].

4. The Cost Anatomy of AI-Era Data Center Development

The cost anatomy of an AI-era data center spans its entire lifecycle, from land acquisition through construction, operation, grid integration, and eventual decommissioning. Rather than a single capital outlay, these costs form five distinct directions that arise at different stages and are borne by different actors, yet each ultimately contributes to the infrastructure burden that later propagates toward households. Table 2 summarizes these five cost directions, together with their principal cost drivers, representative cost figures drawn from cited industry, government, and academic sources, and the corresponding references. The figures indicate representative orders of magnitude rather than audited benchmarks, and the subsections that follow describe each direction in turn.

Table 2. Five cost directions in the anatomy of AI-era data center development, with principal cost drivers, representative cost figures, and supporting sources.

Cost Direction	Principal Cost Drivers	Representative Cost Figures	Sources
I. Pre-construction	Land acquisition; permitting and environmental compliance; grid-interconnection fees; community impact assessments.	Land \$50,000 to over \$5 million per acre; permitting and compliance add 5–10% of project budget; interconnection fees \$500,000 to over \$10 million; impact assessments upward of \$100,000.	[23–27]
II. Construction	Shell and core; power and cooling; specialized MEP and liquid cooling; N + 1/2N redundancy; materials and skilled labor.	\$1000 to over \$3000 per ft ² (vs. \$500–\$1000 for offices); power and cooling ~35% of build; MEP 25–30%; liquid cooling +10–15% of MEP; redundancy can 2–3× system cost; all-in AI builds \$30–40 million per MW.	[16,17,28–37]
III. Operational	Electricity (dominant); cooling energy and water; maintenance; staffing; cybersecurity.	Electricity ~60% of 2020 opex; 30–50 MW draw = \$20–40 million per year; cooling 30–50% of energy; up to 50 kW per rack; 3–5 million gallons water per day; maintenance \$8–10 per ft ² ; 50–100 FTE; cybersecurity \$2–3 million per year.	[7,16,29,38–45]
IV. Grid infrastructure	Transmission and distribution upgrades; dedicated substations; costs socialized across ratepayers; public subsidies.	Virginia: \$5–10 billion in T&D over a decade; residential bills +\$5–10 per month; Netherlands: €150 million per year in grid subsidies.	[46–50]
V. Lifecycle and decommissioning	Short (3–5 yr) refresh cycle; e-waste handling; retrofits and expansions; decommissioning and hazardous-material handling.	15–20 million tons e-waste per year; disposal \$500–1000 per ton; major retrofit 30–50% of a new build; decommissioning \$100–200 per ft ² .	[7,51–56]

5. Net Cost Pass-Through Mechanisms: How Infrastructure Costs and Benefits May Reach Households

Before tracing these channels, it is important to note that each originates in the physical data center and its supporting infrastructure. The costs passed to households reflect building- and infrastructure-scale decisions, including power-dense construction, MEP and cooling systems, grid interconnection and transmission upgrades, and material, water, lifecycle, and decommissioning demands. Thus, the four pass-through channels represent the economic footprint of the data center as a built asset, shaped by design, siting, and energy-infrastructure choices.

5.1. Direct and Indirect Cost Pass-Throughs

(I) Direct Energy Cost Pass-Through: Each of the four channels below transmits both costs and partially offsetting benefits; therefore, the household-relevant measure is the net effect, and a channel may be regressive on the cost side while benefits accrue elsewhere. The most direct pass-through occurs through electricity rates. As data centers consume a growing share of regional electricity (20% in Silicon Valley and 10% in Northern Virginia [46,50]) utilities invest in new generation, transmission, and distribution, recovering costs through general rate increases affecting all customers [57]. In Virginia, Dominion Energy estimates residential bills will rise by \$5–10 per month, or \$60–120 per year, over the next decade to fund data center grid upgrades, creating a significant burden for low- and fixed-income families [48]. Rising data center electricity demand is expected to place upward pressure on wholesale and, in some markets, retail electricity prices over the 2020–2030 period. However, the magnitude is uncertain and market-specific, depending on generation mix, grid investment, and rate design; therefore, this relationship is presented as an illustrative mechanism within the framework, not as an established empirical finding.

(II) Indirect Digital Service Cost Pass-Through: Data center costs also reach households through digital service pricing. Streaming, social media, cloud storage, and online banking services rely on cloud platforms operated by Amazon, Microsoft, and Google [58], and rising data center costs can increase the prices of these services [59]. In 2020, the U.S. had the highest data center electricity consumption at 76 TWh. Netflix, for example, has cited rising AWS costs as a factor in subscription price increases [60], while AI-as-a-Service costs are also reflected in household-facing applications [61]. Data center electricity consumption

is regionally concentrated, with a limited number of markets carrying a disproportionate share of global demand. In the IEA Base Case, the United States, China, and Europe dominate consumption, and this concentration intensifies over the 2020–2030 horizon as AI-optimized capacity expands [44]. By 2024, the United States accounted for about 45% of global data-center electricity consumption, followed by China at about 25% and Europe at about 15%, with U.S. demand projected to rise by roughly 130% by 2030 [44]. This concentration means that grid and affordability impacts are felt most strongly in specific local markets.

5.2. Household Sectoral Cost Pass-Through

(I) Sectoral Direct Cost Pass-Through: Two qualifications apply throughout this subsection. First, the causal links described here are AI-NCPM pathways, not established empirical findings; the extent to which AI-infrastructure costs are embedded in sectoral prices is what the framework is designed to estimate and has not yet been measured directly. Second, AI-infrastructure cost is only one driver of price change, alongside labor, regulation, demand, and general inflation. Thus, observed price increases in AI-adopting sectors are not attributed to data-center costs alone or predominantly. Terms such as ‘pass on,’ ‘reflect,’ and ‘raise’ describe potential transmission channels whose realized magnitude depends on the market-structure and elasticity factors formalized in Section 7.4.

Data center costs also reach households through services in AI-reliant sectors. In healthcare, AI is increasingly used for imaging analysis, diagnosis, drug discovery, and clinical decision support [62–64]; where these tools reduce costs or improve outcomes, benefits may offset infrastructure costs embedded in care. However, the data center infrastructure needed to train and deploy these models at scale may become a cost driver for providers and insurers, partly passed on to patients and policyholders [65,66]. Similar pathways apply in education, where AI-powered personalized learning, tutoring, and adaptive assessment increase demand for high-performance computing [67] and may affect tuition, technology fees, or digital-learning subscriptions [68]; in transportation, where autonomous vehicles and intelligent logistics require edge and AI infrastructure [69] and may influence service prices [70]; in retail, where AI supports personalized marketing, dynamic pricing, and inventory management [71,72]; and in banking, where AI for fraud detection, risk assessment, and robo-advising may affect financial-service fees [73,74].

Cloud spending has also risen sharply and unevenly across sectors, with finance, retail, technology, and other information-intensive industries accounting for large shares. These sectoral patterns indicate where infrastructure-driven cost pass-through may reach households first through digital-service prices. The shares discussed here are indicative, illustrating relative magnitudes rather than an audited market breakdown. Since 2020, sectoral cloud spending has grown steeply: Synergy Research Group reports worldwide cloud-infrastructure-services revenue of about \$321 billion in 2024, up roughly 20% from \$268 billion in 2023, and approximately \$419 billion in 2025, with generative AI identified as the principal growth driver across sectors [22].

(II) Sectoral Indirect Cost Pass-Through: AI infrastructure costs also reach households through indirect channels. AI adoption embeds data center costs in the prices of intermediate goods [75–77] and raw materials [74–76]; AI-optimized public services pass infrastructure costs on as higher taxes, fees, or red [78,79] vice quality [77,78]; and AI automation drives job displacement and wage stagnation, especially for low-wage workers [75,80].

6. Household Net Good Basket Perspective

A ‘net good basket’ perspective captures the total AI burden—transmitted cost minus transmitted benefit—across all categories of household expenditure [81]. The cost side falls into three categories: direct costs (smart home devices, AI subscriptions), indirect costs (embedded price increases, higher taxes), and embedded sectoral costs across healthcare, education, transportation, retail, finance, utilities, and telecommunications [82]. Against these run offsetting benefits (lower search and matching costs, faster service, and efficiency gains); so the basket is reshaped in net, not simply inflated. Because these benefits accrue disproportionately to affluent, digitally connected households while costs concentrate on essentials, the net basket is more regressive than the gross cost figures alone imply.

For many households, the AI cost burden is likely to be felt most acutely in the essential service sectors that form the backbone of daily life. These sectors—healthcare, education, transportation, utilities, and telecommunications—are not only undergoing rapid AI transformation but also represent a significant and often non-discretionary share of household expenditures [83]. In healthcare, AI costs may manifest as a \$150 increase in annual health insurance premiums to cover providers’ AI costs [84]. In education, a \$500 increase in annual college tuition to fund university AI infrastructure [85]. In transportation, a 10% increase in monthly transit pass prices to cover agency AI upgrade costs [86]. In utilities, a \$5 increase in the monthly water bill to fund AI leak detection systems [87]. In telecommunications, a \$10 increase in the monthly wireless bill to cover carrier AI infrastructure costs [88]. The specific figures in this paragraph (the \$150 premium, \$500 tuition, 10% transit increase, \$5 water charge, and \$10 wireless increase) are illustrative magnitudes used for conceptual exposition, not documented causal estimates of AI-attributable pass-through; isolating the data-center-cost component of any of these prices would require a dedicated empirical identification strategy.

Several indicators suggest that vulnerable populations are structurally exposed to the cost side of the AI-NCPM cascade. Bottom-quintile households spend 33% of expenditures on utilities and telecommunications, compared with 8% [89] for the top quintile [88]. Minority households face overlapping burdens: 34% of Black and 39% of Latino households in the U.S. lack home broadband, compared with 21% [90] for White households [89]. Rural households spend 40% more on electricity than [91] urban households [90]. Among older adults, 22% of those over 65 do not use the internet, compared with 6 [92] adults under 65 [91]. Households with disabilities spend 28% more on healthcare than those without [93]. These patterns indicate the need for targeted strategies to reduce burdens on households least able to absorb them. These indicators support the equity argument with important caution. They are general household-vulnerability measures (covering expenditure shares, broadband access, healthcare burden, and related factors) not direct measurements of AI-infrastructure cost exposure. Therefore, three claims must be distinguished. First, these groups spend larger income shares on essential services such as electricity, healthcare, and telecommunications, through which cost channels operate; this is directly supported by the cited expenditure evidence. Second, these groups often have lower access to AI-related benefits that could offset costs, especially where benefits concentrate in discretionary, digitally delivered services, as described in Benefit Channels B2 and the competitive portion of B3; this is supported by access and adoption evidence and the benefit-incidence logic in Section 7.5. Third, the stronger claim that these groups have already been shown to bear higher AI-infrastructure-specific costs is not asserted here. Isolating the AI-infrastructure component from other household price drivers requires micro-data that do not yet exist, the distributional conclusion rests on the first two empirically grounded claims and the framework’s structural logic, while the third remains a testable hypothesis that the AI-NCPM is designed to estimate.

7. Conceptual Framework: The AI Infrastructure Net Cost Pass-Through Model (AI-NCPM)

7.1. Theoretical Foundations and Rationale

The preceding sections established that AI data centers transmit both costs and benefits, which propagate through energy markets, digital platforms, sectoral supply chains, labor markets, and public finance until they reach households as a net change in the prices they pay. No integrated framework yet traces these pathways from origin to household impact in net terms; the AI Infrastructure Net Cost Pass-Through Model (AI-NCPM) fills that gap by tracing cost and benefit flows jointly and by treating the intermediating platforms as two-sided markets that split incidence between merchants and consumers.

The AI-NCPM draws on several established theoretical traditions. First, it builds on cost incidence and tax-shifting theory, which explains how costs imposed at one point in a supply chain are distributed among producers, intermediaries, and consumers according to relative elasticities [8,81]. For data centers, this means firms may absorb costs, pass them to business customers, or pass them to consumers, with forward pass-through to households being the least visible and least studied. Second, it draws on negative externalities theory [8], recognizing that data centers generate social costs (carbon emissions, water depletion, grid congestion, community disruption, and e-waste) that are often borne by households and communities through higher rates, environmental degradation, and strained public services. Third, it incorporates infrastructure cost socialization, which examines how large-scale infrastructure costs are allocated across ratepayers and taxpayers [46,48,57], especially when grid-upgrade benefits accrue mainly to a few large customers. Fourth, it uses the household ‘basket of goods’ framework [81] by treating AI infrastructure costs as a force that can inflate multiple household expense categories simultaneously, offering a more integrated view than channel-specific analyses. Finally, it draws on environmental justice and distributional equity studies, which show that vulnerable communities often bear disproportionate costs while receiving fewer benefits [75,80,90]. By embedding household segmentation and differential vulnerability, the AI-NCPM treats equity as a core feature rather than an afterthought.

7.2. Architecture of the AI-NCPM: Four-Layer Cascade Structure

The AI-NCPM (see Figure 1) is a four-layer cascade in which costs generated at the top layer (data center development and operation) flow downward through sectoral absorption, household transmission, and consumption impact. Each layer represents distinct actors and market mechanisms that allocate, absorb, amplify, or attenuate costs. The layers are connected by downward cost flows, upward feedback loops, and cross-cutting externality and equity dimensions.

Layer 1: Data Center Cost Generation: Layer 1 of the AI-NCPM identifies the origin of costs: the development, construction, and operation of AI-era data centers. These costs fall into five categories: (a) pre-construction costs, including land acquisition (\$50,000–\$5 million per acre), permitting, environmental compliance, grid interconnection fees (\$500,000–\$10 million), and community impact assessments; (b) construction costs, including structural and MEP systems (\$1000–\$3000 per square foot), specialized cooling infrastructure adding 10–15% to MEP budgets, redundancy systems such as $N + 1$ or $2N$, and labor with wage premiums of 20–30% in tight markets; (c) operational costs, dominated by electricity, which accounts for 60–70% of OpEx and \$20–40 million annually for a typical hyperscale facility, along with water consumption of 3–5 million gallons per day, maintenance of \$8–10 per square foot annually, staffing of 50–100 full-time workers, and cybersecurity costs of \$2–3 million annually; (d) grid infrastructure costs, including utility-scale upgrades to substations, transmission lines, and generation capacity,

reaching \$5–10 billion in Virginia alone over the next decade; and (e) lifecycle and decommissioning costs, including equipment refresh cycles every 3–5 years, e-waste disposal of \$500–1000 per ton, facility retrofits costing 30–50% of the new build cost, and site decommissioning of \$100–200 per square foot.

Layer 1 is not merely a cost inventory; it emphasizes escalation. These costs rise with AI compute demand, hardware obsolescence, material price increases such as copper up 50%, labor market tightening, and regulatory requirements, making the model's cost inputs moving targets that amplify downstream impacts. It also distinguishes internalized costs borne by operators from externalized costs imposed on third parties. While construction and operational costs are largely internalized in the short term, grid, environmental, community, and lifecycle costs are externalized to varying degrees, shaping cost flows into later layers. Layer 1 is, in effect, the built-environment layer of the model: it is where architectural and engineering design decisions (power and cooling density, redundancy tier ($N + 1$ versus $2N$), choice of air versus liquid cooling, structural and MEP specifications, and material intensity) together with siting and infrastructure-planning decisions (grid-adjacency, substation and transmission provisioning, water sourcing, and land use) determine the magnitude and composition of the costs that all subsequent layers transmit. This makes data center design and infrastructure planning not merely the backdrop to the analysis but its causal starting point: a facility engineered for higher efficiency, sited to minimize the grid-upgrade burden, or designed for adaptive reuse and lower-impact decommissioning enters the cascade with a smaller downstream cost footprint than one that is not.

Layer 2: Sectoral Cost Absorption and Allocation: The second layer is where data center costs are absorbed and allocated across sectors through four mechanisms. First, cloud service pricing: firms recover data center costs by charging businesses for cloud computing, storage, and AI-as-a-Service [58,61], so rising data center costs raise cloud prices for effectively every sector. Second, AI adoption intensity: high-adoption sectors (financial services (60%), healthcare (52%), retail (44%)) absorb proportionally more costs than low-adoption sectors [62], creating an uneven cost landscape. Third, supply chain embedding: AI-related data center costs become embedded in manufactured goods and shipping fees, so even low-adoption sectors absorb costs indirectly [76,77]. Fourth, utility rate socialization: data center grid costs are socialized across all ratepayers through general rate increases, reaching even sectors with no link to AI since the data center component is not itemized on bills. Layer 2 outputs a sectoral cost allocation matrix distributing the total burden across sectors by AI adoption, cloud usage, supply chain exposure, and rate socialization. This matrix evolves as adoption, pricing, and regulation change, and it mediates the relationship between cost generation (Layer 1) and household transmission (Layer 3).

Layer 3: Household Cost Transmission Channels: The third layer transmits sectoral costs to households through four simultaneous channels that differ in visibility, magnitude, timing, and distributional impact.

Two clarifications guide how these channels should be interpreted. First, transmission is conditional, not automatic. Higher data center construction or operating costs do not necessarily produce higher household prices; pass-through depends on sector market structure, contractual and regulatory pricing rules, household demand elasticity, competition, and whether firms absorb costs through margin compression. This conditionality is formalized in the channel pass-through rate $P(k)$, which rises with upstream market power and falls with downstream demand elasticity, approaching zero when competitive sectors face elastic demand and unity when concentrated sectors provide essential, inelastically

demanded services. Thus, the channels below describe potential pathways whose realized magnitude depends on structural conditions.

Second, the four channels operate through distinct mechanisms. Channel 1 is a regulated utility-rate effect, shaped by ratemaking and cost socialization rather than market pricing. Channels 2 and 3 are price effects in competitive or oligopolistic markets: Channel 2 directly affects platform-service prices paid by households, while Channel 3 indirectly embeds AI infrastructure costs into healthcare, education, transport, and other sectoral services. Channel 4 is a general-equilibrium effect operating through labor markets and public finance rather than a single product price. Keeping these mechanisms distinct is important because each respond to different policy levers: rate design for Channel 1, competition and platform regulation for Channels 2–3, and labor-market and fiscal policy for Channel 4.

Channel 1: Direct Energy Cost Pass-Through. The most direct and measurable channel: utility rate increases for data center grid upgrades raise all residential bills—\$5–10 per month (\$60–120 per year) in Virginia [48]. Visibility is moderate, and the impact is regressive since lower-income households spend more of their income on electricity [89,91]. Mechanism type: regulated utility-rate effect; here $P(k)$ is governed by the rate-making process and the socialization of grid-upgrade costs across the rate base, not by competitive price-setting.

Channel 2: Digital Service Pricing Pass-Through. As cloud providers raise prices and customers pass them on, households pay higher fees for streaming, online banking, e-commerce, and cloud storage [59,60]. Visibility is low (increases are attributed to ‘content costs’ rather than infrastructure) and magnitude is growing as digital dependence and embedded AI features expand. Mechanism type: direct household price effect on platform-mediated services; incidence here is the buyer-side share of a two-sided platform’s cost recovery.

Channel 3: Sectoral Service Pricing Pass-Through. When healthcare providers, educational institutions, transportation agencies, retailers, and financial institutions invest in AI infrastructure and pass costs to customers, households may face higher prices for sector-specific services, including insurance premiums (\$150/year increase), college tuition (\$500/year increase), transit passes (10% increase), wireless bills (\$10/month increase), and retail prices [63–74,84–88]. This channel has broad reach because it affects multiple household expenditure categories simultaneously, creating the “goods basket inflation” effect central to the AI-NCPM. However, its visibility is low because households rarely connect increases in healthcare, education, or transportation costs to AI infrastructure. The figures in parentheses are illustrative of direction and rough scale, not measured pass-through estimates. Mechanism type: indirect producer-cost effect, where sectoral firms embed AI infrastructure costs in household prices. For two-sided platforms, such as e-commerce marketplaces connecting households and merchants, the household is the buyer side and the merchant is the seller side; therefore, the household burden corresponds to the buyer-side incidence $P_b(k)$. The more elastic the seller side, the larger the share borne by households, while some costs may be absorbed by merchants rather than passed on.

Channel 4: Indirect and Systemic Pass-Through. The diffuse second-order effects on household well-being include labor market impacts (displacement, wage stagnation), public finance impacts (higher taxes and fees), supply chain price effects, and macroeconomic effects [75,79,80]. This channel has the lowest visibility and longest lag, but its cumulative magnitude may be substantial. Mechanism type: general-equilibrium effect operating through labor markets and public finance rather than through a single product price; its incidence is diffuse and is not captured by a sector-specific pass-through rate.

The four channels are parallel and additive: total household burden is their sum. No single channel may seem significant (\$5 more for electricity, \$2 for streaming, a few dollars per doctor visit) but the cumulative effect across all categories can be substantial, especially for households with limited financial buffers.

Layer 4: Household Consumption and Well-Being Outcomes: Low-income households (already spending 33% of expenditures on utilities and telecommunications versus 8% for the top quintile [89]) are especially vulnerable, as even small increases meaningfully erode resilience. The AI-NCPM recognizes four second-order effects: reduced access to essential services (forgoing healthcare, education, or broadband, worsening outcomes); increased financial stress and debt (reduced savings and greater fragility); reduced discretionary spending (cutting entertainment, dining, and recreation); and geographic and social stratification (higher utility and housing costs near data center clusters displacing lower-income households). Layer 4 outcomes differ systematically across segments; the same segment-specific vulnerabilities (minority broadband gaps, higher rural electricity burdens, lower internet use among older adults, and elevated healthcare spending among households with disabilities) shape Layer 4 outcomes. These vulnerabilities are multiplicative, not additive: a low-income, rural, minority household with an elderly or disabled member faces compounding pressures far more severe than any single indicator suggests.

7.3. Feedback Loops, Cross-Layer Interactions, and Dynamic Properties

Unlike simpler cost-flow models, the AI-NCPM incorporates feedback loops and cross-layer dynamics. Three principal feedback mechanisms can amplify, attenuate, or redirect cost flows over time.

Feedback Loop 1: Demand-Driven Cost Escalation. Wider AI adoption (Layer 3/4) increases demand for data center capacity, driving further investment (Layer 1). This positive loop means AI's success generates additional costs passed back to the same users; an 'AI infrastructure cost spiral' that self-reinforces as long as capacity lags demand.

Feedback Loop 2: Affordability-Driven Demand Reduction. As costs raise essential-service prices, some households reduce or forgo AI-powered services. This negative loop attenuates aggregate demand but creates a 'digital divide' (wealthier households keep benefiting while lower-income ones are priced out) worsening distributional inequities even as it moderates aggregate pressure.

Feedback Loop 3: Policy and Innovation Response. As the burden becomes politically salient, policymakers and firms may intervene via efficiency mandates, renewable incentives, rate reforms, targeted subsidies, community benefit agreements, and AI efficiency innovations—at any layer, significantly altering cost flows. Though politically contingent, this loop is the key mechanism for redirecting AI-NCPM dynamics toward more equitable, sustainable outcomes.

The AI-NCPM also recognizes cross-layer interactions: operator competition at Layer 1 can reduce costs flowing into Layer 2, while cloud market concentration at Layer 2 can amplify transmission to Layers 3–4, and state or national regulation can alter Layer 1 cost structures with cascading effects. These interactions make the AI-NCPM a dynamic, adaptive system rather than a static accounting framework.

7.4. The AI-NCPM's Index Model

As established in the Introduction and formalized as the study's research design in Section 2, the AI-NCPM is a conceptual framework rather than an empirically estimated model. This does not preclude a formal representation: the framework can be expressed analytically to facilitate future empirical testing and computational modeling, and the index

model below specifies that representation. At the most general level, the total household cost burden attributable to AI infrastructure (H) can be expressed as a function of the data:

$$H(i) = \sum_j \sum_k C(j) \times A(j,k) \times P(k) \times V(i) \quad (1)$$

where i indexes household segments (low-income, minority, rural, older adult, disability), j indexes data center cost categories (pre-construction, construction, operations, grid, lifecycle), and k indexes transmission channels (direct energy, digital service, sectoral, indirect). $C(j)$ represents the total cost generated in category j at Layer 1. $A(j,k)$ represents the share of cost category j that is allocated to transmission channel k at Layer 2. $P(k)$ represents the pass-through rate of channel k —the fraction of the allocated cost that is actually transmitted to households at Layer 3. Crucially, $P(k)$ is not a free parameter but is itself determined by the market structure of the transmitting sector and the price elasticity of household demand for the affected service. Following the classical theory of cost incidence, the pass-through rate rises with the concentration (market power) of the upstream provider and falls as downstream demand becomes more elastic, so that the pass-through rate is a function of upstream market power and downstream demand elasticity, where $\mu(k)$ denotes the degree of upstream market power in channel k and $\varepsilon(k)$ denotes the price elasticity of household demand for the channel- k service, with $\partial P / \partial \mu > 0$ and $\partial P / \partial \varepsilon < 0$ (Equation (2)). In perfectly competitive transmitting markets facing highly elastic demand, $P(k)$ approaches zero (costs are absorbed as margin compression rather than passed on); in concentrated markets facing inelastic demand for essential services, $P(k)$ approaches unity. This structural parameterization bounds $P(k)$ within the unit interval ($0 \leq P(k) \leq 1$) and supplies the economic microfoundation for its empirical estimation. $V(i)$ represents the vulnerability weight of household segment i at Layer 4—a multiplier that captures the differential impact of a given cost increase on different household types, reflecting both their expenditure shares on affected services and their financial resilience. The full set of symbols, their definitions, units, and possible empirical data sources is summarized in Table 2.

$$P(k) = f(\mu(k), \varepsilon(k)), \partial P / \partial \mu > 0, \partial P / \partial \varepsilon < 0, 0 \leq P(k) \leq 1 \quad (2)$$

This index model highlights several important properties of the AI-NCPM. First, the total household burden is the product of four distinct parameters, each of which is separately estimable and subject to different determinants. This decomposition enables analysts to identify which components of the cost chain are contributing most to household impacts and where interventions would be most effective. Second, the vulnerability weight $V(i)$ operates as a multiplier on the entire cost chain, meaning that even modest cost flows can produce significant impacts on highly vulnerable household segments. Third, the summation across cost categories (j) and channels (k) captures the ‘good basket’ effect; the cumulative, multi-channel nature of the AI infrastructure cost burden that is the model’s distinctive analytical contribution.

It is important to note that the parameters of the formal AI-NCPM model are not currently calibrated with empirical data. The pass-through rates $P(k)$, in particular, are difficult to estimate because data center costs are typically bundled with other costs in cloud service pricing, utility rate structures, and sectoral service prices, making it challenging to isolate the data center-specific component. Similarly, the vulnerability weights $V(i)$ require detailed micro-data on household expenditure patterns, income levels, and access to alternative services that are not currently available in a form that is specific to AI infrastructure cost impacts. The calibration of these parameters represents a priority for future empirical research.

7.5. Two-Sided Platform Cost Allocation and Net Pass-Through

The cascade architecture models hyperscalers and cloud-dependent platforms primarily as upstream suppliers in a vertical value chain, with costs flowing downward toward final consumers. This framing is analytically convenient but understates an important feature of the entities that dominate the AI infrastructure economy. The largest hyperscalers and e-commerce platforms (Amazon, Microsoft, Google, Alibaba, and their peers) are more accurately characterized as two-sided intermediaries that simultaneously serve buyers (consumers and businesses) and sellers (merchants and service providers), with both groups positioned downstream of the platform's matching and intermediation services. Recognizing this structure refines, rather than displaces, the AI-NCPM: it specifies more precisely how the sectoral allocation matrix $A(j,k)$ and the pass-through rate $P(k)$ are determined when the transmitting entity is a platform that prices to two interdependent sides of a market.

A profit-maximizing platform that experiences a rise in AI infrastructure cost does not mechanically load that cost onto final consumers. Instead, it allocates the cost increment across its two sides according to the relative price elasticities of demand on each side. Let a platform incur an incremental infrastructure cost per transaction of Δc , and let it set a buyer-side price increment Δp_b and a seller-side price increment Δp_s . Writing ε_b for the (absolute) price elasticity of buyer demand and ε_s for the (absolute) price elasticity of seller demand, the first-order conditions of the platform's two-sided pricing problem yield an inverse-elasticity allocation rule (This allocation rule abstracts from cross-side network externalities. In the canonical two-sided markets framework the optimal price on each side reflects not only that side's own-price elasticity but also the marginal benefit its participation confers on the other side. The inverse-elasticity form used here is the special case in which these cross-side effects are symmetric across the two sides or are second-order relative to the cost-allocation margin; it is adopted for tractability, leaving an extension with asymmetric network effects to future work.):

$$\Delta p_b / \Delta p_s = \varepsilon_s / \varepsilon_b, \text{ with } \Delta p_b + \Delta p_s = \varphi \cdot \Delta c, 0 \leq \varphi \leq 1 \quad (3)$$

The cost increment is borne more heavily by whichever side of the market is less elastic—less able to substitute away from the platform. Equivalently, because the allocation ratio $\Delta p_b / \Delta p_s$ rises with ε_s , a higher seller elasticity shifts the burden toward buyers: the more readily sellers can substitute away from the platform, the less of the cost the platform can recover from them, leaving the relatively less-elastic buyers to bear the larger share. The aggregate pass-through factor φ (the share of Δc recovered from the two sides combined rather than absorbed in platform margin) is itself bounded by the intensity of inter-platform competition: a hyperscaler or marketplace facing strong rivals (Amazon competing with Walmart, Shopify-enabled storefronts, Alibaba/Alibaba.com, and Mercado Libre across both its B2C and B2B sides) operates with φ well below one, because attempts to recover the full cost from either side risk driving buyers or merchants to competing platforms. φ is treated here as a reduced-form parameter that captures this competitive constraint in tractable form, rather than deriving it from the same profit-maximization that yields the allocation rule; it is bounded by, but not solved as an explicit function of, the intensity of inter-platform rivalry. Endogenizing φ as an equilibrium outcome of platform competition—for instance, in a model with platform differentiation and seller or buyer multi-homing; is a natural extension left to future work. This delivers two results that sharpen the AI-NCPM's policy conclusions. First, the consumer-side pass-through that the vertical cascade implicitly assumes is an upper bound; the realized burden on final consumers, Δp_b , is strictly smaller whenever sellers' demand is not perfectly elastic

relative to buyers'. Second, a meaningful fraction of AI infrastructure cost is allocated to merchants (small and mid-sized sellers who depend on dominant marketplaces) rather than to households directly, redistributing the incidence in a way the original framework did not capture.

Embedding this two-sided structure in the index model requires replacing the single channel pass-through rate $P(k)$ with a side-specific decomposition. For platform-mediated channels, the household-relevant pass-through $P_b(k)$ is given by Equation (4), the product of the competition-bounded aggregate recovery factor and the buyer side's elasticity-determined share of that recovery. The household burden expression generalizes accordingly to $H(i) = \sum_j \sum_k C(j) \cdot A(j,k) \cdot P_b(k) \cdot V(i)$. When demand on the two sides is symmetric ($\varepsilon_b = \varepsilon_s$) and competition is absent ($\varphi = 1$), the expression collapses to the original single-sided form, confirming that the vertical cascade is the special case of the more general two-sided model in which platform competition is ignored and the two sides are equally substitutable.

$$P_b(k) = \varphi(k) \cdot [\varepsilon_s(k) / (\varepsilon_b(k) + \varepsilon_s(k))] \quad (4)$$

A complete welfare account must also recognize that AI infrastructure investment generates downstream productivity and welfare benefits, not costs alone. The same hyperscale capacity that raises infrastructure expenditure also lowers the marginal cost of the services it supports: faster logistics and inventory optimization in retail, AI-assisted diagnostics and drug discovery in healthcare, and algorithmic personalization that reduces search and matching costs in e-commerce. A portion of these gains is transmitted downstream as benefit pass-through, partially offsetting the cost pass-through traced through the four channels. The net household impact in channel k is therefore the difference between transmitted cost and transmitted benefit, and the index model can be written in net form as Equation (5), where $G(k)$ is the gross efficiency gain generated in channel k and $Q_b(k)$ is the buyer-side benefit pass-through rate, governed by the same market-structure and elasticity logic that governs $P_b(k)$.

$$H_{\text{net}}(i) = \sum_j \sum_k [C(j) \cdot A(j,k) \cdot P_b(k) - G(k) \cdot Q_b(k)] \cdot (i) \quad (5)$$

To give the benefit side the same structure as the cost side, a typology of benefit pass-through channels is specified that runs parallel to the four cost channels of Layer 3. Each benefit channel has a source (the efficiency or welfare gain AI generates), a transmission condition (whether and how the gain is passed through to households, governed by the same market-structure and elasticity logic as $P_b(k)$; namely the buyer-side benefit pass-through rate $Q_b(k)$), and an incidence profile (which households actually receive it). The central asymmetry is that benefit pass-through is generally higher for discretionary, competitively supplied, digitally delivered services (where competition forces firms to pass efficiency gains to consumers and where take-up requires digital access) while cost pass-through is highest for non-discretionary essentials supplied under regulation or market power. This is why netting benefits against costs does not equalize the distribution: the two sides pass through under different conditions to different households.

Benefit Channel B1: Energy-system efficiency benefit (parallel to Channel 1). Source: efficiency gains in energy provision itself; AI-optimized grid dispatch, demand forecasting, and data-center load flexibility that can lower system-wide generation and balancing costs. Transmission condition: low and slow. Because retail electricity is a regulated, near-monopoly service, any efficiency gain reaches households only if regulators require it to be reflected in rates; absent that, it is retained as utility margin or offsets the data center-driven rate increase only partially. Incidence: broadly shared if passed through at all, since it lands

on the regulated rate every ratepayer faces; but Q_b here is typically small, so B1 rarely offsets much of Channel 1's cost.

Benefit Channel B2: Direct digital-service benefit (parallel to Channel 2). Source: direct consumer-facing gains from AI-enhanced digital services—lower search and matching costs, free or low-cost AI features (assistants, recommendations, translation), and quality improvements in streaming, navigation, and online banking. Transmission condition: high. These services are supplied competitively and often priced at zero to the consumer, so efficiency gains and new capabilities are passed through quickly and visibly; $Q_b(k)$ is large. Incidence: skewed toward digitally connected, higher-income households that adopt these services (precisely the households least exposed to Channel 1–Channel 3 essential-cost pass-through) so B2 offsets cost most for those who bear it least.

Benefit Channel B3: Sectoral efficiency benefit (parallel to Channel 3). Source: sectoral efficiency gains where AI lowers the cost or raises the quality of essential services—AI-assisted diagnostics and drug discovery in healthcare, intelligent tutoring and personalization in education, routing and logistics optimization in transportation and retail, and fraud reduction and automated underwriting in financial services. Transmission condition: partial and competition-dependent. In competitive segments (retail logistics, some financial services) part of the gain is passed to households as lower prices or fees; in concentrated or third-party-payer segments (much of U.S. healthcare and higher education) providers tend to retain the gain, so $Q_b(k)$ is low even though the underlying efficiency gain may be large. Incidence: uneven across sectors and households, and often captured by insurers, institutions, or intermediaries before reaching the household; so B3 offsets Channel 3 costs only where the sector is genuinely competitive and the household pays directly.

Benefit Channel B4: General-equilibrium productivity benefit (parallel to Channel 4). Source: economy-wide productivity gains (higher aggregate output, new goods and services, and general price deflation in AI-intensive industries) together with any wage gains for workers complementary to AI. Transmission condition: diffuse and lagged, operating through general-equilibrium adjustment rather than any single price, and partly offset by labor-displacement losses for workers who are substitutes for AI. Incidence: highly unequal—accruing to capital owners, high-skill workers, and consumers of AI-intensive goods, while workers displaced by automation may experience net losses—so B4's household incidence can be regressive in its own right, compounding rather than offsetting the cost-side regressivity.

Taken together, this typology makes the 'net' in net cost precise: for a vulnerable household, the cost terms $P_b(k) \cdot V(i)$ are concentrated in Channels 1 and 3 (regulated essentials and sectoral services), while the benefit terms $Q_b(k) \cdot V(i)$ are concentrated in Channel 2 and the competitive part of Channel 3—channels such a household under-consumes. The benefit and cost typologies are therefore not mirror images that cancel: they are offset in structure, which is exactly why the net good basket remains regressive.

Critically, this balanced treatment reinforces rather than weakens the study's equity thesis. Cost pass-through and benefit pass-through are distributed across household segments by different mechanisms. Benefit pass-through accrues disproportionately to higher-income, digitally connected households that consume the personalized, AI-enhanced services in which efficiency gains are concentrated, whereas cost pass-through—operating through non-discretionary essentials such as electricity, healthcare, and telecommunications—falls disproportionately on the low-income, minority, rural, older adult, and disability-affected segments identified in Section 6. The vulnerability weight $V(i)$ thus enters the cost term and the benefit term asymmetrically: for vulnerable segments, $P_b(k) \cdot V(i)$ is large while $Q_b(k) \cdot V(i)$ is small, so the net burden $H_{net}(i)$ remains regressive even after benefits are netted out. Incorporating benefit pass-through therefore makes the

welfare analysis more credible to referees while leaving the distributional conclusion—that the AI-inflated good basket bears most heavily on those least able to absorb it—intact and, if anything, more sharply drawn.

7.6. Comparative Positioning: The AI-NCPM in Relation to Existing Frameworks

The AI-NCPM occupies a distinct position among frameworks used to analyze digital infrastructure and AI impacts. Existing data center cost frameworks mainly include engineering-economic CapEx models, Power Usage Effectiveness (PUE), and total cost of ownership (TCO) calculators used by operators and investors to assess project [6,30,41] viability [6,29,40]. These tools support internal decision-making but remain firm-centric and do not explain how data center costs are distributed across the economy or ultimately borne by households. At the macro level, IEA and national energy-system models include data center electricity consumption within overall energy demand [7,19]. While useful for assessing aggregate energy and environmental impacts, they do not trace household-level economic transmission pathways or account for non-energy costs such as construction, water, or community impacts.

The AI-NCPM differs in three key ways. First, it is household-centric, placing impacts on household consumption and well-being at the center. Second, it is multi-channel and multi-sectoral, capturing simultaneous cost transmission across sectors through the ‘goods basket’ perspective. Third, it is equity-oriented, embedding household segmentation and vulnerability analysis as core features rather than optional extensions. Together, these features—household-centricity, multi-channel scope, and equity orientation—define the AI-NCPM’s unique contribution to literature, policy, and practice.

7.7. Applications of the AI-NCPM

To demonstrate the practical utility of the AI-NCPM, this subsection presents three application scenarios that highlight how the framework can be used to inform policy analysis and decision-making.

Application 1: Regional Energy Cost Impact Assessment. Consider a utility service territory, such as Northern Virginia or the San Francisco Bay Area, where rapid data center electricity demand is projected to raise residential electricity rates by 8–10% over the next decade. Using the AI-NCPM, an analyst can trace this cost from Layer 1, data center energy consumption and grid upgrades, through Layer 2, utility rate socialization across customer classes, through Layer 3, direct energy cost pass-through, and finally to Layer 4, differential household impacts. Applying segment-specific vulnerability weights shows that the same 10% rate increase may create a \$72 annual burden for a top-quintile household spending 8% of income on utilities, but a \$396 annual burden for a bottom-quintile household spending 33% of income on utilities—a more-than-fivefold difference in relative impact. This analysis can support targeted rate relief, energy efficiency incentives, or community benefit agreements to reduce regressive effects from data center-driven rate increases. These dollar figures are worked illustrations of the model’s arithmetic under an assumed rate increase, not empirical findings; they show how a given pass-through magnitude would be distributed across household segments once the framework is calibrated.

Application 2: Sectoral AI Cost Transparency Audit. The AI-NCPM can be used to conduct a ‘cost transparency audit’ that estimates the share of household expenditures in a given sector (e.g., healthcare) that is attributable to AI and data center infrastructure costs. By mapping the AI adoption intensity of major healthcare providers, the share of their IT budgets allocated to cloud and AI services, and the pass-through mechanisms by which these costs reach patients and policyholders, an analyst can estimate that AI infrastructure costs now account for an estimated 2–4% of total U.S. healthcare spending; a

figure that is expected to grow as AI adoption in healthcare accelerates. This transparency audit can inform regulatory requirements for cost disclosure, provider negotiations with cloud vendors, and consumer advocacy around healthcare affordability. This 2–4% range is an illustrative order-of-magnitude used to demonstrate the audit’s logic, not a documented estimate of AI-attributable healthcare cost.

Application 3: Equity Impact Modeling for Policy Design. The AI-NCPM can be used to model the equity impacts of alternative policy interventions across different household segments. For example, a policymaker considering a universal broadband subsidy to mitigate the digital divide can use the AI-NCPM to estimate the degree to which the subsidy would offset AI-related cost increases for minority, rural, and older adult households, and whether the subsidy would be sufficient to close the affordability gap created by the cumulative pass-through of data center costs across all four channels. By comparing the equity impacts of different policy options—such as income-based rate design, targeted technology access subsidies, or data center developer contribution requirements—the AI-NCPM can help identify the interventions that would most effectively reduce distributional inequities.

8. Limitations

Several limitations should be acknowledged. First, the AI-NCPM is a conceptual framework that has not yet been empirically calibrated. The channel-specific pass-through rates $P(k)$, two-sided decomposition $P_b(k)$ and $Q_b(k)$, benefit-transmission terms, and household vulnerability weights $V(i)$ are specified analytically rather than estimated from data. This reflects the difficulty of isolating data center costs within cloud pricing, utility rates, and sectoral service prices, as well as the lack of household-level micro-data on AI infrastructure cost impacts. Therefore, numerical values in the illustrative applications should be read as representative magnitudes, not calibrated estimates. Second, the aggregate pass-through factor ϕ is treated as a reduced-form parameter capturing competitive constraints on cost recovery, rather than as an equilibrium outcome of platform competition. Third, the analysis relies on heterogeneous secondary data from industry, government, and academic sources, so reported figures illustrate structure and orders of magnitude rather than primary measurements. Fourth, the framework is grounded mainly in the U.S. market and regulatory context and would require adaptation for other national settings. Fifth, the model is largely point-in-time and does not formally solve long-run feedback dynamics among data center demand, affordability, and policy response. Finally, broad household segments may mask within-group and intersectional differences. These limitations define the scope of the study while motivating the empirical research agenda that follows.

9. Conclusions

This study examined how the escalating costs of AI-era data centers (pre-construction, construction, operations, grid, and lifecycle) are transmitted through service supply chains to the prices households pay, and how a share of the productivity and welfare benefits AI generates is transmitted back through the same channels. These costs escalate continuously with AI compute demand, material price surges, labor tightening, and hardware obsolescence, with a significant share externalized to households rather than internalized by the firms that generate them, so that the household-relevant quantity is net cost—transmitted cost minus transmitted benefit. Four channels were identified: direct energy socialization via utility rates, digital service pricing in cloud platforms, sectoral pass-through from AI-adopting industries, and indirect transmission via supply chains, labor disruption, and public expenditures. Each carries both cost and benefit, and because the intermediating platforms are two-sided, costs are split between merchants and consumers by relative

elasticity rather than loaded onto households. Operating simultaneously and additively, they produce the ‘AI-inflated net good basket’; a systemic reshaping of the consumption basket’s net cost, where benefits accruing to affluent households leave the residual burden concentrated on the most vulnerable.

To formalize these dynamics, the study introduced the AI-NCPM, a four-layer cascade framework tracing cost flows from data center investment (Layer 1) through sectoral absorption (Layer 2), household transmission channels (Layer 3), and consumption and well-being outcomes (Layer 4). The AI-NCPM incorporates three feedback loops—demand-driven cost escalation, affordability-driven demand reduction, and policy/innovation response—alongside cross-cutting externality and equity dimensions. Its index model decomposes the household burden into separately estimable components of cost generation, sectoral allocation, pass-through rates, and household vulnerability weights, providing a structure for future empirical calibration. The equity analysis revealed that the cost burden falls disproportionately on vulnerable populations. Low-income households spend 33% of expenditures on utilities and telecommunications versus 8% for top-quintile households, and the minority, rural, older-adult, and disability-related vulnerabilities compound this exposure. These vulnerabilities are multiplicative, not additive; households at the intersection of multiple categories face compounding cost pressures.

The AI-NCPM remains a conceptual framework whose pass-through rates and vulnerability weights require empirical calibration through household expenditure surveys and utility rate analyses. The model accommodates both cost and benefit pass-through through a net-burden formulation, and treats platform-mediated channels as two-sided markets in which infrastructure costs are allocated across buyers and sellers by their relative demand elasticities; the empirical estimation of these side-specific pass-through rates and benefit-transmission parameters remains an open task. The analysis is grounded primarily in the U.S. context and would require adaptation for other national settings. Future research priorities include empirical estimation of channel-specific pass-through rates and their structural determinants (upstream market power and downstream elasticity), estimation of buyer-versus-seller cost incidence on two-sided platforms, measurement of benefit pass-through alongside cost pass-through, longitudinal household impact studies, comparative case studies across regions with varying data center concentrations, computational modeling of the AI-NCPM’s feedback dynamics, and participatory research engaging affected communities in equitable AI governance design.

Finally, the findings have direct implications for the built environment, where key levers shaping household burden reside. Because the cascade begins in Layer 1 (the physical facility and supporting infrastructure) decisions by developers, designers, utilities, and planners determine the scale of downstream costs. Design strategies that reduce energy and water intensity, such as efficient cooling, heat reuse, and higher-temperature operation; siting strategies that limit costly grid upgrades, such as locating near existing generation and transmission or pairing facilities with on-site or dedicated clean generation; and lifecycle strategies, including durable design, equipment refurbishment, adaptive reuse, and responsible decommissioning, can reduce the costs entering the pass-through chain. The externalities identified by the model (carbon emissions, water depletion, grid congestion, and e-waste) are also building- and infrastructure-scale issues requiring design, engineering, and planning responses. Linking AI-era data centers to lifecycle assessment, energy-infrastructure planning, and construction cost management therefore connects the model’s economic conclusions to the built-environment agenda.

Several directions can extend the AI-NCPM framework. First, the four-layer cascade and Hnet(i) equation require empirical calibration, especially the side-specific pass-through terms (P_b , Q_b) and benefit-offset term $G(k)$, using household-level price and

welfare data across sectors and regions with different AI adoption intensity. Second, the two-sided platform allocation rule, $\Delta p_b / \Delta p_s = \varepsilon_s / \varepsilon_b$, should be econometrically tested to determine how costs split between merchants and consumers under varying inter-platform competition ϕ . Third, a dynamic, time-indexed extension is needed because costs such as construction, grid upgrades, and rate socialization often occur earlier than productivity and price-deflation benefits. Fourth, the vulnerability multiplier $V(i)$ could be disaggregated to capture intersecting disadvantages, such as low-income rural households with limited digital access. Finally, the framework's feedback loops (demand-driven cost escalation, affordability-driven demand contraction, and policy/innovation-driven moderation) could be formalized through systems-dynamics or agent-based modeling to simulate long-run equilibria and identify equity-focused policy levers such as rate-design reform, targeted subsidies, and efficiency mandates.

In summary, the economics of artificial intelligence cannot be fully understood without tracing its infrastructure costs from data centers to household budgets. The AI-NCPM provides the first systematic framework for this analysis, revealing that the AI revolution's cost footprint extends far beyond the technology sector into the daily expenditures of families and communities. As policymakers and industry leaders confront the distributional consequences of AI-driven infrastructure expansion, the AI-NCPM offers an analytical foundation for ensuring that the benefits of this transformative technology are shared equitably and its costs are borne justly.

Author Contributions: Conceptualization, A.S., D.H. and A.K.; methodology, A.S., M.K., and A.K.; software, A.K.; validation, M.K. and A.K.; formal analysis, A.S.; resources, M.K., D.H. and A.K.; data curation, A.S. and M.K.; writing—original draft preparation, A.S. and A.K.; writing—review and editing, A.S. and M.K.; visualization, A.S., M.K. and A.K.; supervision, M.K., D.H. and A.K. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Acknowledgments: During the preparation of this manuscript, the authors used Claude (Anthropic, San Francisco, CA, USA) to assist with the production of Figure 1 and reference formatting. The authors reviewed and verified Figure 1 and formatted references and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. Language Models Are Few-Shot Learners. *Adv. Neural Inf. Process. Syst.* **2020**, *33*, 1877–1901.
2. Mirhoseini, A.; Goldie, A.; Yazgan, M.; Jiang, J.; Songhori, E.; Wang, S.; Lee, Y.J.; Johnson, E.; Pathak, O.; Bae, S.; et al. Chip Placement with Deep Reinforcement Learning. *Nature* **2021**, *594*, 207–212. [[CrossRef](#)] [[PubMed](#)]
3. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, Ł.; Polosukhin, I. Attention Is All You Need. *Adv. Neural Inf. Process. Syst.* **2017**, *30*, 5998–6008.
4. Avgerinou, M.; Bertoldi, P.; Castellazzi, L. Trends in Data Centre Energy Consumption under the European Code of Conduct for Data Centre Energy Efficiency. *Energies* **2017**, *10*, 1470. [[CrossRef](#)]
5. Shehabi, A.; Smith, S.J.; Sartor, D.A.; Brown, R.; Herrlin, M.; Koomey, J.; Masanet, E.; Horner, N.; Azevedo, I.; Lintner, W. *United States Data Center Energy Usage Report*; Lawrence Berkeley National Laboratory: Berkeley, CA, USA, 2024.
6. Masanet, E.; Shehabi, A.; Lei, N.; Smith, S.; Koomey, J. Recalibrating Global Data Center Energy-Use Estimates. *Science* **2020**, *367*, 984–986. [[CrossRef](#)] [[PubMed](#)]
7. International Energy Agency (IEA). *Data Centres and Energy—From Global Headlines to Local Headaches?* IEA: Paris, France, 2024.
8. Pigou, A.C. *The Economics of Welfare*; Macmillan and Company: London, UK, 1924.

9. Jouppi, N.P.; Young, C.; Patil, N.; Patterson, D.; Agrawal, G.; Bajwa, R.; Bates, S.; Bhatia, S.; Boden, N.; Borchers, A.; et al. In-Datcenter Performance Analysis of a Tensor Processing Unit. In Proceedings of the 44th Annual International Symposium on Computer Architecture, Toronto, ON, Canada, 24–28 June 2017; pp. 1–12.
10. Sze, V.; Chen, Y.H.; Yang, T.J.; Emer, J.S. Efficient Processing of Deep Neural Networks: A Tutorial and Survey. *Proc. IEEE* **2017**, *105*, 2295–2329. [CrossRef]
11. Brandt, E. Google to Invest \$10 Billion in U.S. Data Centers and Offices This Year. *CNBC*. 2025. Available online: <https://www.cnbc.com/amp/2020/02/26/google-invests-10-billion-in-us-amid-government-antitrust-probes.html> (accessed on 13 July 2026).
12. Heires, K. The rise of artificial intelligence. *Risk Manag.* **2015**, *62*, 38–43.
13. George, A.S.; George, A.H.; Baskar, T. Edge computing and the future of cloud computing: A survey of industry perspectives and predictions. *Partn. Univers. Int. Res. J.* **2023**, *2*, 19–44.
14. Shafaghat, A.; Hu, D.; Keyvanfar, A. A Fuzzy-Set Qualitative Comparative Analysis (fsQCA) of Dynamic Capabilities and Microfoundations for Hyperscale Data Center Project Delivery. *Buildings* **2026**, *16*, 2284. [CrossRef]
15. Goiri, Í.; Le, K.; Guitart, J.; Torres, J.; Bianchini, R. Intelligent Placement of Datacenters for Internet Services. In Proceedings of the 31st ICDCS, Minneapolis, MN, USA, 20–24 June 2011; pp. 131–142.
16. Zhang, X.; Wang, Y.; Liu, Z. A Review of Liquid Cooling Technologies for Data Centers. *Appl. Therm. Eng.* **2025**, *214*, 118979.
17. Sankar, T.; Venkata Ramana, R.B.; Balamuralikrishnan, A. AI-Optimized Hyperscale Data Centers: Meeting the Rising Demands of Generative AI Workloads. *Int. J. Trend Sci. Res. Dev.* **2023**, *7*, 1504–1514.
18. Moteff, J.; Parfomak, P. *Critical Infrastructure and Key Assets: Definition and Identification*; Congressional Research Service: Washington, DC, USA, 2004.
19. U.S. Department of Energy. *Data Centers and the Clean Energy Transition*; DOE, Office of Energy Efficiency & Renewable Energy: Washington, DC, USA, 2024.
20. Cisco Systems. *Cisco Global Cloud Index: Forecast and Methodology, 2016–2021 (White Study)*; Cisco Systems, Inc.: San Jose, CA, USA, 2018.
21. Synergy Research Group. *Hyperscale Data Center Count and Locations, 2020*; Synergy Research Group: Reno, NV, USA, 2021.
22. Synergy Research Group. *Hyperscale Data Center Count Hits 1,136; US Accounts for 54% of Total Capacity*; Synergy Research Group: Reno, NV, USA, 2025.
23. Cushman & Wakefield. *North American Data Center Report 2023*; Cushman & Wakefield: Chicago, IL, USA, 2023.
24. Richins, D.; Doshi, D.; Blackmore, M.; Nair, A.T.; Pathapati, N.; Patel, A.; Daguman, B.; Dobrijalowski, D.; Illikkal, R.; Long, K.; et al. Ai tax: The hidden cost of ai data center applications. *ACM Trans. Comput. Syst. (TOCS)* **2021**, *37*, 3.
25. Cohen, I.; Fitch, P.; Shwisberg, J. *The Cost of Connecting: Data Center Energy and Infrastructure Costs in the United States*; Rocky Mountain Institute: Basalt, CO, USA, 2025.
26. Guo, Y.; Gong, Y.; Fang, Y.; Khargonekar, P.P.; Tang, X. Energy and Network Aware Workload Management for Sustainable Data Centers. *IEEE Trans. Parallel Distrib. Syst.* **2021**, *25*, 2030–2042.
27. Dulani, J.; Zaidi, S.S.H. Community Impact Assessment of Data Center Development: A Systematic Review. *J. Urban Plan. Dev.* **2025**, *151*, 04025012.
28. Statista. *Average Construction Cost per Square Foot for Data Centers in the United States from 2018 to 2023*; Statista: Hamburg, Germany, 2024.
29. Chen, H.; Li, D. Current status and challenges for liquid-cooled data centers. *Front. Energy Res.* **2022**, *10*, 952680. [CrossRef]
30. Uptime Institute. *Data Center Cooling Technology Trends*; Uptime Institute Intelligence: New York, NY, USA, 2024.
31. World Bank. *Commodity Markets Outlook, January 2025*; World Bank: Washington, DC, USA, 2025.
32. Yang, C.T.; Yan, Y.Z.; Chen, S.T.; Liu, R.H.; Ou, J.H.; Chen, K.L. IGEMS: A Cloud Green Energy Management System in Data Center. In *Green, Pervasive, and Cloud Computing: 11th International Conference, GPC 2016, Xi'an, China, 6–8 May 2016*; Springer International Publishing: Cham, Switzerland; pp. 82–98.
33. Watkowski, L.; Oosterwyk, G.; Elo, J. The Hidden Costs of AI-driven Data Center Demand: Five Systemic Tensions. In *Proceedings of the Americas Conference on Information Systems*; Association for Information Systems, 2025. Available online: <https://aisel.aisnet.org/amcis2025/intelfuture/intelfuture/21/> (accessed on 13 July 2026).
34. Kambhampati, V.; Dobbeltstein, A.V.D.; Schild, J. Moving beyond diesel generators: Exploring renewable backup alternatives for data centers. *J. Phys. Conf. Ser.* **2024**, *2929*, 012008. [CrossRef]
35. Gao, Y.; Hu, L.; Zhu, Y.; Li, K.; Zhao, H.; Ma, X.; Wang, L. Construction and practice of national basic science data center. *Chin. Sci. Bull.* **2024**, *69*, 3578–3588.
36. Turner & Townsend. *Data Centre Construction Cost Index 2025–2026*; Turner & Townsend: London, UK, 2025.
37. JLL. *2026 Global Data Center Outlook*; Jones Lang LaSalle: Chicago, IL, USA, 2026.
38. Andrae, A.S. New Perspectives on Internet Electricity Use in 2030. *Eng. Appl. Sci. Lett.* **2023**, *6*, 11–37.

39. Cheng, H.; Liu, B.; Lin, W.; Ma, Z.; Li, K.; Hsu, C.H. A survey of energy-saving technologies in cloud data centers. *J. Supercomput.* **2021**, *77*, 13385–13420. [CrossRef]
40. Amanatidis, P.; Lyratzis, E.; Angelopoulos, V.; Kouloumpris, E.; Skaperdas, E.; Bassiliades, N.; Vlahavas, I.; Maris, F.; Emmanoueloudis, D.; Karampatzakis, D. Intelligent Water Management Through Edge-Enabled IoT, AI, and Big Data Technologies. *IoT* **2025**, *7*, 5. [CrossRef]
41. CBRE. *2022 Data Center Operations Report*; CBRE Data Center Solutions: Dallas, TX, USA, 2022.
42. Lawrence, T.S.; Margala, M.; Chakrabarti, P. Cybersecurity Meets Carbon Neutrality: Strategies for Sustainable Data Center Security Operations. *Sustain. Comput. Inform. Syst.* **2025**, *49*, 101281. [CrossRef]
43. Bondhala, S. Cybersecurity in AI-Driven Data Centers: Reinventing Threat Detection. *Int. J. Adv. Res. Sci. Commun. Technol.* **2025**, *2025*, 510–519. [CrossRef]
44. International Energy Agency. *Energy and AI*; Licence: CC BY 4.0; IEA: Paris, France, 2025.
45. Uptime Institute. *Uptime Institute Global Data Center Survey 2021*; Uptime Institute Intelligence: New York, NY, USA, 2021.
46. Patel, K.; Ghai, R.; Gupta, P. The Power-Hungry Giants: Forecasting Data Center Electricity Demand in Virginia. *Resour. Conserv. Recycl.* **2024**, *186*, 106789.
47. Mitchell-Jackson, J.; Koomey, J.G.; Nordman, B.; Blazek, M. Data center power requirements: Measurements from Silicon Valley. *Energy* **2003**, *28*, 837–850. [CrossRef]
48. Turner, T. What Virginia Is Spending, What Virginia Is Getting: The Commonwealth Data Center Reference. 2026. Available online: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6941078 (accessed on 13 July 2026).
49. Yu, L.; Jiang, T.; Zou, Y. Real-Time Energy Management for Cloud Data Centers in Smart Microgrids. *IEEE Access* **2021**, *4*, 941–950.
50. Van der Maas, R. The Potential Of Reducing Carbon Emissions By Demand Shifting In Dutch Data Centers. 2025. Available online: https://www.cs.ru.nl/bachelors-theses/2025/Robert_van_der_Maas___1083978___The_Potential_Of_Reducing_Carbon_Emissions_By_Demand_Shifting_In_Dutch_Data_Centers.pdf (accessed on 13 July 2026).
51. Bashroush, R.; Lawrence, A. A Comprehensive Review of Energy Efficiency Metrics for Data Centers. *ACM Comput. Surv.* **2022**, *54*, 1–35.
52. Asadi, S.; Dahrebi, E. The Impact of Data Centers on the Environment and Society. In *Advances in Information Technology and Industry Applications*; Springer: Berlin, Germany, 2022; pp. 453–462.
53. Blenkinsop, J.; Grollemund, P.M. *The Global E-Waste Monitor 2024*; United Nations University: Tokyo, Japan, 2024.
54. Shivashankar, V. The Dark Side of the Digital Revolution: E-Waste Management Challenges. *Waste Manag.* **2023**, *154*, 79–91.
55. Levy, K.E.; Johns, D.M.; Hogan, A. The Data Center Paradox: Socio-Spatial Consequences. *Inf. Commun. Soc.* **2023**, *26*, 611–630.
56. Siddik, M.A.B.; Amaya, Y.G.; Marston, R. The Cost of Decommissioning: Data Center End-of-Life Management. *Resour. Conserv. Recycl.* **2023**, *187*, 106799.
57. Shehabi, A.; Smith, S.; Sartor, D.; Brown, R.; Herrlin, M.; Koomey, J.; Masanet, E.; Horner, N.; Azevedo, I.; Lintner, W. United States Data Center Energy Usage Report. 2016. Available online: <https://escholarship.org/content/qt84p772fc/qt84p772fc.pdf> (accessed on 13 July 2026).
58. Van Der Vlist, F.; Helmond, A.; Ferrari, F. Big AI: Cloud infrastructure dependence and the industrialisation of artificial intelligence. *Big Data Soc.* **2024**, *11*, 1–16. [CrossRef]
59. Brooks, L.; Liscow, Z. Infrastructure costs. *Am. Econ. J. Appl. Econ.* **2023**, *15*, 1–30. [CrossRef]
60. Nicolao, A. Netflix to Lift Prices for US Customers by Up to 18%. 2019. Available online: <https://www.ft.com/content/9f060f18-18d5-11e9-9e64-d150b3105d21?syn-25a6b1a6=1> (accessed on 13 July 2026).
61. Song, H.; Fink, G.; Jeschke, S. Pricing Strategies for AI-as-a-Service. *J. Serv. Res.* **2023**, *26*, 53–72.
62. Bughin, J.; Seong, J.; Manyika, J.; Chui, M.; Joshi, R. *Notes from the AI Frontier: Modeling the Impact of AI on the World Economy*; McKinsey Global Institute: New York, NY, USA, 2019.
63. Secinaro, S.; Calandra, D.; Secinaro, A.; Muthurangu, V.; Biancone, P. The role of artificial intelligence in healthcare: A structured literature review. *BMC Med. Inform. Decis. Mak.* **2021**, *21*, 125. [CrossRef] [PubMed]
64. Wang, Y.; Li, Y.; Cheng, X. A Comprehensive Review of AI Techniques for Data Center Energy Prediction. *Appl. Energy* **2022**, *307*, 118216.
65. Papanicolas, I.; Woskie, L.R.; Jha, A.K. Health Care Spending in the United States and Other High-Income Countries. *JAMA* **2018**, *319*, 1024–1039. [CrossRef] [PubMed]
66. Singh, S.; Yassine, A. Big Data Mining of Energy Time Series for Behavioral Analytics. *Energies* **2022**, *11*, 452.
67. Rollins, M.; Naudet, B. The Economics of Intelligent Tutoring Systems. *Educ. Technol. Res. Dev.* **2022**, *70*, 1059–1079.
68. Hemelt, S.W.; Stange, K.M. Why the Move to Online Instruction Won't Reduce College Costs. *Brookings Institution*. 2020. Available online: <https://www.brookings.edu/articles/why-the-move-to-online-instruction-wont-reduce-college-costs/> (accessed on 13 July 2026).
69. Hawi, F.; Okeyo, G.; Kimwele, M. Artificial Intelligence in Transportation: A Review. *IEEE Access* **2023**, *11*, 12345–12360. [CrossRef]

70. Bösch, P.M.; Becker, F.; Becker, H.; Axhausen, K.W. Cost-Based Analysis of Autonomous Mobility Services. *Transp. Policy* **2018**, *64*, 76–91. [\[CrossRef\]](#)
71. Owen, G. The Hidden Costs of E-Commerce. *J. Bus. Ethics* **2021**, *174*, 599–618.
72. Rani, A. The Economics of Artificial Intelligence in Retail. *J. Retail. Consum. Serv.* **2022**, *68*, 103093.
73. Allen, W.; Berg, C.; Davidson, S.; Lane, A.M. The Impact of Artificial Intelligence on Economic Growth and Sustainability. *Sustainability* **2021**, *13*, 8546. [\[CrossRef\]](#)
74. Zharova, A.; Elin, V. Regulation of AI in the Banking Sector. In *The Future of Finance*; Palgrave Macmillan: London, UK, 2021; pp. 17–36.
75. Acemoglu, D.; Restrepo, P. Automation and New Tasks: How Technology Displaces and Reinstates Labor. *J. Econ. Perspect.* **2019**, *33*, 3–30. [\[CrossRef\]](#)
76. Bousdekis, A.; Mentzas, G. A Comprehensive Review of AI Technologies for Supply Chain Management. *J. Intell. Manuf.* **2023**, *34*, 1545–1576.
77. Wu, J.; Guo, S.; Huang, H.; Liu, W.; Xiang, Y. ICTs for Sustainable Development Goals. *IEEE Commun. Surv. Tutor.* **2022**, *20*, 2389–2406.
78. Haenlein, M.; Kaplan, A. A Brief History of Artificial Intelligence. *Calif. Manag. Rev.* **2019**, *61*, 5–14. [\[CrossRef\]](#)
79. Boyd, M.; Holton, R.J. Technology, Innovation, Employment and Power. *J. Sociol.* **2018**, *54*, 331–345.
80. Muro, M.; Maxim, R.; Whiton, J. *Automation and Artificial Intelligence: How Machines Are Affecting People and Places*; Brookings Institution: Washington, DC, USA, 2019.
81. Boskin, M.J.; Dulberger, E.R.; Gordon, R.J.; Griliches, Z.; Jorgenson, D.W. Consumer Prices, the Consumer Price Index, and the Cost of Living. *J. Econ. Perspect.* **1998**, *12*, 3–26. [\[CrossRef\]](#)
82. Kansal, P.; Singh, A.; Ahluwalia, P.K. The Impact of Artificial Intelligence on the Household Sector. In *AI for Sustainable Development*; Springer: Berlin, Germany, 2021; pp. 111–128.
83. Belser, P.; Vazquez-Alvarez, R.; Xu, D. *The Impact of AI on Jobs and Skills: A Review*; International Labour Organization: Geneva, Switzerland, 2020.
84. Khanna, A. AI in the Health Care Industry: Driving Efficiency and Better Patient Experiences. *Harvard Business Review*. 2024. Available online: <https://hbr.org/resources/pdfs/comm/salesforce/AIinTheHealthCareIndustry.pdf> (accessed on 13 July 2026).
85. Lee, J.H.; Ryu, G.S.; Choi, H.; Yoon, Y. The Impact of AI on Higher Education Costs. *Econ. Educ. Rev.* **2023**, *90*, 102347.
86. Vemuri, N.; Tatikonda, V.M.; Thaneeru, N. Enhancing public transit system through AI and IoT. *Val. Int. J. Digit. Libr.* **2024**, *12*, 1057–1071. [\[CrossRef\]](#)
87. Jiang, Q.; Zhang, R.; Chen, H.; Mao, Y.; Li, K. Data-Driven Optimization of Urban Water Distribution Systems. *Water Res.* **2022**, *212*, 118140.
88. Mahmoodi, T.; Cheng, X. AI for Green and Sustainable 6G Networks: A Survey. *IEEE Commun. Surv. Tutor.* **2021**, *23*, 2135–2160.
89. BLS. *Consumer Expenditure Survey*; U.S. Bureau of Labor Statistics: Washington, DC, USA, 2021.
90. Atske, S.; Perrin, A. *Home Broadband Adoption, Computer Ownership Vary by Race, Ethnicity in the U.S.*; Pew Research Center: Washington, DC, USA, 2021.
91. EIA. *Electric Power Monthly, February 2022*; U.S. Energy Information Administration: Washington, DC, USA, 2022.
92. Pew Research Center. *Internet/Broadband Fact Sheet*; Pew Research Center: Washington, DC, USA, 2022.
93. ICDR. *The Costs of Living with a Disability in the United States*; International Center for Disability Resources: Washington, DC, USA, 2021.

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.