Flexible Computing Loads & Carbon Emission Reduction: A Tale of Bitcoin Mining and AI Training

March 26, 2026

Abstract

Although Bitcoin mining consumes lots of electricity, does it necessarily increase carbon emissions? Inspired by this inquiry, we develop an analytical framework for *flexible computing loads* (FCLs), that is, intensive electricity demands that can be elastically scaled up or down in real time without disrupting later operations. In a model of integrated electricity grid with endogenous renewable, fossil, and FCL capacity, we show that integrating FCLs via a demand response program can generate a *triple win*: lower carbon emissions, higher utility profit, and lower blackout probability. The key intuition builds on two reinforcing channels: (1) FCLs monetize renewable surplus that would otherwise be curtailed, raising returns on renewable investments; (2) FCLs substitute for fossil backups during peak hours, reducing fossil usage. The triple win obtains if and only if the amortized capital cost of FCL hardware falls below a threshold (which we derive in our framework)—a condition comfortably satisfied under current Bitcoin mining economics. We analyze this mechanism across multiple market structures, finding that rate-of-return regulation amplifies the benefit by passing mining revenue through to consumers as lower retail prices, while pre-installed legacy capacity can reverse the carbon effect in the short run. While Bitcoin mining is the archetypal FCL, AI training has the potential to generate similar benefits (though less perfectly). Our analysis lays a rigorous foundation for evaluating the carbon-reduction potentials of FCLs, and provides guidance for designing demand response programs to harness them.

Keywords: AI, Bitcoin, blockchain, carbon emissions, demand response, electricity pricing policy and regulation, flexible computing loads, mining, renewable energy, sustainability

1 Introduction

Bitcoin mining, and more broadly proof-of-work blockchain operations, have long been criticised for their high electricity consumption. Media headlines frequently compare Bitcoin mining's energy use to that of a mid-sized country (Forbes, 2022), and some studies have made the bold claim that Bitcoin mining alone could contribute to a 2°C rise in global temperatures (Mora, Rollins, Taladay, Kantar, Chock, Shimada and Franklin, 2018). Some scholars have even pointed to their policy analyses linking Bitcoin mining to carbon emissions as a primary justification for China's ban of Bitcoin mining in September 2021.[1] Given the recent political shifts, rising Bitcoin prices, and increasing adoption by governments and corporations—including recent initiatives for strategic Bitcoin reserves (U.S. Senator Lummis, 2024; The Washington Post, 2024) and the emergence of Bitcoin treasury companies—the energy demand from Bitcoin mining is likely to grow further, making a rigorous assessment of its environmental impact more urgent than ever.

Meanwhile, the explosive growth of artificial intelligence has created an even larger class of electricity-intensive computing operations. Training a frontier AI model today can require tens of thousands of specialized processors running continuously for weeks, with energy demands rivalling those of a mid-sized data center. Like Bitcoin mining, such batch training workloads are elastically scalable: the GPU cluster can be to large extent scaled up or down in real time without disrupting later operations, so the jobs can in principle be scheduled around grid conditions. This shared property—the ability to absorb or release electricity demand in real time without disrupting later operations—defines what we call *flexible computing loads* (FCLs). Bitcoin mining is the archetypal FCL; AI training is a close but imperfect cousin, with important qualifications we examine in Section 8. Developing a rigorous analytical framework for FCLs as a class—rather than treating Bitcoin mining and

[1] Private communication. On China's ban, see https://www.reuters.com/world/china/china-central-bank-vows-crackdown-cryptocurrency-trading-2021-09-24/.

AI in isolation—is one of the contributions of this paper.

While it is indisputable that Bitcoin mining (and FCLin general) consumes vast amounts of electricity, we argue that a closer analysis of its economic structure reveals a strikingly different possibility. With appropriately designed demand response programs, the high electricity demand from Bitcoin mining may actually *lower* net carbon emissions. To see why, note two characteristics of the electricity system that our model makes precise. First, renewable energy sources—wind, solar—feature high fixed costs but near-zero variable costs and imperfect correlation between supply and demand. This supply-demand mismatch forces grids to rely on fossil "peaker" plants for balancing, imposing carbon costs and undermining the commercial case for renewable investment. Second, Bitcoin mining is *memoryless*: every joule of hashing contributes equally to expected rewards regardless of when it occurs, so miners are perfectly indifferent between running now and running later. This temporal flexibility—unique among large electricity consumers—makes Bitcoin mining an ideal shock absorber for intermittent renewable supply. Mining can be dialed up during renewable surplus (earning revenue on power that would otherwise be curtailed) and dialed down during scarcity (freeing fossil capacity for consumers), simultaneously raising the profitability of renewable investment and reducing the need for fossil backup.

It is important to stress that this carbon-reduction effect is not automatic. It requires that the demand response program be appropriately designed, and it depends on key parameters—most importantly the amortized hardware cost of mining equipment relative to the expected per-megawatt revenue from flexible dispatch. Our paper characterizes these conditions precisely. When they are satisfied, mining generates a *triple win*: lower carbon emissions, lower blackout probability, and higher utility profit. When they fail, the reverse holds. The even-handed characterization of both cases is, we argue, the central contribution of a formal analysis that informal advocacy cannot provide.[2]

[2]Some practitioners and policy advocates have argued informally that Bitcoin mining can act as a flexible

This paper develops the first rigorous analytical framework for studying these questions. Unlike conventional electricity uses—heating, lighting, industrial processes—where curtailment imposes an immediate service disruption, FCLs can be elastically scaled up or down in real time without disrupting later operations.

Our baseline model features a regulated electricity system in which a profit-maximizing utility jointly decides how much renewable capacity, fossil backup capacity, and FCL capacity to install. The key trade-off is between the high fixed cost of renewable capacity (which generates electricity at zero marginal cost but intermittently) and the high variable cost of fossil backup (which generates reliably but emits carbon). We show that integrating FCLs via a demand response program solves this trade-off through two reinforcing channels. First, FCLs monetize renewable surplus that would otherwise be curtailed, raising the revenue from renewable investment and thus incentivizing more capacity. Second, FCLs substitute for fossil backup during hours when renewables just fall short of consumer demand, reducing the installed fossil capacity at the optimum. Together, these channels can lower total carbon emissions while increasing utility profit.

We characterize the exact conditions under which this *triple win* occurs—higher renewable investment, lower carbon emissions, and higher utility profit—and show that these conditions coincide. The mining preference condition takes the intuitive form that the amortized capital cost of FCL hardware must be below the expected per-megawatt value of flexible dispatch, which equals the per-MWh profit margin times the fraction of hours the grid has spare fossil capacity.

Beyond the baseline model, we analyze four extensions that are critical for external validity and policy relevance. First, we consider *timing variations*: many grids have legacy fossil or renewable capacity that is not freely adjustable. We show that the carbon improvement

grid resource (Ibañez and Freier, 2023; Bitcoin Policy Institute, 2024), and sporadic empirical evidence supports this channel (Forbes, 2024; Halaburda and Yermack, 2023). Our paper provides the analytical foundation for evaluating when and why the mechanism operates, and when it does not.

channel critically relies on endogenous fossil capacity choice—with fully fixed capacities, the demand response program can increase carbon by raising the fossil plant's utilization rate. Second, we embed the model in the *Averch–Johnson rate-of-return regulation* framework, where the retail electricity price is determined endogenously. Under ROR regulation, mining revenue creates a "double win": Third, we analyze *deregulated competitive electricity markets*, showing that the competitive equilibrium replicates the regulated outcome with the market's scarcity price playing the role of the regulated penalty, and that FCLs help address the "missing money" problem facing fossil generators. Fourth, we carefully examine the *scope of the grid*: whether the model describes a price-taking regional grid or one with market power in the global hash market, and how coordination failures arise when many grids simultaneously adopt demand response programs.

Finally, we quantify the model's predictions using current market data. Prevailing Bitcoin prices, reward per block (the "coinbase" reward), and falling renewable capital costs imply that the viable range for FCL adoption today is roughly three times wider than under our baseline calibration. For AI training, the economics are less favorable but still promising: the viable range is approximately 20% wider than under our baseline. These quantitative assessments suggest that the carbon-reduction potential of FCLs is not just a theoretical possibility but a practical reality under current market conditions.

Key contributions. This paper makes three contributions to the intersection of energy economics, blockchain, and operations management.

First, we provide the first analytical characterization of FCLs as a grid resource, unifying Bitcoin mining and AI training within a common framework that identifies when and why flexible computing loads can reduce carbon emissions. Previous work on demand response (Aïd, Possamaï and Touzi, 2022; Fattahi, Dasu and Ahmadi, 2023) does not consider the specific properties of computing loads (memoryless, globally competitive, transparent rev-

enue), and previous work on Bitcoin mining (Halaburda and Yermack, 2023; Cong, He and Li, 2021) does not develop a full welfare analysis within an electricity system model.

Second, we provide tractable comparative-static results across a wide range of market structures—regulated utilities, rate-of-return regulation, and competitive markets—enabling a unified policy analysis of FCL programs across different electricity system designs.

Third, we identify and formalize the key differences between Bitcoin mining and AI training as FCL technologies, including the inference latency constraint, revenue opacity, and power surge behavior, providing a theoretic foundation for the emerging debate about whether data centers can serve as flexible grid resources.

The rest of the paper is organized as follows. Section 2 surveys related literature. Section 3 presents the model. Section 4 analyzes the baseline simultaneous optimization. Section 5 studies timing variations with pre-installed capacity. Section 6 analyzes rate-of-return regulation. Section 7 analyzes deregulated competitive markets. Section 8 compares Bitcoin mining and AI training as FCL technologies. Section 9 presents the quantitative assessment. Section 10 concludes. Appendices contain proofs, the global mining Nash equilibrium extension, the analysis of the $R \leq c$ parameter region, and alternative regulatory objective functions.

2 Related Literature

Our work contributes to the intersection of four distinct streams of literature: research on Bitcoin mining and blockchain, demand response programs, renewable energy investments and electricity markets, and the economics of artificial intelligence and data centers.

Research on Bitcoin mining. On the theory side, Biais, Bisiere, Bouvard and Casamatta (2019) model Bitcoin mining as a coordination game and identify the possibility of multiple

equilibria. Cong, He and Li (2021) studies the dynamics of mining pool formation and shows that the rise of mining pools may lead to significant electricity consumption. Prat and Walter (2021) propose a model that uses the Bitcoin–USD exchange rate to predict the computing power of the network. Arnosti and Weinberg (2022) show that Bitcoin mining has a tendency toward significant centralization given even small initial differences in variable costs. Bertucci, Bertucci, Lasry and Lions (2024) models Bitcoin mining as a mean-field game.

On the empirical side, the most directly related paper is Halaburda and Yermack (2023), who study the operations and financial valuations of 13 cryptocurrency mining companies listed on NASDAQ and find that miners using Texas wind power are offline more than other miners in a more erratic pattern, yet are more profitable due to extremely low prices during renewable surplus periods. Their findings directly support the mechanism we model. Benetton, Compiani and Morse (2023) document that cryptocurrency miners' use of local electricity implies higher electricity prices for households and small businesses—an effect that disappears when demand response programs are designed to price-discriminate between mining and non-mining loads, as in our model.

Several industry reports provide supporting evidence. Ibañez and Freier (2023) contain numerous verbal arguments on why Bitcoin mining may be environmentally friendly that echo key aspects of this paper. Bitcoin Policy Institute (2024) document data patterns consistent with Bitcoin mining's flexible electricity demand. There is also extensive anecdotal evidence of grids using Bitcoin mining to absorb surplus electricity, including documented cases in Texas, Pakistan, and Alberta.

The flexibility of Bitcoin mining at the hardware level is grounded in ASIC design. Miners have been reported to switch operations "at a minute's notice" (Bloomberg, 2021; CNBC, 2022; Forbes, 2021), and technical product manuals (BITMAIN Support, 2022) confirm that rapid on/off cycling is a design feature. Yaish, Tochner and Zohar (2022) show that miners

can repeatedly switch machines to manipulate difficulty in ways that increase profits, confirming that rapid cycling is economically feasible. Yaish and Zohar (2023) further document two firms—Layer1 Technologies (2022) and Blockstream (2022)—that designed ASICs *explicitly* as grid shock absorbers, underscoring that hardware-level flexibility is an intentional design feature rather than an artifact. Before Bitcoin mining was banned in China, miners routinely migrated between hydro-rich regions as the wet and dry seasons shifted renewable availability (Bitcoin.comNews, 2019; decrypt, 2020; Cointelegraph, 2020)—an early large-scale demonstration of FCL demand flexibility in practice. Kawaguchi, Noda and Komiyama (2022) provide econometric evidence that miners adjust their electricity consumption in response to profitability shocks, confirming elastic demand response in the field.

Research on demand response programs. A stream of research has explored demand response programs in electricity markets. Peura and Bunn (2015) examine how peak electricity producers set prices under uncertain demand. Aïd, Possamaï and Touzi (2022) analyze optimal demand response contracts. Fattahi, Dasu and Ahmadi (2023) propose a scenario-based approach to optimize direct load control contracts. Lohmann and Rebennack (2017) develop a long-term power generation expansion model that factors in short-term demand response. Agrawal and Yücel (2022) investigate the role of baseline adjustments in demand-response programs. Together, these studies highlight the importance of demand response in balancing supply and demand, improving grid stability, and reducing costs. However, none examines how such programs can leverage the specific properties of computing loads to facilitate renewable investment—the central contribution of this paper.

Research on electricity markets and renewable investment. Electricity supply has historically been treated as a natural monopoly subject to rate-of-return (ROR) regulation, under which utilities recover prudently incurred costs plus a regulated return on invested capital (Train, 1991; Jamison, 2008). The landmark pathology of this regime is the Averch–

Johnson overcapitalisation effect (Averch and Johnson, 1962): because allowed profits scale with the rate base, regulated firms over-invest in capital relative to the social optimum. This observation directly informs our Section 6.

Beginning in the 1990s, dissatisfaction with ROR regulation motivated electricity sector restructuring: separating generation from transmission, introducing bid-based wholesale markets, and creating independent system operators (Joskow, 1997, 2003). Evidence on restructuring's efficiency effects is largely positive (Fabrizio, Rose and Wolfram, 2007; Davis and Wolfram, 2012), though market power remains a concern (Wolak, 2003). Electric Power Research Institute (2024) project significant near-term load growth from data centers and identify grid reliability risks. Satchwell, Mims Frick, Cappers, Sergici, Hledik, Kavlak and Oskar (2025) survey large-load tariffs— real-time pricing, demand charges, and interruptible service agreements— through which utilities manage financially and operationally challenging loads.

Among academic work on renewable investment, Aflaki and Netessine (2017) show that carbon pricing may sometimes discourage renewable adoption due to intermittency challenges. Kök, Shang and Yücel (2018) show that flat pricing can encourage solar adoption. Kök, Shang and Yücel (2020) find that renewable and flexible energy sources complement each other, whereas inflexible sources act as substitutes. Peng, Wu and Souza (2024) study the interplay between renewable generation and energy storage. Andonov and Rauh (2024) document the increasing entry of private equity in U.S. electricity markets, especially renewable investment. Darmouni, Lehrer and Sastry (2025) point out that the high fixed cost of renewables justifies ex-ante power purchase agreements. These papers provide context for understanding FCLs as a complementary mechanism for enabling renewable investment.

Research on AI, data centers, and electricity. The electricity demands of artificial intelligence have received growing attention. Electric Power Research Institute (2024) project

that U.S. data centers will require 90–100 GW of additional generation capacity by 2030, placing AI on a scale comparable to current Bitcoin mining. International Energy Agency (2024) document that global data centers consumed approximately 415 TWh in 2024, and project this to exceed 800 TWh by 2030 under a high-growth scenario.

The flexibility of AI training loads has attracted emerging interest. Modern distributed training frameworks support checkpointing—pausing mid-run without losing progress—which is directly analogous to Bitcoin miners powering down ASICs and restarting (Jiang, Lin et al., 2024). Hyperscalers (Google, Microsoft, Amazon) have begun using temporal flexibility to reduce energy costs and participate in demand response programs (Acun et al., 2023; Radovanović, Koningstein et al., 2023). The feasibility of demand-response participation by data centers is documented in Texas, where ERCOT's large flexible load programme includes multiple data center operators alongside Bitcoin miners.

The key difference between AI training and Bitcoin mining lies in *AI inference*: serving live queries is latency-constrained in ways that hashing is not, reducing the curtailable fraction of AI electricity demand as inference becomes a larger share of total AI compute (International Energy Agency, 2024; Sachs, 2024). Our Section 8 formally models this distinction and characterizes the conditions under which each technology is the superior flexible load.

3 Model

This section presents the model, which captures a regulated electricity system with stochastic renewable generation, fossil backup, and an optional demand response program that integrates flexible computing loads.

Our baseline model has one strategic decision maker: a regulated utility. The utility chooses three capacity levels—renewable capacity K, fossil backup capacity k, and flexible

computing load capacity δ—to maximize expected profit. Its payoff is revenues from consumer electricity sales and computing revenue, minus variable fossil costs, blackout penalties (which the regulator requires it to internalize), and capital costs for all three capacity types. In the baseline model the utility sets δ directly; we relax this assumption in Section 8. The regulator sets the retail price p exogenously and imposes a blackout penalty β; we will endogenize p in Section 6 and Section 7 under alternative market structures.

3.1 Key model ingredients

The model features a market for electricity within a representative period. The key state variable is *residual demand*

$$\tilde{x} \equiv \tilde{D} - K\tilde{S}, \tag{1}$$

where $\tilde{D}$ is stochastic consumer (non-computing) demand drawn from distribution F_D, K is installed renewable capacity, and $\tilde{S} \in [0, 1]$ is the stochastic renewable capacity factor. When $\tilde{x} > 0$, renewables fall short of consumer demand and fossil backup is needed. When $\tilde{x} < 0$, there is a renewable surplus. Let $F(x; K)$ denote the CDF of $\tilde{x}$ given renewable capacity K, with density $f(x; K)$. More renewable capacity shifts the distribution leftward: $\partial F/\partial K > 0$ for any given x, meaning a higher K reduces the probability of high residual demand.

The supply side consists of: (1) renewable capacity K (MW) with zero variable cost and capital cost C per MW; (2) fossil backup capacity k (MW) with variable cost c per MWh and capital cost γ per MW;[3] and (3) an optional FCL program of capacity δ (MW) generating revenue R per MWh consumed, with capital cost κ per MW. A blackout occurs when residual demand exceeds total dispatchable capacity and imposes social penalty β per

[3]We model the fossil source as a flexible, dispatchable plant (e.g., a natural gas peaker). Inflexible conventional sources such as nuclear cannot be rapidly ramped up or down, so they act as *substitutes* for renewables rather than complements (Kök, Shang and Yücel, 2020): their fixed output reduces the need for renewable balancing but does not enable the demand-flexibility mechanism we study. Including nuclear would not alter the key renewable-versus-fossil trade-off that drives our results.

MWh of unserved demand. Consumers pay retail price p per MWh.

Assumption 1 (Parameter regularity)**.** *We maintain throughout: $R > c$ (computing is profitable per MWh above variable cost), $\beta + p > c + \gamma$ (it is worth building any fossil backup), and $\beta + p > R$ (consumer welfare exceeds computing revenue per MWh, so optimal fossil capacity is finite). The case $R \leq c$ is analyzed separately in Appendix D.*

Within the regime $R > c$, one further comparison is economically meaningful: The relationship between R and $c+\gamma+\kappa$ determines whether the utility's optimal fossil capacity is interior or at a corner:

- *Interior regime* ($R < c + \gamma + \kappa$): the joint optimum has $k^{**} \in (0, \bar{D})$ with a positive residual blackout probability;
- *Corner regime* ($R \geq c + \gamma + \kappa$): mining revenue is so high relative to fossil and ASIC costs that the utility finds it optimal to eliminate blackouts entirely ($k^{**} = \bar{D}$).

Both regimes are analyzed in Section 4.2.

The role of the retail price p differs by market structure. In the baseline model and the timing-variations section, p is set exogenously by the regulator. Under rate-of-return regulation (Section 6), p is determined endogenously to ensure the utility earns a fair return. In deregulated markets (Section 7), p is replaced by a spot price that clears the wholesale market.

3.2 Demand response program for flexible computing loads

The utility operates a Demand Response Program that dispatches registered FCL operators. FCL demand under the Demand Response Program is given by

$$d(x) = \min\Big\{\max\{k - x,\, 0\},\, \delta\Big\}, \tag{2}$$

where x is the realized value of $\tilde{x}$. This formulation captures the key properties of FCL demand response: (i) when residual demand $x \in (k-\delta, k)$, FCLs absorb the exact capacity gap $k-x$, running on fossil power (the utility pays them revenue R and incurs variable cost c per MWh); (ii) when $x \leq k-\delta$ (fossil has more spare capacity than δ, or renewables are in surplus), FCLs run at full capacity δ; (iii) FCLs shut down during scarcity events when $x > k$, providing downward demand response that partially alleviates shortage.

The elastic scalability of FCLs is what makes this dispatch rule sensible. Because FCLs can be scaled up or down in real time without disrupting later operations, curtailment imposes no service disruption. For AI training, a similar logic applies to batch workloads with flexible scheduling horizons (see Section 8 for a detailed analysis of inference versus training).

3.3 Utility company's problem

A profit-maximizing utility simultaneously chooses renewable capacity K, fossil backup capacity k, and FCL capacity δ to maximize expected profit. Incorporating the demand response program (2):

$$\max_{K,k,\delta \geq 0} \mathbb{E}\Big[p\,\tilde{D} + R \cdot d(\tilde{x}) - c \cdot \Big[\min\{k, \max\{\tilde{x} + d(\tilde{x}),\, 0\}\}\Big] - \beta \cdot (\tilde{x} - k)^{+} - CK - \gamma k - \kappa\delta\Big]. \tag{3}$$

In words, the utility earns revenue from consumer electricity sales $p\tilde{D}$, mining revenue $R \cdot d(\tilde{x})$, incurs variable fossil costs, absorbs blackout penalties (which the regulator requires it to internalize), and pays capital costs for all three types of capacity.

The Demand Response Program ensures that whenever $\tilde{x} \in (k-\delta, k)$ the fossil plant runs at its full capacity k (the FCL absorbs the residual gap $k-\tilde{x}$). When $\tilde{x} \in (0, k-\delta)$, the fossil plant runs at $\tilde{x} + \delta < k$ (FCL is at full capacity δ but consumer demand is not large

enough to require the full fossil plant). Expected fossil output is therefore $\mathbb{E}\big[\min\{k,\,(\tilde{x}+\delta)^+\}\big] = k - \int_{-K}^{k-\delta} F(x;K)\,dx$, which is strictly greater than the no-FCL expected output $\mathbb{E}[\min\{k,\tilde{x}^+\}] = \int_0^k (1-F(x;K))\,dx$. The mining programme thus raises fossil utilization and earns net computing revenue $R-c$ per MWh of computing load.

Throughout the analysis we treat the per-MWh mining revenue R as exogenous, which is appropriate for any price-taking regional grid. The full discussion of grid scope—including the endogenous-R case, Tullock mining competition, and multi-grid coordination externalities—is deferred to Appendix B.

4 Baseline Analysis: Simultaneous Capacity Choice

This section analyzes the full joint optimization in which the utility simultaneously chooses renewable capacity K, fossil backup capacity k, and FCL capacity δ. This is the natural benchmark when all three investments can be made simultaneously (e.g., when designing a new electricity system from scratch or executing a comprehensive grid upgrade). Throughout this section, we work under Case c ($\delta^{**} > K^{**}$, Assumption 1), the empirically relevant regime in which optimal FCL capacity exceeds renewable capacity; the small-FCL case (Case b: $\delta^{**} \le K^{**}$) is analyzed in Appendix F. We also maintain the interior regime ($R < c+\gamma+\kappa$, Assumption 1); the corner regime is discussed in Section 4.2 below. We first characterize the first-order conditions and optimal allocations (Section 4.1), then compare blackout risk (Section 4.2), carbon emissions (Section 4.3), and the conditions under which all three outcomes improve simultaneously (Section 4.4).

4.1 Optimal allocations: K^{**}, k^{**}, δ^{*}

Objective. The utility chooses (K,k,δ) to maximize expected profit (3). Plugging in the Demand Response Program dispatch rule (2) and evaluating expectations using the residual-

demand CDF $F(x;K)$, the profit function (under Case c, as assumed throughout this section) reduces to:

$$\Pi(K,k,\delta) = p\,\mathbb{E}[\tilde{D}] + R\int_{k-\delta}^{k} F(x;K)\,\mathrm{d}x - c\int_{-K}^{k-\delta}\Big(1-F(x;K)\Big)\,\mathrm{d}x \\ - (\beta+p)\int_{k}^{\bar{D}}\Big(1-F(x;K)\Big)\,\mathrm{d}x - \gamma k - CK - \kappa\delta. \quad (4)$$

The three integral terms represent expected net mining revenue, expected fossil variable cost (the plant runs whenever $\tilde{x} \in (-K, k-\delta)$ or $\tilde{x} \in (k-\delta, k)$), and expected blackout cost, respectively. Setting $\delta = 0$ gives the no-FCL benchmark:

$$\Pi_0(K,k) = p\,\mathbb{E}[\tilde{D}] - c\int_{0}^{k}\Big(1-F(x;K)\Big)\,\mathrm{d}x - (\beta+p)\int_{k}^{\bar{D}}\Big(1-F(x;K)\Big)\,\mathrm{d}x - \gamma k - CK. \quad (5)$$

Without the FCL, the fossil plant never runs during renewable-surplus hours, so the fossil cost integral starts from 0 rather than $-K$.

Without FCL program (benchmark). Differentiating (5) with respect to k and setting equal to zero gives the FOC:

$$(\beta+p-c)\Big[1-F(k^*;K^*)\Big] = \gamma \quad\Longrightarrow\quad F(k^*;K^*) = \frac{\beta+p-c-\gamma}{\beta+p-c}. \quad (6)$$

Blackout probability at the no-FCL optimum is $\gamma/(\beta+p-c) \in (0,1)$.

Fossil capacity k^{**}. Differentiating (4) with respect to k and substituting the FOC for δ (equation (10) below, which gives $F(k^{**}-\delta^*;K^{**}) = \kappa/(R-c)$ in the interior regime) yields:

$$(\beta+p-R)\,F(k^{**};K^{**}) = \beta+p-c-\gamma-\kappa \quad\Longrightarrow\quad F(k^{**};K^{**}) = \frac{\beta+p-c-\gamma-\kappa}{\beta+p-R}. \quad (7)$$

Note the presence of the ASIC cost κ in the numerator: the harder it is to install ASIC capacity (high κ), the lower the optimal $F(k^{**})$ (lower fossil reliability), and vice versa. When $\kappa \to 0$, the formula reduces to a version with R replacing the no-mining benchmark's effective "c", reflecting that mining revenue fully offsets fossil costs at the margin. The equilibrium blackout probability with mining is therefore:

$$1 - F(k^{**}; K^{**}) = \frac{c + \gamma + \kappa - R}{\beta + p - R}. \tag{8}$$

Two regimes for optimal fossil capacity. Equation (7) is derived under the *interior regime* $R < c + \gamma + \kappa$, in which the right-hand side lies in $(0, 1)$ and an interior optimum $k^{**} \in (0, \bar{D})$ exists. When instead $R \geq c + \gamma + \kappa$, the right-hand side equals or exceeds one, meaning the marginal profit of fossil capacity never turns negative before the support upper bound $\bar{D}$: the optimum is the *corner solution* $k^{**} = \bar{D}$ with zero blackout probability. Economically, when the net mining revenue per MW of fossil capacity ($R-\kappa$) covers all fossil costs ($c+\gamma$), the utility maximizes computing uptime by building enough fossil to guarantee no blackout ever occurs.

In the corner regime the FCL adoption condition $\kappa < \bar{\kappa} \equiv (R-c)(\beta+p-c-\gamma)/(\beta+p-c)$ is automatically satisfied—anticipating the formal result in Proposition 1, which shows that $\kappa < \bar{\kappa}$ is necessary and sufficient for the triple win (lower blackouts, lower carbon, and higher profit). To see this, note that $\kappa \leq R - c - \gamma$ (the condition $R \geq c + \gamma + \kappa$ rearranged) and $R - c - \gamma < \bar{\kappa}$ whenever $\beta + p > R$ (Assumption 1). Hence the corner regime is a *degenerate but favorable* special case in which zero blackouts, lower carbon, and higher profit are all achieved simultaneously.

For the remainder of the paper we focus on the empirically relevant *interior regime* $R < c + \gamma + \kappa$; all results carry over to the corner regime with the blackout probability replaced by zero.

Renewable capacity K^{**}. Applying the envelope theorem to (4) and differentiating with respect to K gives:

$$R\int_{k^{**}-\delta^*}^{k^{**}} \frac{\partial F}{\partial K}(x;K^{**})\,\mathrm{d}x + c\int_{-K^{**}}^{k^{**}-\delta^*} \frac{\partial F}{\partial K}(x;K^{**})\,\mathrm{d}x + (\beta+p)\int_{k^{**}}^{\bar{D}} \frac{\partial F}{\partial K}(x;K^{**})\,\mathrm{d}x = C, \quad (9)$$

where $\partial F/\partial K > 0$ (more renewables increase the probability of low residual demand states). The key insight is that the weight in the FCL dispatch region $(k^{**}-\delta^*, k^{**})$ is R, whereas in the no-FCL case it is c. Since $R > c$, computing revenue *amplifies* the shadow value of renewable investment, inducing higher K^{**} than the no-FCL optimum.

FCL capacity δ^*. Differentiating (4) with respect to δ, the marginal FCL unit is dispatched whenever residual demand falls below $k^{**}-\delta^*$ (spare fossil capacity exceeds δ^*):

$$(R-c)\cdot F(k^{**}-\delta^*;\, K^{**}) = \kappa. \quad (10)$$

Inverting gives $\delta^* = k^{**} - F^{-1}\Big(\kappa/(R-c);\, K^{**}\Big)$. The infra-marginal FCL units are dispatched with total probability $F(k^{**};K^{**}) > \kappa/(R-c)$ (whenever residual demand is below fossil capacity), but it is the *marginal* dispatch probability $F(k^{**}-\delta^*;K^{**}) = \kappa/(R-c)$ that pins down the equilibrium δ^*.

Uniform-distribution closed form. Under the uniform distribution $F(x;K) = (x+K)/(\bar{D}+K)$ on $[-K,\bar{D}]$, inverting $F(k^{**};K^{**}) = (\beta+p-c-\gamma-\kappa)/(\beta+p-R)$ gives:

$$k^{**} = \frac{(\beta+p-c-\gamma-\kappa)(\bar{D}+K^{**})}{\beta+p-R} - K^{**}, \quad (11)$$

$$\delta^* = k^{**} - F^{-1}\Big(\frac{\kappa}{R-c};K^{**}\Big) = k^{**} - \frac{\kappa(\bar{D}+K^{**})}{R-c} + K^{**}, \quad (12)$$

with K^{**} determined from (9). Note that $\delta^* > 0$ requires $F(k^{**}; K^{**}) > \kappa/(R-c)$, i.e., $\beta + p - c - \gamma - \kappa > \kappa(\beta + p - R)/(R-c)$, which simplifies to $\kappa < \bar{\kappa}$ (the adoption threshold defined above), confirming that the FCL program is positive precisely when it is also socially beneficial.

4.2 Blackout comparison

A blackout occurs when $\tilde{x}$ exceeds fossil capacity k. Comparing (6) and (8), straightforward algebra shows:

$$\underbrace{1 - F(k^{**}; K^{**})}_{\text{with FCL}} = \frac{c + \gamma + \kappa - R}{\beta + p - R} \lessgtr \frac{\gamma}{\beta + p - c} = \underbrace{1 - F(k^{*}; K^{*})}_{\text{without FCL}} \quad \Longleftrightarrow \quad \kappa \gtrless \bar{\kappa}. \tag{13}$$

Specifically, cross-multiplying and simplifying, the difference LHS $-$ RHS $= \Big[\kappa(\beta + p - c) - (R-c)(\beta + p - c - \gamma)\Big]/[(\beta + p - R)(\beta + p - c)]$, which is negative exactly when $\kappa < \bar{\kappa}$.

When $\kappa < \bar{\kappa}$ (the FCL adoption condition), the equilibrium blackout probability is strictly lower with the FCL program than without. The intuition is that the mining revenue $R > c$ earned from spare fossil capacity raises the per-MW profitability of fossil investment, inducing the utility to build more fossil backup. The net effect (when the ASIC cost κ is not too high) is a higher equilibrium $F(k^{**}; K^{**})$, meaning lower blackout probability, rather than a substitution away from fossil.

4.3 Carbon emissions

Carbon emissions equal c times expected fossil fuel output. The Demand Response Program creates two opposing effects on carbon.

utilization effect (carbon-increasing). Under the Demand Response Program, whenever $\tilde{x} \in (0, k)$ the fossil plant runs at full capacity k (the FCL absorbs the gap $k - x$),

whereas without the program fossil output equals only x. The Demand Response Program strictly increases the utilization rate of the installed fossil plant.

Capacity effect (carbon-ambiguous). At the joint optimum, the mining programme changes both k and K. More fossil capacity (k^{**} vs. k^{*}) tends to raise carbon, while more renewable capacity (K^{**} vs. K^{*}) shifts the residual demand distribution leftward, tending to lower fossil utilization. Under the uniform distribution on $[-K, \bar{D}]$, inverting the FOC for k (equation (7)) gives:

$$k^{**} = \frac{(\beta + p - c - \gamma - \kappa)(\bar{D} + K^{**})}{\beta + p - R} - K^{**}, \tag{14}$$

compared with $k^{*} = (\beta+p-c-\gamma)(\bar{D}+K^{*})/(\beta+p-c) - K^{*}$ without mining (the $-K$ terms arise from inverting the uniform CDF). Whether $k^{**} \gtrless k^{*}$ depends on how much K changes; the direction of the carbon effect is governed by the net balance of the two channels.

Net effect. The carbon outcome depends on which effect dominates. Under the uniform distribution:

$$\text{Carbon falls with FCL} \quad \Longleftrightarrow \quad \kappa < \bar{\kappa} \;=\; \frac{(R-c)(\beta+p-c-\gamma)}{\beta+p-c}. \tag{15}$$

Strikingly, this is the same threshold $\bar{\kappa}$ that governs the utility's profit preference for FCLs (defined in Section 4.1 and established formally in Proposition 1 below).

4.4 Conditions for FCL adoption preference

Proposition 1 (Mining preference condition and triple win)**.** *The utility strictly prefers allowing the* FCL *program if and only if*

$$\kappa < \bar{\kappa} = \frac{(R-c)(\beta+p-c-\gamma)}{\beta+p-c}, \tag{16}$$

where $\bar{\kappa}$ is the adoption threshold introduced in Section 4.1. When $\kappa < \bar{\kappa}$, allowing the FCL *program simultaneously (i) increases utility profit, (ii) reduces carbon emissions, and (iii)* lowers *the equilibrium blackout probability—a triple win. When $\kappa > \bar{\kappa}$, none of the three improvements occurs: the utility would not adopt the program, and if forced to it would face higher costs, higher carbon, and a higher blackout probability.*

The threshold $\bar{\kappa}$ has an instructive factorisation. Define $\phi^* \equiv F(k^*; K^*) = 1 - \gamma/(\beta + p - c)$ as the no-FCL probability of *no* blackout. Then

$$\bar{\kappa} = (R-c) \cdot \phi^*.$$

The per-MW value of FCL investment equals the per-MWh profit margin $(R-c)$ times the fraction of time the grid is not blacked out (i.e., the probability that spare fossil capacity is available for computing). Intuitively, computing hardware is worth building if its amortized cost κ is below its expected revenue per MW of capacity.

Comparative statics of $\bar{\kappa}$. The viability threshold $\bar{\kappa}$ is:

- *Increasing in R*: higher computing revenue (e.g., higher BTC price) widens the viable range for ASIC capital cost;
- *Decreasing in c*: cheaper fossil variable cost raises the effective per-MWh computing margin $R - c$;

- *Increasing in β and p*: higher blackout penalty or retail price raises the no-FCL reliability (ϕ^* rises), so the grid has more spare capacity hours available for computing;
- *Decreasing in γ*: cheaper fossil capital raises no-FCL reliability ϕ^* (the utility builds more fossil cheaply), again increasing spare capacity hours.

Remark 1. A necessary condition for $\bar{\kappa} > 0$ is $\beta + p > c + \gamma$, which is exactly the condition that it is worthwhile to build any fossil backup at all without the FCL program. If the grid would not build fossil capacity absent mining, the demand response mechanism has nothing to work with.

4.5 Summary table: Baseline results

Table 1 summarizes the direction of all key effects when the FCL program is introduced at the simultaneous optimum, distinguishing the two sub-cases $R - c > \gamma$ and $R - c < \gamma$.

[Table 1 about here.]

The key insight from Table 1 is that the mining-preference condition $\kappa < \bar{\kappa}$ is simultaneously necessary and sufficient for all three elements of the triple win: higher profit, lower carbon, and lower blackout probability. There is no trade-off between reliability and the other objectives—the same condition that makes the program financially viable also makes it environmentally and socially beneficial. This stands in contrast to the false intuition that FCLs substitute for fossil capacity and thereby raise blackout risk; in fact, mining revenue raises the per-MW value of fossil capacity, inducing the utility to build more fossil backup alongside more renewables.

5 Timing Variations: Pre-Installed Capacity

In practice, capacity decisions are made at different points in time. Many grids have legacy fossil fleets or mandated renewable targets that cannot be freely adjusted in the short or medium run. This section analyses three cases of pre-existing (sunk) capacity, characterizing how the results of Section 4 change when the substitution channel is constrained.

5.1 Case A: Pre-installed fossil capacity

Suppose k_0 MW of fossil capacity is already installed (a legacy coal or gas fleet). The utility takes k_0 as given and optimizes only over K and δ.

Mining and carbon. With k_0 fixed, the Demand Response Program raises fossil utilization: the plant runs at full capacity whenever there is any consumer demand ($\tilde{x} > 0$), whereas without the Demand Response Program it would run at partial capacity ($\tilde{x}$) for intermediate states. For the *same* K, the Demand Response Program increases carbon by fully utilising the pre-installed fossil plant.

However, the joint optimization of K now internalizes the higher marginal value of renewable investment (since renewables now generate revenue via computing as well as displacing fossil), leading to $K^{**} > K_0$. The net carbon effect depends on whether additional renewable investment reduces fossil use sufficiently to offset the higher utilization rate:

$$\Delta\text{Carbon} < 0 \quad \Longleftrightarrow \quad (R - c) \cdot \left.\frac{\partial\, \mathbb{E}[\text{fossil output}]}{\partial K}\right|_{K=K_0} > C.$$

When renewable capital cost C is low (as in current markets), this condition is readily satisfied.

Key insight. Carbon improvement requires that the additional renewable investment induced by computing revenue generates enough surplus hours to offset the higher fossil utilization. Grids that are already heavily renewable (large K_0) have many surplus hours where FCLs run on clean power at zero carbon cost—making the net effect unambiguously positive in such cases.

5.2 Case B: Pre-installed renewable capacity

Here K_0 MW of renewables are fixed (e.g., from prior policy mandates), and the utility optimizes only over k and δ. The FOCs are identical to those of Section 4 (equations (7) and (10)), evaluated at $K = K_0$:

$$F(k^{**}; K_0) = \frac{\beta + p - c - \gamma - \kappa}{\beta + p - R}, \qquad (R - c) \cdot F(k^{**} - \delta^*; K_0) = \kappa. \tag{17}$$

Effect of renewable over-installation. If $K_0 > K^{**}$ (the grid has more renewables than the jointly optimal level), the fossil dispatch probability is lower and FCLs spend more time running on clean renewable power. The viability threshold rises:

$$\bar{\kappa}(K_0) = (R - c) \cdot \phi(K_0) > \bar{\kappa}(K^{**}),$$

making the Demand Response Program more attractive even at higher hardware capital cost. This has direct policy relevance: grids with large solar/wind buildout and significant curtailment (ERCOT-Texas, Germany, California) have more surplus hours, making the Demand Response Program beneficial even for expensive computing hardware.

Effect of renewable under-installation. Conversely, if $K_0 < K^{**}$, fossil dispatch is more frequent, the Demand Response Program threshold falls, and the carbon benefit of FCLs shrinks. Mining is still viable if $\kappa < \bar{\kappa}(K_0)$, but the condition is tighter.

5.3 Case C: Both capacities pre-installed

In the short run, both k_0 and K_0 are fixed and only FCL capacity is chosen. The single FOC is:

$$(R - c) \cdot F(k_0; K_0) = \kappa, \tag{18}$$

yielding $\delta^* = k_0 - F^{-1}\big(\kappa/(R-c); K_0\big)$ if $\kappa/(R-c) < F(k_0; K_0)$, and $\delta^* = 0$ otherwise.

Short-run carbon impact. Since neither capacity adjusts, the only carbon effect is through changed fossil utilization, which is unambiguously positive (the Demand Response Program increases carbon):

$$\Delta\text{Carbon} = c \cdot \mathbb{E}\Big[\min\{\delta^*,\, k_0 - \tilde{x}\} \cdot \mathbf{1}(\tilde{x} \in (0, k_0))\Big] > 0. \tag{19}$$

Remark 2. Equation (19) highlights a crucial distinction: the paper's positive carbon result critically relies on the *capacity substitution channel*. With fully fixed capacity, the Demand Response Program is carbon-neutral only if FCLs run exclusively on renewable surplus ($\tilde{x} < 0$), requiring additional design constraints on the demand response program (e.g., a renewable-only Demand Response Program that curtails computing when $\tilde{x} > 0$). Critics who point to high fossil utilization rates at mining facilities near gas plants are typically observing Case C environments where no capacity adjustment has occurred.

5.4 Summary table: Timing variations

Table 2 summarizes the carbon and profit outcomes across the three cases, and the key condition for carbon improvement.

[Table 2 about here.]

The progression from Case C to the baseline captures the key insight: carbon improve-

ment requires the capacity substitution channel. The more freely capacity can be adjusted, the more powerful the mechanism.

6 Endogenous Electricity Prices (Rate-of-Return Regulation)

The baseline model treats the retail price p as set exogenously by the regulator. In regulated markets the regulator typically sets p endogenously to ensure a *fair return* on investment via *rate-of-return* (ROR) regulation. This section embeds the model in the Averch–Johnson framework (Averch and Johnson, 1962).

6.1 The Averch–Johnson framework

Let $\bar{r}$ denote the allowed rate of return on the rate base, and ρ the utility's cost of capital, with $\bar{r} \geq \rho$. The *rate base* includes all capital investments: $\mathcal{B} = \gamma k + CK + \kappa\delta$. The regulator sets p such that the utility earns exactly $\bar{r}$ on its rate base:

$$p \cdot \mathbb{E}[D_{\text{served}}] + R \cdot \mathbb{E}[d(\tilde{x})] = \bar{r} \cdot (\gamma k + CK + \kappa\delta) + c \cdot \mathbb{E}[\text{fossil output}], \tag{20}$$

where $\mathbb{E}[D_{\text{served}}]$ is expected demand served. Solving for the endogenous price:

$$p(k, K, \delta) = \frac{\bar{r}\,\mathcal{B} + c \cdot \mathbb{E}[\text{fossil}] - R \cdot \mathbb{E}[d]}{\mathbb{E}[D_{\text{served}}]}. \tag{21}$$

6.2 The Averch–Johnson over-investment effect

The classic AJ result states that when $\bar{r} > \rho$, the regulated firm over-invests in capital relative to the cost-minimizing benchmark: additional capital raises the allowed rate base and hence the permitted price, relaxing the effective cost constraint.

In the present model, the AJ effect modifies the FOC for fossil capacity:

$$(\beta + p - c)\Big[1 - F(k_{\text{AJ}}; K)\Big] = \gamma(1 - \bar{r}). \tag{22}$$

Because $\bar{r} > 0$, the effective net cost of fossil capital is $\gamma(1 - \bar{r}) < \gamma$, implying $k_{\text{AJ}} > k^*$: the AJ-regulated utility builds more fossil capacity than the unregulated optimum. The same logic applies to renewable capacity: $K_{\text{AJ}} > K^*$ (over-investment in renewables under ROR). Accidentally, AJ over-investment benefits consumers through lower blackout probability.

6.3 Mining revenue and the double win

Under ROR regulation, FCL revenue has a direct pass-through effect on the retail price. From (21), a marginal increase in expected computing output lowers the required retail price:

$$\frac{\partial p}{\partial\, \mathbb{E}[d]} = \frac{-R}{\mathbb{E}[D_{\text{served}}]} < 0. \tag{23}$$

This is the *double win* of FCLs under ROR regulation:

1. **Capacity dividend**: FCL revenue subsidises capital investment, allowing more renewable and fossil capacity to be built (the paper's main channel in the baseline).
2. **Price dividend**: FCL revenue is passed through to consumers as a lower retail electricity price, benefiting households directly. At current calibration values ($R \approx$ \$120/MWh, computing approximately 10% of grid output), the implied price reduction is roughly \$12/MWh—material in regulated retail markets with average prices of \$50–120/MWh.

The double win strengthens the welfare case for allowing regulated utilities to participate in FCL programs. It also suggests that the socially optimal hardware investment under ROR regulation may exceed the utility's private optimum, creating a potential role for regulatory encouragement.

6.4 Rate-base classification of FCL hardware

A key regulatory question is whether FCL hardware (e.g., ASICs, GPU clusters) qualifies as "rate base." Rate base traditionally includes assets used to serve regulated consumers. If regulators *exclude* computing hardware from rate base (treating FCLs as an ancillary unregulated activity), the AJ over-investment effect applies only to K and k, and the hardware choice is governed by a separate unregulated maximization. This creates a two-tier optimization structure not present in the baseline model. Whether regulators include or exclude computing hardware from the rate base is currently an open question in most jurisdictions, with significant implications for the size of the double win.

7 Deregulated Competitive Electricity Markets

The paper's baseline assumes a vertically integrated regulated utility. Many electricity markets (ERCOT in Texas, PJM in the Mid-Atlantic, CAISO in California) are deregulated, with competitive generation and a system operator that clears spot markets. This section develops the appropriate competitive framework.

7.1 Spot price mechanism and FCL participation

In a competitive deregulated market, the system marginal price equals the short-run marginal cost of the marginal generator:

$$P(\tilde{x}) = \begin{cases} 0 & \text{if } \tilde{x} \leq 0 \quad \text{(renewable surplus)}, \\ c & \text{if } \tilde{x} \in (0, k) \quad \text{(fossil marginal)}, \\ P_{\text{cap}} & \text{if } \tilde{x} > k \quad \text{(scarcity)}. \end{cases} \tag{24}$$

The price cap P_{cap} represents either a regulatory ceiling or the value of lost load (VOLL). Fossil generators recover capital costs γ only through scarcity rents during shortage events.

FCL operators participate as price-responsive buyers: they run at full capacity when $P \leq R$ and shut down when $P > R$. Given spot prices (24): FCLs run when $\tilde{x} \leq 0$ (earning R per MWh at zero electricity cost) and when $\tilde{x} \in (0, k)$ (earning $R - c > 0$ per MWh), and shut down during scarcity events when $P_{\text{cap}} > R$. This provides downward demand response that partially alleviates scarcity.

7.2 Competitive equilibrium without FCL

Competitive fossil generators enter until expected profits are zero:

$$(P_{\text{cap}} - c) \cdot \Pr(\tilde{x} > k) = \gamma \quad \Longrightarrow \quad \Pr(\text{blackout}) = \frac{\gamma}{P_{\text{cap}} - c}. \tag{25}$$

This is structurally identical to (6) with P_{cap} playing the role of $\beta + p$. The competitive equilibrium replicates the regulated outcome when $P_{\text{cap}} = \beta + p$ (the price cap reflects the full social cost of lost load).

7.3 Competitive equilibrium with FCL

With FCL operators present, computing loads shut down whenever the spot price reaches P_{cap}, reducing effective demand during scarcity by δ. A blackout occurs when residual demand net of FCL curtailment exceeds fossil capacity: $\tilde{x} - \delta > k$, i.e., $\tilde{x} > k + \delta$. The zero-profit condition for fossil entry becomes:

$$(P_{\text{cap}} - c) \cdot \Pr\Big(\tilde{x} > k + \delta\Big) = \gamma, \tag{26}$$

implying a lower equilibrium blackout probability than (25) (since $\Pr(\tilde{x} > k+\delta) < \Pr(\tilde{x} > k)$ for $\delta > 0$). Simultaneously, the zero-profit condition for FCL investment accounts for revenue in both the fossil-dispatch region and the renewable-surplus region:

$$R \cdot \Pr(\tilde{x} \leq 0) + (R - c) \cdot \Pr(0 < \tilde{x} < k) = \kappa, \tag{27}$$

the competitive analogue of the regulated FOC (10), with the R term reflecting zero electricity cost during renewable surplus hours.

The "missing money" problem and FCLs. In energy-only deregulated markets without capacity payments, fossil generators may not recover their fixed costs if P_{cap} is set below VOLL—the so-called "missing money" problem (Joskow, 2008). FCL operators partially mitigate this: by buying electricity at the competitive price c during off-peak hours, they provide revenue that supports fossil capacity viability. Specifically, FCL purchases during hours $\tilde{x} \in (0, k)$ raise the effective load factor of fossil generators, improving their economics without raising the market price for consumers.

7.4 Carbon taxes and FCL programs as complements

A carbon tax τ per unit of fossil output raises the effective marginal cost to $c + \tau$, shifting the spot price during dispatch. The carbon tax interacts with FCLs in two ways:

1. **Direct effect**: The higher dispatch price $(c + \tau)$ reduces the FCL profit margin to $R - (c + \tau)$ during fossil dispatch hours. If $\tau > R - c$, FCL operators shut down whenever fossil is the marginal generator, and compute *only* on renewable surplus $(\tilde{x} \leq 0)$. This endogenously tilts computing toward zero-carbon hours without an explicit renewable mandate.
2. **Indirect effect**: The higher scarcity premium $P_{\text{cap}} - (c+\tau)$ improves revenue prospects

for fossil generators, attracting more entry and reducing blackout probability (partially offsetting the higher variable cost).

Carbon taxes and FCL demand response programs are therefore *complementary instruments*: a sufficiently high τ pushes computing exclusively onto renewable surplus, while tax revenues can fund renewable capacity subsidies.

7.5 Comparison with regulated market

Table 3 compares the key equilibrium conditions and FCL effects across the two market structures.

[Table 3 about here.]

8 Flexible Computing Loads: Bitcoin Mining and AI Training

The model developed above applies to any FCL technology. This section compares Bitcoin mining and AI training as specific implementations, identifying their structural similarities and the key differences that determine which technology provides more effective grid flexibility.

8.1 Bitcoin mining as the paradigm FCL

Bitcoin mining exhibits the three properties that make an ideal FCL:

Temporal flexibility. Bitcoin mining is a memoryless process: each microsecond of hashing delivers an identical expected reward, with no "memory" of previous computation. Switching ASICs off and back on imposes zero productivity cost. This makes Bitcoin min-

ing the ideally flexible load: curtailment is costless from the miner's perspective (modulo hardware wear), and computing intensity can be adjusted in real time without any state loss.

Revenue transparency. The per-MWh mining revenue R is observable, real-time, and publicly verifiable:

$$R = \frac{\text{block reward} \times P_{\text{BTC}}}{\text{total network energy per block}}.$$

Both numerator and denominator are published in real time by the Bitcoin network. This transparency is essential for Demand Response Program contract design: the utility and regulator can independently verify the revenue rate, set incentive-compatible pricing, and monitor program compliance.

Scalability and modularity. Bitcoin mining is modular at any scale: a Demand Response Program can enrol participants ranging from a single ASIC (~3 kW) to a 500 MW mining facility, providing smooth scaling of Demand Response Program capacity across a utility's service territory. This is not purely theoretical: programmes in Texas, Alberta, and several other jurisdictions demonstrate workable large-scale deployment.

8.2 AI training as an emerging FCL

We now analyse AI training as a potential FCL, in comparison to Bitcoin mining.

Similarities. Large-language-model (LLM) training runs last days to weeks and support checkpointing: pausing mid-run without losing progress, then resuming— analogous to powering down ASICs and restarting. A 50-hour annual curtailment (Duke Energy's estimated threshold for accommodating large data center loads without new generation) is imperceptible in a three-week training run. Hyperscalers (Google, Microsoft, Amazon) can also reroute

training workloads across data centers in different grid regions, providing spatial demand response unavailable to geographically fixed Bitcoin miners.

At scale, AI compute is comparable to Bitcoin mining: global data centers consumed approximately 415 TWh in 2024, projected to exceed 800 TWh by 2030 (IEA). U.S. AI-focused data centers alone are forecast to require 90–100 GW of additional capacity by 2030, comparable to Bitcoin's current 16 GW continuous draw.

The inference latency constraint. The most important structural difference between AI and Bitcoin mining is the *inference* workload. AI inference—serving live queries to products like ChatGPT—is latency-constrained: users expect responses within seconds, making inference non-curtailable. As inference increasingly dominates AI electricity consumption (already the majority in 2024–2025), the fraction of AI demand that is truly curtailable shrinks. Bitcoin mining has no latency constraint whatsoever.

Revenue opacity. Bitcoin mining revenue R is transparent, real-time, and market-determined. AI training revenue is the expected future commercial value of the model being trained—speculative, internal, and unverifiable by a grid operator or regulator. There is no direct observable per-MWh analogue of R for AI training, complicating Demand Response Program contract design.

Power surge behavior. GPU clusters exhibit rapid power fluctuations: as GPUs transition between compute phases, power draw can surge or drop 30–60% within milliseconds (IEA 2024 data center report). This creates frequency stability challenges absent from Bitcoin mining's near-constant ASIC power draw, and imposes additional costs on grid operators.

8.3 Model modification for AI loads

The paper's model accommodates AI by decomposing data center demand into a curtailable portion (training and batch workloads) of fraction $\theta \in (0, 1)$ and a must-run portion (inference, live services) of fraction $1 - \theta$:

$$d_{\text{AI}}(x) = \theta \cdot \min\{k - x,\, \delta_{\text{AI}}\}, \tag{28}$$

where δ_{AI} is total AI data center GPU/ASIC capacity. The effective per-MWh revenue rate for the Demand Response Program operator is θR_{AI} (only the curtailable fraction generates Demand Response Program revenue), and the viability threshold becomes:

$$\kappa_{\text{AI}} < \frac{(\theta R_{\text{AI}} - c)(\beta + p - c - \gamma)}{\beta + p - c}. \tag{29}$$

Grids with 50% inference workloads ($\theta = 0.5$) obtain roughly half the carbon benefit of pure Bitcoin mining, all else equal.

Comparison. The structural comparison reveals that Bitcoin mining dominates AI training as an FCL on three dimensions (temporal flexibility, revenue transparency, hardware modularity), while AI training potentially dominates on two dimensions (spatial flexibility via rerouting, and scale of available capacity if R_{AI} is high). The latter advantage is heavily conditioned on whether training revenue R_{AI} can be contracted and verified—a practical obstacle that does not apply to Bitcoin.

A grid operator choosing between Bitcoin mining and AI training as Demand Response Program participants should prefer Bitcoin mining when transparency and contract verifiability are paramount (regulated utilities, public interest programs), and may prefer AI if large-scale data centers offer credible interruptible service commitments and their inference fraction $1 - \theta$ is demonstrably low.

9 Quantitative Assessment

This section updates the model's calibration to current market conditions and quantifies the implications for the viability of FCL demand response programs.

9.1 Mining revenue R per MWh

Table 4 updates the mining revenue calibration from the model's original parameters to current (2026) values.

[Table 4 about here.]

Calculation. At \$95,000 $\times$ 3.125 = \$296,875 per block, and with global mining power approximately 16,300 MW (143 TWh/year), a 10-minute block consumes $16{,}300 \times 600/3600 \approx$ 2,717 MWh. Hence R = \$296,875/2,717 $\approx$ \$109/MWh at current conditions. Using R = \$120/MWh as a central estimate, our baseline calibration at \$70/MWh is roughly 40% below current market value.

9.2 Updated viability threshold $\bar{\kappa}$

Table 5 shows how the viability threshold scales with the updated mining margin $R - c$.

[Table 5 about here.]

Current Antminer S21+ machines (16.5 J/TH efficiency) cost \$25–35/TH, or approximately \$2,500–3,500/MW of compute. Amortized over a 2–3 year hardware lifecycle, $\kappa \approx$ \$1,000–1,750/MW-year—well within the revised threshold, confirming strong current viability.

9.3 Renewable and fossil capacity costs

Table 6 presents renewable and fossil capital cost data. Renewable CapEx has fallen roughly 40–50% since 2020, reducing the parameter C in the model and shifting the jointly optimal renewable capacity K^{**} sharply upward. This implies more surplus hours, more zero-carbon computing, and a stronger capacity substitution effect.

[Table 6 about here.]

9.4 Summary: Revised calibration assessment

All key market movements since our baseline calibration strengthen the model's conclusions. First, mining revenue R has approximately doubled, tripling the viable ASIC capital cost range $\bar{\kappa}$. Second, renewable CapEx C has fallen roughly 40%, making high-K^{**} scenarios economically realistic and dramatically increasing surplus hours. Third, ASIC efficiency has improved approximately 12% per year (J/TH), reducing effective κ. Together, the model's calibrated carbon reduction estimates likely understate current potential by a factor of 2–3.

A cross-sectional observation from the calibration exercise is also worth noting. The quantitative benefit of FCL adoption scales with the *correlation between residual demand $\tilde{x}$ and grid conditions*: in regions where renewable supply and consumer demand are poorly synchronized (high variance of $\tilde{x}$), the shock-absorbing role of mining is more valuable and the carbon reduction larger. Conversely, in grids where supply and demand move closely together, the FCL mechanism adds less value. This suggests that FCL programs will be most impactful in grids with volatile, intermittent renewables—precisely the environments where fossil backup costs are highest.

9.5 Implications for AI data center loads

For AI training with curtailable fraction $\theta = 0.5$ and training-compute revenue $R_{\text{AI}} =$ \$200/MWh (reflecting the commercial value of frontier model training), the effective margin is $\theta R_{\text{AI}} - c = 0.5 \times 200 - 45 =$ \$55/MWh—roughly comparable to the current Bitcoin mining margin. However, at current inferred R_{AI} values the opacity problem means this revenue cannot be directly contracted, making practical Demand Response Program implementation with AI training loads substantially more challenging than with Bitcoin mining.

10 Conclusion

This paper has developed and analyzed an analytical model of flexible computing loads (FCLs) as a mechanism for reducing carbon emissions through demand response in electricity markets. The core result is that, under economically plausible conditions, integrating FCLs via a demand response program simultaneously increases utility profit, reduces carbon emissions, and enables more renewable capacity investment—a triple win that critically relies on the capacity substitution channel.

Key findings. The main analytical results can be summarized as follows.

In the *baseline simultaneous optimization*, the mining preference condition $\kappa < \bar{\kappa} = (R - c)\phi^*$ is both necessary and sufficient for the triple win. The threshold $\bar{\kappa}$ is informative: it equals the per-MWh profit margin times the probability of non-blackout, reflecting the expected revenue per MW of computing hardware.

Timing variations reveal that the carbon improvement channel critically depends on the endogeneity of fossil capacity. With fixed legacy capacity, the demand response program raises the utilization rate of existing fossil plants and can increase carbon in the short run. The policy implication is that Demand Response Program programs should be evaluated

over a time horizon long enough for capacity adjustment to take place.

Rate-of-return regulation introduces a double win: computing revenue reduces the required retail electricity price under ROR pass-through, benefiting consumers directly in addition to subsidising renewable capacity. The Averch–Johnson over-investment in capital inadvertently benefits consumers through higher reliability, and computing revenue amplifies this effect.

Deregulated competitive markets replicate the regulated outcomes with the market's scarcity price replacing the regulatory penalty. Computing loads address the missing money problem facing fossil generators and are complementary to carbon taxes, which endogenously tilt computing toward zero-carbon hours.

Comparing Bitcoin mining and AI training, Bitcoin mining dominates on the three key flexibility dimensions: temporal flexibility (memoryless hashing), revenue transparency (publicly verifiable block rewards), and hardware modularity (scalable from kilowatts to gigawatts). AI training potentially offers higher revenue and spatial rerouting flexibility, but is constrained by inference latency and revenue opacity.

Quantitatively, current market conditions (Bitcoin near \$95,000, post-halving rewards, dramatically lower renewable CapEx) imply the viable range for FCL adoption is approximately three times wider today than under our baseline calibration.

Policy implications. The analysis suggests several policy actions. Utility commissions should consider streamlined rate treatment for demand response programs that integrate FCL operators, particularly given the double win under ROR regulation. Regulators should evaluate FCL program timelines over a multi-year horizon to allow the capacity substitution channel to operate. Carbon tax levels above the mining margin $R - c$ would endogenously restrict computing to renewable surplus hours, a simpler and more robust mechanism than prescriptive renewable-only Demand Response Program requirements. Finally, given the

difficulty of verifying AI training revenues, blockchain-based computing (Bitcoin mining or similar transparent-revenue loads) is currently the more practical FCL candidate for publicly accountable demand response programs.

Limitations and future research. Several important extensions remain. First, stochastic BTC prices (correlated or uncorrelated with electricity prices) introduce real-options considerations in both the FCL investment decision and the renewable capacity choice. Second, multi-grid competition with endogenous hash-rate pricing creates a strategic interaction not captured in the price-taking model. Third, a continuous-time model with capital adjustment costs and learning curves would capture the dynamic aspects of the renewable transition. Fourth, empirical validation—particularly using ERCOT panel data on Bitcoin miners' curtailment behavior and grid outcomes—would provide direct tests of the model's comparative-static predictions.

References

Acun, Bilge, et al. 2023. "Carbon-aware computing for AI training." *arXiv preprint arXiv:2302.08437*.

Aflaki, Sam, and Serguei Netessine. 2017. "Strategic Investment in Renewable Energy Sources: The Effect of Supply Intermittency." *Manufacturing Service Oper. Management*, 19(3): 489–507.

Agrawal, Vishal V, and Şafak Yücel. 2022. "Design of Electricity Demand-Response Programs." *Management Sci.*, 68(10): 7441–7456.

Aïd, René, Dylan Possamaï, and Nizar Touzi. 2022. "Optimal Electricity Demand Response Contracting with Responsiveness Incentives." *Math. Oper. Res.*, 47(3): 2112–2137.

Andonov, Aleksandar, and Joshua Rauh. 2024. "The Shifting Finance of Electricity Generation." National Bureau of Economic Research Working Paper 32733.

Arnosti, Nick, and S Matthew Weinberg. 2022. "Bitcoin: A natural oligopoly." *Management Sci.*, 68(7): 4755–4771.

Averch, Harvey, and Leland L. Johnson. 1962. "Behavior of the Firm Under Regulatory Constraint." *American Economic Review*, 52(5): 1052–1069.

Benetton, Matteo, Giovanni Compiani, and Adair Morse. 2023. "When Cryptomining Comes to Town: High Electricity-Use Spillovers to the Local Economy." National Bureau of Economic Research.

Bertucci, Charles, Louis Bertucci, Jean-Michel Lasry, and Pierre-Louis Lions. 2024. "A Mean Field Game Approach to Bitcoin Mining." *SIAM J. Financial Math.*, 15(3): 960–987.

Biais, Bruno, Christophe Bisiere, Matthieu Bouvard, and Catherine Casamatta. 2019. "The blockchain folk theorem." *Rev. Financ. Stud.*, 32(5): 1662–1715.

Bitcoin.comNews. 2019. "Chinese Bitcoin Miners Migrate North After Wet Season." *https://web.archive.org/web/20220427014036/https://news.bitcoin.com/chinese-bitcoin-miners-migrate-north-after-wet-season/*, Accessed on 10 February 2025.

Bitcoin Policy Institute. 2024. "How Bitcoin Mining Reduces Carbon Emissions." *https://www.btcpolicy.org/articles/bitcoin-mining-reduces-carbon-emissions*, Accessed on 10 February 2025.

BITMAIN Support. 2022. "Difference between Low Power Mode and Low Power Enhanced Mode." *https://web.archive.org/web/20220309092700/https://support.bitmain.com/hc/en-us/articles/360019738593-Difference-between-Low-Power-Mode-and-Low-Power-Enhanced-Mode*, Accessed on 10 February 2025.

Blockstream. 2022. "Instant energy demand from the Bitcoin network." *https://web.archive.org/web/20220401125045/https://layer1.com/*, Accessed on 24 August 2021.

Bloomberg. 2021. "Bitcoin 'Hash Crash' Rebound Points to Miners Plugging Back In." *https://www.bloomberg.com/news/articles/2021-09-03/bitcoin-hash-crash-rebound-points-to-miners-plugging-back-in*, Accessed on 10 February 2025.

CNBC. 2022. "Crypto Decoded Bitcoin Miners are Helping the Texas Grid Brace for Winter Storm Impact." *https://web.archive.org/web/20220203143819/https://www.cnbc.com/2022/02/03/winter-storm-descends-on-texas-bitcoin-miners-shut-off-to-protect-ercot.html*, Accessed on 10 February 2025.

Cointelegraph. 2020. "BTC Hash Rate Slumps Amid Seasonal Miner Migration in China." *https://cointelegraph.com/news/btc-hash-rate-slumps-amid-seasonal-miner-migration-in-china*, Accessed on 10 February 2025.

Cong, Lin William, Zhiguo He, and Jiasun Li. 2021. "Decentralized mining in centralized pools." *Rev. Financ. Stud.*, 34(3): 1191–1235.

Darmouni, Olivier, Clemens Lehrer, and Parinitha Sastry. 2025. "Financing Investment in Electricity."

Davis, Lucas W., and Catherine Wolfram. 2012. "Deregulation, Consolidation, and Efficiency: Evidence from US Nuclear Power." *American Economic Journal: Applied Economics*, 4(4): 194–225.

decrypt. 2020. "Bitcoin's Hash Rate Drops As China's Rainy Season Ends." *https://web.archive.org/web/20211117210216/https://decrypt.co/46601/bitcoin-hash-rate-drop-attributed-to-chinese-rainy-season*, Accessed on 10 February 2025.

Electric Power Research Institute. 2024. "Powering Data Centers: U.S. Energy System and Emissions Impacts of Growing Loads." Electric Power Research Institute (EPRI) White Paper.

Fabrizio, Kira R., Nancy L. Rose, and Catherine D. Wolfram. 2007. "Do Markets Reduce Costs? Assessing the Impact of Regulatory Restructuring on US Electric Generation Efficiency." *American Economic Review*, 97(4): 1250–1277.

Fattahi, Ali, Sriram Dasu, and Reza Ahmadi. 2023. "Peak-Load Energy Management by Direct Load Control Contracts." *Management Sci.*, 69(5): 2788–2813.

Forbes. 2021. "How This Billionaire-Backed Crypto Startup Gets Paid To Not Mine Bitcoin." *https://web.archive.org/web/20200524160256/https://www.forbes.com/sites/christopherhelman/2020/05/21/how-this-billionaire-backed-crypto-startup-gets-paid-to-not-mine-bitcoin/#7bdc51947596*, Accessed on 10 February 2025.

Forbes. 2022. "Why Does Bitcoin Use So Much Energy?" *https://www.forbes.com/advisor/investing/cryptocurrency/bitcoins-energy-usage-explained/*, Accessed on 10 February 2025.

Forbes. 2024. "The New Power Couple: Renewable Energy And Bitcoin Mining." *https://www.forbes.com/sites/beccabratcher/2024/09/03/the-new-power-couple-renewable-energy-and-bitcoin-mining/*, Accessed on 10 February 2025.

Halaburda, Hanna, and David Yermack. 2023. "Bitcoin Mining Meets Wall Street: A Study of Publicly Traded Crypto Mining Companies." National Bureau of Economic Research.

Ibañez, Juan Ignacio, and Alexander Freier. 2023. "Bitcoin's Carbon Footprint Revisited: Proof of Work Mining for Renewable Energy Expansion." *J. Chall.*, 14(3): 35.

International Energy Agency. 2024. "Electricity 2024: Analysis and forecast to 2026." IEA. See also the follow-up special report: Energy and AI (2024).

Jamison, Mark A. 2008. "Rate of Return Regulation." In *Encyclopedia of Law and Economics*. Cheltenham, UK:Edward Elgar. Public Utility Research Center, University of Florida.

Jiang, Ziheng, Haibin Lin, et al. 2024. "MegaScale: Scaling Large Language Model Training to More Than 10,000 GPUs."

Joskow, Paul L. 1997. "Restructuring, Competition and Regulatory Reform in the U.S. Electricity Sector." *Journal of Economic Perspectives*, 11(3): 119–138.

Joskow, Paul L. 2003. "Electricity Sector Restructuring and Competition: Lessons Learned." *Cuadernos de Economía*, 40(121): 548–558.

Joskow, Paul L. 2008. "Capacity payments in imperfect electricity markets: Need and design." *Utilities Policy*, 16(3): 159–170.

Kawaguchi, Kohei, Shunya Noda, and Junpei Komiyama. 2022. "Miners' Reward Elasticity and Stability of Competing Proof-of-Work Cryptocurrencies." Center for Open Science.

Kök, A Gürhan, Kevin Shang, and Şafak Yücel. 2018. "Impact of Electricity Pricing Policies on Renewable Energy Investments and Carbon Emissions." *Management Sci.*, 64(1): 131–148.

Kök, A Gürhan, Kevin Shang, and Şafak Yücel. 2020. "Investments in Renewable and Conventional Energy: The Role of Operational Flexibility." *Manufacturing Service Oper. Management*, 22(5): 925–941.

Layer1 Technologies. 2022. "Building Bitcoin Batteries." *https://web.archive.org/web/20220401125045/https://layer1.com/*, Accessed on 1 April 2022.

Lohmann, Timo, and Steffen Rebennack. 2017. "Tailored Benders Decomposition for a Long-Term Power Expansion Model with Short-Term Demand Response." *Management Sci.*, 63(6): 2027–2048.

Mora, Camilo, Randi L Rollins, Katie Taladay, Michael B Kantar, Mason K Chock, Mio Shimada, and Erik C Franklin. 2018. "Bitcoin emissions alone could push global warming above 2 C." *Nature Climate Change*, 8(11): 931–933.

Peng, Xiaoshan, Owen Q Wu, and Gilvan C Souza. 2024. "Renewable, Flexible, and Storage Capacities: Friends or Foes?" *Manufacturing Service Oper. Management*, 26(5): 1730–1749.

Peura, Heikki, and Derek W Bunn. 2015. "Dynamic Pricing of Peak Production." *Oper. Res.*, 63(6): 1262–1279.

Prat, Julien, and Benjamin Walter. 2021. "An equilibrium model of the market for bitcoin mining." *J. of Political Econom.*, 129(8): 2415–2452.

Radovanović, Ana, Ross Koningstein, et al. 2023. "Carbon-aware computing for datacenters." *IEEE Transactions on Power Systems*, 38(2): 1270–1281.

Sachs, Goldman. 2024. "Generational Growth: AI, Data Centers and the Coming US Power Demand Surge." Goldman Sachs Research.

Satchwell, Andrew, Natalie Mims Frick, Peter Cappers, Sanem Sergici, Ryan Hledik, Goksin Kavlak, and Glenda Oskar. 2025. "Electricity Rate Designs for Large Loads: Evolving Practices and Opportunities." Lawrence Berkeley National Laboratory Technical Brief.

The Washington Post. 2024. "Crypto Leaders Press Trump to Create Federal Bitcoin Reserve." *https://www.washingtonpost.com/business/2024/11/27/trump-strategic-bitcoin-reserve-plan/*, Accessed on 10 February 2025.

Train, Kenneth E. 1991. *Optimal Regulation: The Economic Theory of Natural Monopoly.* Cambridge, MA:MIT Press.

U.S. Senator Lummis. 2024. "Lummis Introduces Strategic Bitcoin Reserve Legislation." *https://www.lummis.senate.gov/press-releases/lummis-introduces-strategic-bitcoin-reserve-legislation/*, Accessed on 10 February 2025.

Wolak, Frank A. 2003. "Measuring Unilateral Market Power in Wholesale Electricity Markets: The California Market, 1998–2000." *American Economic Review*, 93(2): 425–430.

Yaish, Aviv, and Aviv Zohar. 2023. "Correct Cryptocurrency ASIC Pricing: Are Miners Overpaying?" Schloss Dagstuhl-Leibniz-Zentrum für Informatik.

Yaish, Aviv, Saar Tochner, and Aviv Zohar. 2022. "Blockchain Stretching & Squeezing: Manipulating Time for your Best Interest." 65–88.

A Proofs

A.1 Proof of Proposition 1

Step 1: Profit comparison. Since $(K^{**}, k^{**}, \delta^{**})$ is the jointly optimal solution, by revealed preference $\Pi^{**} \geq \Pi$ evaluated at any feasible point, including $(K^*, k^*, \delta = 0)$ (the no-FCL optimum). Therefore $\Pi^{**} \geq \Pi^*$. The inequality is strict if and only if $\delta^* > 0$, which holds iff the marginal profit of FCL capacity at $\delta = 0$ is positive: $(R - c) \cdot F(k^{**}; K^{**}) > \kappa$ (the infra-marginal dispatch probability exceeds $\kappa/(R-c)$, so the first unit of capacity is strictly profitable). Substituting (7): $(R-c) \cdot (\beta + p - c - \gamma - \kappa)/(\beta + p - R) > \kappa$, which rearranges to $\kappa < \bar{\kappa}$. Hence $\Pi^{**} > \Pi^*$ iff $\kappa < \bar{\kappa}$.

Step 2: Blackout probability comparison. The equilibrium blackout probabilities are:

$$1 - F(k^*; K^*) = \frac{\gamma}{\beta + p - c}, \qquad 1 - F(k^{**}; K^{**}) = \frac{c + \gamma + \kappa - R}{\beta + p - R}.$$

The difference $(1 - F^{**}) - (1 - F^*)$ equals

$$\frac{c + \gamma + \kappa - R}{\beta + p - R} - \frac{\gamma}{\beta + p - c} = \frac{\kappa(\beta + p - c) - (R - c)(\beta + p - c - \gamma)}{(\beta + p - R)(\beta + p - c)}.$$

The denominator is positive (both $\beta + p > R$ and $\beta + p > c$). The numerator is negative iff $\kappa < (R-c)(\beta + p - c - \gamma)/(\beta + p - c) = \bar{\kappa}$. Hence $1 - F(k^{**}; K^{**}) < 1 - F(k^*; K^*)$ iff $\kappa < \bar{\kappa}$. $\square$ (blackout probability is lower with FCL)

Step 3: Carbon comparison. Expected fossil fuel output (proportional to carbon) with the FCL program equals $k^{**} - \int_{-K^{**}}^{k^{**} - \delta^*} F(x; K^{**})\, \mathrm{d}x$ (derived by integrating the Demand Response Program dispatch rule over the distribution of $\tilde{x}$; see the simplified profit function in Appendix B of the underlying paper). Expected carbon without the FCL program is $\int_0^{k^*} (1 - F(x; K^*))\, \mathrm{d}x$. Under the uniform distribution, substituting $F(x; K) = (x + K)/(\bar{D} +$

K) and using the FOC solutions ($F(k^{**} - \delta^*; K^{**}) = \kappa/(R - c)$, $F(k^{**}; K^{**}) = (\beta + p - c - \gamma - \kappa)/(\beta + p - R)$, and the analogous no-FCL expressions) a direct calculation shows that the condition for carbon to fall reduces to $\kappa < \bar{\kappa}$, the same threshold as for profit and blackout probability. □

B Grid Scope and the Role of Local versus Global Mining

A critical modeling question is whether the utility's grid is large enough to affect the global hash market, determining whether the per-MWh mining revenue R should be treated as exogenous or endogenous.

Price-taking benchmark. Any single regional grid is small relative to global mining infrastructure. A 1 GW demand response program represents approximately 0.6% of the global mining fleet (roughly 160 GW in 2025). For a single grid acting alone, treating R as exogenous is therefore appropriate, and this is our baseline assumption. Under this assumption, the results of Sections 4–7 hold without modification.

Exogenous versus endogenous mining revenue. Two polar assumptions bracket the space of mining market structures. In the *exogenous-R* case (our baseline, maintained throughout Sections 4–7), the utility's FCL program serves only a small fraction of the global mining fleet, so changes in δ have negligible effect on global hash rate and R is determined by global Bitcoin markets outside the model. In the *endogenous-R* case (polar opposite), the utility is the sole provider to *all* miners worldwide. With competitive mining and a fixed total block reward $\bar{R}$ per period, the per-unit mining revenue is $R(\delta) = \bar{R}/\mathbb{E}[d(\tilde{x}; \delta)]$, so that aggregate mining income $R \cdot \mathbb{E}[d] = \bar{R}$ is constant, independent of program size.

The marginal mining-revenue incentive for FCL adoption therefore vanishes in the global monopolist case: $\partial\Big(R\,\mathbb{E}[d]\Big)/\partial\delta = 0$. Appendix C formalizes this contrast and shows that the global monopolist utility optimally sets $\delta^G = 0$.

The endogenous-R case can be further rationalized via the Tullock structure of miner competition. To see this, notice that with n symmetric oligopolistic miners each maximizing $R\,d_i/(d_i + \sum_{j\neq i} d_j) - d_i P$, the Nash equilibrium aggregate demand is

$$d = \frac{n-1}{n} \cdot \frac{R}{P}, \tag{30}$$

so total miner expenditure $Pd = \frac{n-1}{n}R$ is independent of the electricity price P. In the competitive limit $n \to \infty$, expenditure converges to R, still independent of P. The regulator can therefore implement the Demand Response Program dispatch rule (2) through contingent pricing without distorting the aggregate computing revenue.

Remark 3 (Revenue independence from local grid decisions)**.** In practice, R itself depends on factors external to the grid: the Bitcoin price, the block-subsidy schedule (which halves approximately every four years), and transaction fees. The regulator can adjust the electricity price P contingently on these external factors to hold the aggregate mining demand d fixed at the grid's desired dispatch level, regardless of Bitcoin market conditions. This is feasible because almost all mining is conducted at large farms with individual grid contracts, allowing the utility to price-discriminate between mining and non-mining loads and to work out contingent pricing schedules in advance (Kawaguchi, Noda and Komiyama, 2022).

Carbon accounting. The grid's carbon reduction is a local externality: the utility reduces its own fossil investment regardless of what miners on other grids do. No cross-jurisdictional coordination is required, making implementation feasible under existing regulatory structures.

Multi-grid coordination and strategic interaction. If many large grids simultaneously adopt Demand Response Programs, the aggregate increase in computing load raises global hash rate, which—since Bitcoin's difficulty adjustment targets a fixed block production rate—proportionally reduces the per-MWh revenue R for all participants. This creates a "mining commons" externality: each grid's Demand Response Program expansion imposes a negative revenue externality on other grids.

With N grids each running programs of size δ_i:

$$R \approx \frac{\text{block reward} \times P_{\text{BTC}}}{\left(\sum_i \delta_i + m\right) \alpha / H_0 \cdot \text{energy per block}},$$

where m is the hash rate of out-of-grid miners and α is hash rate per MW. A social planner would internalize this externality and prescribe lower individual Demand Response Program sizes. With U.S. mining accounting for roughly 38% of global hash rate as of 2024, North American grids acting in concert would have significant market power. Appendix C formalizes the Nash equilibrium across grid operators and characterizes the coordination failure.

The two polar cases—price-taking local grid and a single global grid operating all mining—bracket the range of realistic scenarios. Our baseline analysis covers the price-taking case; the global grid case is analyzed in Appendix C.

C Global Mining and Nash Equilibrium Extension

This appendix formalizes the multi-grid coordination problem when N grid operators simultaneously run FCL programs, endogenizing the per-MWh mining revenue R.

C.1 Setup

Let grid i operate an FCL program of size δ_i. Total network computing power is $H = (\sum_i \delta_i + m)\alpha$, where m is the hash rate of out-of-grid miners and α is hash rate per MW. The per-MWh mining revenue for grid-i miners is

$$R_i = \frac{\delta_i}{\delta_i + m_i + \sum_{j \neq i} \delta_j} \cdot R_0,$$

where R_0 is the total block reward revenue per block and m_i is grid i's out-of-grid mining allocation. In the Nash equilibrium each grid takes other grids' Demand Response Program sizes as given.

C.2 Nash equilibrium and coordination failure

Proposition 2 (Multi-grid Nash equilibrium)**.** *In the symmetric Nash equilibrium with N identical grids, each grid runs an FCL program of size δ^{NE} satisfying*

$$(R^{NE} - c) \cdot F(k^{**}; K^{NE}) = \kappa, \quad R^{NE} = \frac{R_0}{N\delta^{NE}/m + 1} \cdot \frac{1}{N},$$

where $\delta^{NE} < \delta^{SP}$ (the social planner's optimal program size), reflecting a tragedy of the commons in global hash markets.

Remark 4. The coordination failure is modest when grids are small relative to the global hash market ($N\delta^{NE} \ll m$), as is the case currently for most jurisdictions. It becomes quantitatively important if the U.S. as a whole acts as a single monopsonist buyer of global hash rate.

C.3 Single global utility: monopolist case

The sharpest case of endogenous R arises when a single utility ($N = 1$, no out-of-grid miners: $m = 0$) serves the entire global mining fleet. With total block reward $\bar{R}$ per period and competitive mining, the per-unit revenue is

$$R^G = \frac{\bar{R}}{\mathbb{E}[d(\tilde{x};\delta)]},$$

so aggregate mining income $R^G \mathbb{E}[d] = \bar{R}$ is a constant independent of the program size δ.

Proposition 3 (Global monopolist)**.** *When a single utility serves all global miners and mining is competitive, the marginal revenue from expanding the FCL program is zero:*

$$\frac{\partial}{\partial \delta}\Big[R^G(\delta)\,\mathbb{E}[d(\tilde{x};\delta)]\Big] = \frac{\partial \bar{R}}{\partial \delta} = 0.$$

The modified first-order condition for δ becomes

$$-c\,F(k-\delta;K) = \kappa,$$

which has no interior solution when $\kappa > 0$ and $c > 0$. The global monopolist therefore sets $\delta^G = 0$: no FCL program is offered.

Proof. In the exogenous-R case, the FOC for δ is $(R-c)\,F(k^{**}-\delta^*;K^{**}) = \kappa$ (equation (10)), where the term $R \cdot F(k^{**}-\delta^*;K^{**})$ is the marginal mining revenue from expanding the FCL program by one unit ($F(k^{**}-\delta^*;K^{**})$ is the probability of full FCL dispatch, which is the probability the extra unit is actually used). In the global monopolist case, the marginal mining revenue from δ is $\partial[R^G(\delta)\,\mathbb{E}[d(\tilde{x};\delta)]]/\partial\delta = \partial\bar{R}/\partial\delta = 0$, since total mining income $\bar{R}$ is independent of δ. Replacing the marginal mining revenue $R \cdot F$ with zero in the exogenous FOC leaves only the cost terms: $-c\,F(k^{**}-\delta^*;K^{**}) = \kappa$. Since $c, F, \kappa > 0$, this has no

solution, and the corner solution $\delta^G = 0$ is the unique optimum. □

Remark 5. With $\delta^G = 0$, the global monopolist's capacity choices (K^G, k^G) revert to the no-FCL benchmark (Appendix G.1). All three outcomes—blackout probability, carbon emissions, and utility profit—equal their no-FCL levels, and the "triple win" of Proposition 1 is never achieved.

Propositions 2 and 3 together characterize the full range: the FCL program is largest when grids are small relative to global mining (exogenous R, δ^{**}), shrinks as any single operator acquires market power ($\delta^{NE} < \delta^{**}$), and collapses to zero under global monopoly ($\delta^G = 0$).

C.4 Partial miner coverage: mixed in-grid and out-of-grid miners

An intermediate case—between the price-taking benchmark and the global monopolist—arises when a grid serves n in-grid miners competing against m out-of-grid miners who face an exogenous electricity price P_o. Let $\mathcal{N}$ and $\mathcal{M}$ denote the in-grid and out-of-grid miner sets, respectively. Each in-grid miner i solves

$$\max_{d_i} \; R \frac{d_i}{d_i + \sum_{j \neq i, j \in \mathcal{N}} d_j + \sum_{k \in \mathcal{M}} d_k} - d_i P,$$

while each out-of-grid miner faces price P_o. Aggregating first-order conditions across all miners yields total network demand

$$\sum_{j \in \mathcal{M} \cup \mathcal{N}} d_j = \frac{(m+n-1)R}{nP + mP_o}, \tag{31}$$

and summing only over in-grid miners gives equilibrium in-grid demand:

$$d \equiv \sum_{j \in \mathcal{N}} d_j = \frac{(m+n-1)R}{nP + mP_o} \cdot n \cdot \frac{mP_o - (m-1)P}{nP + mP_o}. \tag{32}$$

Hence total in-grid miner expenditure is

$$Pd = \frac{(m+n-1)R}{n+mP_o/P} \cdot n \cdot \frac{mP_o-(m-1)P}{nP+mP_o}.$$

When P is close to P_o, this is approximately exogenous (independent of the grid's capacity decisions), recovering the price-taking benchmark. In the limit $m/n \to \infty$ (the focal grid is negligible relative to the global mining fleet),

$$d \approx \frac{nR}{P_o}\left(1-\frac{P}{P_o}\right), \qquad Pd \approx nR\,\frac{P}{P_o}\left(1-\frac{P}{P_o}\right),$$

which is maximized at $P = P_o/2$. Importantly, $d \geq 0$ requires $P \leq P_o$; if the grid sets $P > P_o$ all miners relocate.

Remark 6. Equation (32) highlights a complication absent from the main text: when the focal grid is not negligible (finite m/n), total miner revenue Pd is no longer independent of P, so the simplification that mining revenue is exogenous to grid decisions fails. The main-text results remain valid whenever the grid is small relative to the global fleet—the empirically relevant case today—but equation (32) characterizes the correction term for intermediate cases and clarifies the need for multi-grid coordination as FCL programs proliferate.

C.5 Policy implication

The Nash equilibrium each grid reaches overestimates the collective carbon benefit by ignoring the negative externality on other grids' revenues. An international coordination mechanism (e.g., agreed limits on FCL program sizes) could improve on the Nash outcome. The price-taking assumption in the main text is accurate for any single grid in the current environment; as FCL programs proliferate, this assumption should be revisited.

D The Case $R \leq c$

The main text assumes $R > c$ (computing revenue per MWh exceeds fossil variable cost), ensuring that FCL operators willingly consume fossil electricity. This appendix analyzes the boundary case $R = c$ and the unlikely case $R < c$.

D.1 Case $R = c$

When $R = c$, the FCL operator is exactly indifferent between running on fossil electricity and not running during fossil-dispatch hours. The FOC for δ (10) becomes $(R - c) \cdot F = 0 = \kappa$, which is satisfied only if $\kappa = 0$ (free ASIC hardware). In practice, ASIC hardware always has positive cost, so $R = c$ implies no FCL investment. The Demand Response Program program reduces to a pure renewable surplus absorption mechanism (FCLs run only when $\tilde{x} < 0$), which has no effect on fossil capacity decisions but increases renewable investment.

D.2 Case $R < c$

When $R < c$, computing is unprofitable during fossil dispatch hours. No profit-maximizing FCL operator would accept electricity priced at c or above. The FCL program can still be operated if the utility *subsidises* computing during fossil hours (paying miners $c - R$ per MWh to run), but this changes the economics fundamentally: the Demand Response Program now costs the utility money during fossil hours instead of earning revenue. The revised FOC for fossil capacity becomes:

$$(\beta + p + c - R)\Big[1 - F(k; K)\Big] = \gamma,$$

implying that subsidised FCL programs lead to more fossil capacity built (not less), raising the threshold for the utility's preference condition.

The case $R < c$ therefore represents a qualitatively different regime in which the carbon reduction mechanism is reversed: the utility would need to pay FCL operators to curtail, which is equivalent to a conservation or efficiency program rather than a demand response program. We rule out this case in the main text by Assumption 1.

E Alternative Regulatory Objectives

The main text assumes a profit-maximizing regulated utility that internalizes a blackout penalty imposed by the regulator. This appendix considers two alternative regulatory objective functions that are common in practice: (1) minimizing the probability of blackout and (2) minimizing the retail electricity price.

E.1 Objective 1: Minimizing blackout probability

Suppose the regulator (or the utility under explicit mandate) aims to minimize blackout probability $\Pr(\tilde{x} > k; K) = 1 - F(k; K)$ subject to a profit constraint $\Pi \geq \bar{\Pi}$.

Without FCL. Maximizing $F(k; K)$ subject to the profit constraint yields a corner solution: the utility invests the maximum fossil capacity that breaks even. The resulting blackout probability is lower than under profit maximization for any given K, since reliability is the direct objective.

With FCL. Introducing the FCL program changes the profit constraint: computing revenue relaxes it, allowing a higher k (more fossil backup) to be built while remaining profitable. The FCL program therefore reduces blackout probability through two channels: (i) direct demand reduction during scarcity events (FCLs shut down when $\tilde{x} > k$, exactly analogous to the baseline), and (ii) indirect capacity expansion (higher permitted fossil investment from the relaxed profit constraint).

Proposition 4 (Blackout-minimizing objective)**.** *Under a blackout-minimizing objective with binding profit constraint, allowing FCLs unambiguously reduces blackout probability whenever $\kappa < \bar{\kappa}$. The threshold $\bar{\kappa}$ is unchanged from Proposition 1.*

This result suggests that the triple win (profit, carbon, reliability) survives even when reliability is the primary objective—as long as the profit constraint remains binding.

E.2 Objective 2: Minimizing retail electricity price under ROR regulation

Under ROR regulation, the regulator sets the retail price to allow a fair return. Alternatively, the regulator may aim to minimize the retail price subject to the utility earning at least the cost of capital.

From (21), the retail price is

$$p = \frac{\bar{r}\,\mathcal{B} + c \cdot \mathbb{E}[\text{fossil}] - R \cdot \mathbb{E}[d]}{\mathbb{E}[D_{\text{served}}]}.$$

Minimizing p subject to the rate-base constraint gives the regulator an incentive to maximize $R \cdot \mathbb{E}[d]$—the total computing revenue the utility earns. This creates a direct regulatory preference for larger FCL programs, since each additional MWh of computing revenue is directly passed through to consumers.

Proposition 5 (Price-minimizing objective under ROR)**.** *Under a price-minimizing ROR regulatory objective, the regulator prefers the largest FCL program consistent with the utility's hardware investment constraint. The socially optimal FCL program size exceeds the utility's private optimum whenever $\partial p/\partial \delta < 0$, which holds whenever $R > 0$ and $\mathbb{E}[D_{served}] > 0$.*

Implication for rate-base treatment. Proposition 5 provides a rationale for regulators to *include* FCL hardware in the rate base: doing so both encourages over-investment via the

AJ mechanism (Section 6) and passes the resulting computing revenue through to consumers via lower prices. The net effect on consumer welfare is unambiguously positive as long as the computing margin $R - c > 0$ is maintained.

E.3 Objective 3: Social welfare maximization

For completeness, we note that social welfare maximization—which sums consumer surplus, utility profit, and environmental externality—can be decomposed into the profit-maximization problem plus a side condition on the externality. When the regulator requires the utility to internalize the blackout penalty β per MWh of unserved demand and the carbon cost θc per unit of fossil output, the social welfare problem reduces to the baseline utility problem with modified parameters. All results of Sections 4–7 continue to hold with the environmental cost adjustments.

F Case b: Small FCL Capacity ($\delta^{**} \leq K^{**}$)

The main text analyzes the *large-FCL* regime in which optimal FCL capacity exceeds renewable capacity ($\delta^* > K^{**}$, Case c). In this regime the full-capacity dispatch region $(k^{**} - \delta^{**}, k^{**})$ extends below $-K^{**}$, so whenever the FCL runs at full capacity it does so with some fossil power, and the net per-MWh margin is always $R - c$.

When $\delta^* \leq K^{**}$ (Case b), there exist *pure renewable-surplus hours*: states $\tilde{x} \leq -\delta^*$ in which the grid is in such deep surplus that even a fully dispatched FCL of size δ^* leaves renewable power uncurtailed. During those hours the fossil plant is idle, yet the FCL still runs at full capacity δ^* earning R per MWh with zero fossil variable cost—so the net margin is R, not $R - c$.

F.1 Profit function and first-order conditions

The fossil variable cost integral changes because the plant runs only for $\tilde{x} \in (-\delta^*,\, k-\delta^*)$ (not $[-K, k-\delta^*]$ as in Case c). Expected fossil output in Case b is $k - \int_{-\delta^*}^{k-\delta^*} F(x;K)\,dx = \int_{-\delta^*}^{k-\delta^*}(1-F(x;K))\,dx$. The profit function is:

$$\begin{aligned}\Pi^b =& p\,\mathbb{E}[\tilde{D}] + R\int_{k-\delta}^{k} F(x;K)\,\mathrm{d}x - c\int_{-\delta}^{k-\delta}\Big(1-F(x;K)\Big)\,\mathrm{d}x \\ &- (\beta+p)\int_{k}^{\bar{D}}\Big(1-F(x;K)\Big)\,\mathrm{d}x - \gamma k - CK - \kappa\delta.\end{aligned}$$

Differentiating with respect to δ (Leibniz rule, both limits move):

$$(R-c)\,F(k-\delta;\,K) + c\,F(-\delta;\,K) = \kappa. \tag{33}$$

The extra term $cF(-\delta;K) > 0$ relative to the Case c FOC reflects the marginal pure-renewable-surplus hour: a unit increase in δ^* is dispatched during an additional $f(-\delta^*)d\delta$ hours when no fossil runs, earning R per MWh rather than the $R-c$ margin of the fossil-dispatch region.

Using (33) to substitute out $F(k-\delta;K)$ in the FOC for k, one obtains:

$$F(k;\,K) = \frac{\beta+p-c-\gamma-\kappa+c\,F(-\delta;\,K)}{\beta+p-R}. \tag{34}$$

The extra $cF(-\delta;K) > 0$ in the numerator raises $F(k^b;K^b)$ above the Case c value, implying *more* fossil capacity and *lower* blackout probability for the same parameter values.

The expected carbon emission in Case b is proportional to $c\Big(k - \int_{-\delta^*}^{k-\delta^*} F(x;K)\,dx\Big)$; the lower integration limit is $-\delta^*$ instead of $-K^{**}$.

F.2 Blackout probability and the triple win

From (34), the equilibrium blackout probability under Case b is:

$$1 - F(k^b;\, K^b) = \frac{c + \gamma + \kappa - R - c\,F(-\delta^*;\, K^{**})}{\beta + p - R}. \tag{35}$$

Since $cF(-\delta^*; K^{**}) > 0$, this is strictly lower than the Case c blackout probability $(c+\gamma+\kappa-R)/(\beta+p-R)$ for the same parameters. Case b therefore delivers an even larger reliability improvement than Case c.

The adoption condition ($\delta^* > 0$) requires the marginal profit of the first unit of FCL capacity to be positive. Evaluating at $\delta = 0$ using (33):

$$(R-c)\,F(k;\, K)\Big|_{\delta=0} + c\,F(0;\, K)\Big|_{\delta=0} > \kappa.$$

Since $cF(0;K) > 0$, the adoption threshold is *weakly larger* in Case b than in Case c: the pure-renewable-hour revenue expands the viable ASIC cost range. Denoting the Case b adoption threshold by $\bar{\kappa}^b$,

$$\bar{\kappa}^b > \bar{\kappa} = (R-c)\,\phi^*,$$

so grids with high renewable capacity (large $F(0;K)$, many surplus hours) benefit more from FCL programs and can support higher ASIC hardware costs.

F.3 When does Case b arise?

Case b prevails at the optimum when $\delta^* \le K^{**}$. Under the uniform distribution $F(x;K) = (x+K)/(\bar{D}+K)$ on $[-K, \bar{D}]$, using (12) the Case c tentative solution gives:

$$\delta^* \le K^{**} \quad \Longleftrightarrow \quad k^{**} \le \frac{\kappa\,(\bar{D} + K^{**})}{R-c}.$$

Since k^{**} is increasing in scarcity costs $\beta + p$ and decreasing in κ, Case b is the relevant regime when *FCL hardware is cheap and scarcity costs are moderate*—the utility builds a large FCL relative to its renewable fleet. Case c (the main-text regime) applies when k^{**} is large, corresponding to high scarcity costs or high ASIC cost (low δ^* relative to K^{**}). For the current empirical calibration (Section 9), the quantitative check confirms that Case c applies for all baseline parameter values.

G Comprehensive Summary of Results

This appendix consolidates all major analytical results—optimal capacity allocations, blackout probabilities, and carbon comparisons—across every scenario combination analyzed in the paper. Closed-form expressions assume the uniform distribution $F(x; K) = (x + K)/(\bar{D} + K)$ on $[-K, \bar{D}]$ unless otherwise noted. Throughout, the triple-win threshold is $\bar{\kappa} = (R - c)(\beta + p - c - \gamma)/(\beta + p - c)$, and Assumption 1 ($\beta + p > R > c$, $\beta + p - c > \gamma$) is maintained.

G.1 No-FCL Benchmark

Table 7 records the benchmark optimal allocations and carbon outcomes when no FCL program is offered, corresponding to Section **??**.

[Table 7 about here.]

G.2 With FCL Program: Case c (Large FCL, $\delta^{**} > K^{**}$)

Case c is the baseline analyzed in Section 4. Tables 8 and 9 cover the interior and corner regimes respectively.

[Table 8 about here.]

[Table 9 about here.]

G.3 With FCL Program: Case b (Small FCL, $\delta^{**} \leq K^{**}$)

Full derivations are in Appendix F. Table 10 summarizes the first-order conditions, blackout probability, and triple-win threshold under Case b.

[Table 10 about here.]

G.4 Timing Variations: Pre-Installed Capacity (Section 5)

Table 11 collects the optimal allocations, blackout probabilities, and carbon outcomes for each of the three pre-commitment cases (A, B, C) alongside the simultaneous baseline.

[Table 11 about here.]

G.5 Rate-of-Return Regulation (Section 6)

Under ROR regulation the regulated price p is replaced by an endogenously determined price that allows the utility to earn a fair return $\bar{r}$ on rate base. The Averch–Johnson (AJ) over-investment effect raises the optimal k above the unregulated level. Table 12 summarizes how regulation affects the key first-order conditions, the triple-win threshold, and carbon emissions.

[Table 12 about here.]

G.6 Deregulated Competitive Markets (Section 7)

In deregulated markets, p is replaced by a spot price $P(\tilde{x})$. A scarcity price P_{cap} (capacity payment) plays the role of $\beta + p$ in the regulated model. All regulated-market results

are isomorphic with $P_{\text{cap}} \leftrightarrow \beta + p$. Table 13 lays out the zero-profit conditions, blackout probabilities, and triple-win threshold under competitive market equilibrium.

[Table 13 about here.]

G.7 Grid Scope: Local vs. Global Hash Market (Section B)

Table 14 contrasts the key assumption on mining revenue R and its implication for FCL program design across the single local grid, multi-grid Nash, social planner, and global monopolist scenarios.

[Table 14 about here.]

G.8 Master Comparison Table

Table 15 provides a unified comparison of blackout probabilities and the triple-win condition across all main scenarios, evaluated in the interior regime at the simultaneous optimum unless noted. In each cell, the adoption condition for the triple win is the scenario-specific threshold $\bar{\kappa}^{\text{scenario}}$.

[Table 15 about here.]

Tables

Table 1: Effects of the FCL Demand Response Program: Baseline Results

Outcome	**When $\kappa < \bar{\kappa}$ (adopted)**	**When $\kappa > \bar{\kappa}$ (not adopted)**
FCL capacity δ	$\delta^{**} > 0$	$\delta^{**} = 0$
Renewable capacity K	Increases ($K^{**} > K^{*}$)	No change
Blackout probability $1 - F(k; K)$	Decreases	Increases
Carbon emissions	Decreases	Increases
Utility profit	Increases	Decreases

Notes: This table reports the directional effects of introducing the FCL demand response program when all capacities (K, k, δ) are jointly optimized (Case c, interior regime; Assumption 1). The left column describes outcomes when the hardware cost κ falls below the triple-win threshold $\bar{\kappa} \equiv (R - c)(\beta + p - c - \gamma)/(\beta + p - c)$, so that the utility finds it optimal to offer $\delta^{**} > 0$. The right column describes outcomes if the program were forced upon a utility for which $\kappa > \bar{\kappa}$. The triple win—reliability, carbon, and profit all improve simultaneously—holds if and only if $\kappa < \bar{\kappa}$. All comparisons are relative to the no-FCL benchmark (Appendix G.1).

Table 2: Timing Variations: Carbon and Profit Effects by Capacity Flexibility

Case	**Carbon effect**	**Profit effect**	**Condition for net carbon reduction**
A: Fixed k_0, free K, δ	Ambiguous	Positive	$(R-c)\left\|\partial_K \mathbb{E}[\text{fossil}]\|_{K_0}\right\| > C$
B: Fixed K_0, free k, δ	Decreases iff $\kappa < \bar{\kappa}(K_0)$	Positive iff $\kappa < \bar{\kappa}(K_0)$	K_0 large $\Rightarrow$ easier; K_0 small $\Rightarrow$ harder
C: Fixed k_0, K_0; free δ only	Increases (unambiguous)	Positive	Not achievable without renewable-only Demand Response Program constraint
Baseline (§4): All free	Decreases iff $\kappa < \bar{\kappa}$	Positive iff $\kappa < \bar{\kappa}$	$\kappa < \bar{\kappa} = (R-c)\phi^*$; also lowers blackout prob

Notes: This table summarizes the direction of carbon and profit changes when an FCL demand response program is introduced, across four capacity-flexibility scenarios. Case A: fossil capacity k_0 is pre-installed; renewable K and FCL capacity δ are freely optimized. Case B: renewable capacity K_0 is pre-installed; k and δ are freely optimized. Case C: both k_0 and K_0 are pre-installed; only FCL capacity δ is chosen. Baseline: all three capacities are jointly optimized (Section 4). The condition for net carbon reduction in Case A requires that mining-induced renewable investment outweighs the higher fossil utilization. The threshold $\bar{\kappa}(K_0) \equiv (R-c) \cdot F(k^*; K_0)$ rises with K_0, so grids with more renewable capacity face a wider viability range.

Table 3: Regulated versus Deregulated Markets: FCL Equilibrium Conditions

Feature	**Regulated (§4–§6)**	**Deregulated (§7)**
Blackout probability w/o FCL	$\gamma/(\beta + p - c)$	$\gamma/(P_{\text{cap}} - c)$
Blackout probability w/ FCL	$(c + \gamma + \kappa - R)/(\beta + p - R)$ [lower iff $\kappa < \bar{\kappa}$]	$\gamma/(P_{\text{cap}} - c)$ [lower: FCL curtails at scarcity, $\tilde{x} - \delta$]
FCL viability condition	$\kappa < \bar{\kappa} = (R - c)\phi^*$	$R \Pr(\tilde{x} \leq 0) + (R - c) \Pr(0 < \tilde{x} < k) = \kappa$
Carbon channel	Capacity substitution + utilization	utilization + demand reduction
Retail price	Exogenous (p) or ROR-regulated	Spot price $P(\tilde{x})$
FCL revenue	R per MWh from Demand Response Program	$R - c$ per MWh during fossil hours, R on surplus
"Missing money" problem	Not applicable (capacity market)	Partially solved by FCL load factor
Carbon tax complementarity	Carbon tax can substitute for ROR price floor	Carbon tax tilts FCL toward renewable surplus

Notes: This table compares equilibrium conditions and FCL effects across the regulated (Sections 4–6) and deregulated competitive (Section 7) market structures. In the regulated case, β is the value-of-lost-load penalty and p is the exogenous retail price; in the deregulated case, P_{cap} is the competitive scarcity price cap (or capacity payment). Results are structurally isomorphic: replacing $\beta + p$ with P_{cap} maps all regulated expressions to their deregulated counterparts. $\phi^* \equiv F(k^{**} - \delta^*; K^{**})$ is the probability that fossil capacity is not fully loaded at the optimum.

Table 4: Mining Revenue R per MWh: Original Calibration versus Current (2026)

Parameter	**Original calibration**	**Current (2026)**
BTC price	\$22,000	≈\$95,000
Block reward	6.25 BTC	3.125 BTC
Network hash rate	∼400 EH/s	∼800–976 EH/s
Energy per block (average)	≈400 MWh	≈350–400 MWh
R (revenue/MWh)	≈\$70	≈\$100–157

Notes: This table updates the key revenue parameter R (mining income per MWh of electricity consumed) from our baseline to current (2026) market conditions. R is computed as block reward value divided by global energy consumed per block: at \$95,000 × 3.125 BTC per block and global mining power of approximately 16,300 MW, a 10-minute block consumes ∼2,717 MWh, giving $R \approx$ \$109/MWh at the central estimate. The model uses $R =$ \$120/MWh as the calibrated central value.

Table 5: Updated Triple-Win Threshold $\bar{\kappa}$: Original versus Current

Parameter	**Original**	**Updated**
R (mining revenue/MWh)	\$70	\$120
c (fossil variable cost)	\$45	\$45 (stable)
$R - c$ (mining margin)	\$25/MWh	\$75/MWh
$\bar{\kappa}$ (ASIC threshold)	≈\$15/MW-h	≈\$45/MW-h
Viability range in κ	narrow	3× wider

Notes: This table shows how the triple-win threshold $\bar{\kappa} \equiv (R-c)(\beta+p-c-\gamma)/(\beta+p-c)$ changes when the mining revenue parameter R is updated to 2026 conditions. The threefold increase in the mining margin $R - c$ (from \$25 to \$75/MWh) implies a threefold expansion of the range of hardware costs κ for which the FCL program is socially beneficial. All other parameters held at baseline values ($\beta + p = \$200$/MWh, $\gamma = \$10$/MW-h).

Table 6: Renewable and Fossil Capacity Costs: 2020 versus 2025–26

Parameter	**Typical 2020**	**Current (2025–26)**
Solar PV CapEx (utility scale)	\$800–1,200/kW	\$388–600/kW
Onshore wind CapEx	\$1,200–1,600/kW	\$850–1,000/kW
CCGT gas CapEx	\$1,000–1,200/kW	\$800–1,000/kW
Solar LCOE	\$40–60/MWh	\$28–66/MWh
Onshore wind LCOE	\$35–55/MWh	\$25–70/MWh

Notes: This table presents representative capital expenditure (CapEx) ranges and levelized cost of energy (LCOE) for utility-scale solar PV, onshore wind, and combined-cycle gas turbines (CCGT), comparing 2020 benchmarks to current (2025–26) values. Renewable CapEx has declined approximately 40–50% since 2020, corresponding to a fall in the model parameter C (renewable capital cost per MW). A lower C shifts the jointly optimal renewable capacity K^{**} upward, increasing surplus hours and strengthening the capacity-substitution channel. Sources: IRENA, NREL, Lazard LCOE.

Table 7: No-FCL Benchmark: Optimal Allocations and Outcomes

Quantity	**General**	**Uniform distribution**
Fossil capacity k^*	$F(k^*;K^*)=\dfrac{\beta+p-c-\gamma}{\beta+p-c}$	$k^*=\dfrac{\beta+p-c-\gamma}{\beta+p-c}(\bar{D}+K^*)-K^*$
Renewable capacity K^*	$C=(\beta+p)\int_{k^*}^{\bar{D}}\frac{\partial F}{\partial K}\,\mathrm{d}x+c\int_0^{k^*}\frac{\partial F}{\partial K}\,\mathrm{d}x$	Determined by implicit eq. above
Blackout probability	$1-F(k^*;K^*)=\dfrac{\gamma}{\beta+p-c}$	Same (exact)
Expected carbon	$c\cdot\mathbb{E}[\min\{k^*,\,\tilde{x}^+\}]$	$c\cdot\int_0^{k^*}(1-F(x;K^*))\,\mathrm{d}x$

Notes: This table reports optimal fossil capacity k^*, renewable capacity K^*, blackout probability, and expected carbon emissions when no FCL demand response program is offered. Results are given for a general residual-demand distribution $F(x;K)$ (with $\partial F/\partial K>0$) and for the uniform distribution $F(x;K)=(x+K)/(\bar{D}+K)$ on $[-K,\bar{D}]$. The blackout probability $\gamma/(\beta+p-c)$ is distribution-free. Assumption 1 ($\beta+p>c+\gamma$, $\beta+p>R>c$) is maintained. This benchmark serves as the comparison point for all tables with an FCL program.

Table 8: With FCL Program, Case c Interior Regime: Simultaneous Optimum

Quantity	**General**	**Uniform distribution**
Fossil capacity k^{**}	$F(k^{**}; K^{**}) = \dfrac{\beta + p - c - \gamma - \kappa}{\beta + p - R}$	$k^{**} = \dfrac{\beta + p - c - \gamma - \kappa}{\beta + p - R}(\bar{D} + K^{**}) - K^{**}$
Renewable capacity K^{**}	FOC (9): weight R in FCL region	Determined jointly with k^{**}
FCL capacity δ^*	$(R - c)F(k^{**} - \delta^*; K^{**}) = \kappa$	$\delta^* = k^{**} - \dfrac{\kappa(\bar{D} + K^{**})}{R - c} + K^{**}$
Blackout probability	$\dfrac{c + \gamma + \kappa - R}{\beta + p - R}$	Same (exact)
Comparison to no-FCL	Lower iff $\kappa < \bar{\kappa}$	Same
Carbon comparison	Falls iff $\kappa < \bar{\kappa}$	$\bar{\kappa} = \dfrac{(R - c)(\beta + p - c - \gamma)}{\beta + p - c}$
Profit comparison	Rises iff $\kappa < \bar{\kappa}$	Same
Triple win	$\kappa < \bar{\kappa} \Leftrightarrow$ all three improve	Proposition 1

Notes: This table reports optimal allocations and welfare comparisons under Case c ($\delta^{**} > K^{**}$, the empirically relevant regime) and the interior regime ($R < c + \gamma + \kappa$), when K, k, and δ are all chosen simultaneously. Closed-form expressions use the uniform distribution $F(x; K) = (x + K)/(\bar{D} + K)$ on $[-K, \bar{D}]$. The triple-win threshold is $\bar{\kappa} \equiv (R - c)(\beta + p - c - \gamma)/(\beta + p - c)$; the blackout probability $\gamma/(\beta + p - c)$ in the no-FCL benchmark (Table 7) exceeds $(c + \gamma + \kappa - R)/(\beta + p - R)$ iff $\kappa < \bar{\kappa}$. All comparisons are relative to the no-FCL benchmark.

Table 9: With FCL Program, Case c Corner Regime: Simultaneous Optimum

Quantity	**Result**
Fossil capacity k^{**}	$k^{**} = \bar{D}$ (corner solution)
Blackout probability	0 (zero blackout)
Triple win	Automatically satisfied: $\kappa \leq R - c - \gamma < \bar{\kappa}$
Carbon	Strictly lower than no-FCL (zero fossil waste hours)
Profit	Strictly higher

Notes: This table reports the simultaneous optimum under Case c and the corner regime, defined by $R \geq c + \gamma + \kappa$. In this regime the utility builds fossil capacity equal to the maximum demand $\bar{D}$, eliminating blackouts entirely. The triple win is automatically satisfied whenever the corner regime holds, since the adoption condition $\kappa \leq R - c - \gamma$ is strictly stronger than the interior-regime threshold $\bar{\kappa}$. The interior regime ($R < c + \gamma + \kappa$, Table 8) is the empirically relevant case for moderate ASIC costs.

Table 10: With FCL Program, Case b ($\delta^{**} \leq K^{**}$): Simultaneous Optimum

Quantity	**Case b**	**Comparison to Case c**
FOC for δ	$(R - c)F(k - \delta; K) + cF(-\delta; K) = \kappa$	Extra term $cF(-\delta; K) > 0$
FOC for k	$F(k; K) = \frac{\beta+p-c-\gamma-\kappa+cF(-\delta^*;K^{**})}{\beta+p-R}$	Numerator larger by $cF(-\delta^*; K^{**})$
Blackout probability	$\frac{c+\gamma+\kappa-R-cF(-\delta^*;K^{**})}{\beta+p-R}$	**Strictly lower** than Case c
Adoption threshold $\bar{\kappa}^b$	$\bar{\kappa}^b > \bar{\kappa}$	Wider viable range
Condition for Case b	$k^{**} \leq \frac{\kappa(\bar{D}+K^{**})}{R-c}$ (uniform)	Small κ or small K^{**}

Notes: This table summarizes the first-order conditions and key comparisons under Case b ($\delta^{**} \leq K^{**}$), the regime in which the optimal FCL capacity is small enough that FCLs run at full capacity δ only during renewable-surplus hours ($x \leq 0$). In Case b, FCLs consume electricity even during periods of renewable surplus, adding an extra term $cF(-\delta; K) > 0$ to the FOC for δ. This extra term lowers the equilibrium blackout probability below the Case c level and widens the viable range of hardware costs to $\bar{\kappa}^b > \bar{\kappa}$. Full derivations appear in Appendix F. The interior-regime condition ($R < c + \gamma + \kappa$) is maintained.

Table 11: Timing Variations: Optimal Allocations and Carbon Outcomes by Capacity Flexibility

Case	**Optimal k, δ**	**Blackout**	**Carbon**
A: k fixed at k_0, free K, δ	δ from FOC (10) at k_0; K from modified FOC	Unchanged vs. no-FCL (k fixed)	Ambiguous: $(R-c)\lvert\partial_K \mathbb{E}[\text{fossil}]\rvert \gtrless C$
B: K fixed at K_0, free k, δ	$F(k;K_0) = (\beta+p-c-\gamma-\kappa)/(\beta+p-R)$; δ from (10) at K_0	Lower iff $\kappa < \bar{\kappa}(K_0)$	Falls iff $\kappa < \bar{\kappa}(K_0)$
C: both k_0, K_0 fixed, free δ only	δ from (10) at (k_0, K_0); no capacity response	Unchanged (capacity fixed)	**Always rises** (utilization effect only)
Baseline (§4): all free	Table 8	Lower iff $\kappa < \bar{\kappa}$	Falls iff $\kappa < \bar{\kappa}$

Notes: This table reports optimal FCL capacity δ, fossil capacity k, blackout probability, and carbon outcomes for each of the three pre-commitment scenarios analyzed in Section 5, together with the simultaneous baseline (Section 4). Case A: fossil capacity k_0 is pre-installed; renewable K and FCL capacity δ are freely optimized. Case B: renewable capacity K_0 is pre-installed; k and δ are freely optimized. Case C: both k_0 and K_0 are fixed; only δ is chosen. The threshold $\bar{\kappa}(K_0) \equiv (R-c) \cdot F(k^*; K_0)$ uses the no-FCL optimal fossil capacity evaluated at fixed K_0; it satisfies $\bar{\kappa}(K_0) < \bar{\kappa}$ when $K_0 < K^{**}$ (less renewable capacity makes adoption harder).

Table 12: Rate-of-Return (Averch–Johnson) Regulation: Key Results

Quantity	**Without FCL**	**With FCL**
FOC for k (AJ)	$(\beta + p - c)[1 - F(k_{AJ})] = \gamma(1 - \bar{r})$	Same structure with κ in numerator
Over-investment	$k_{AJ} > k^*$ (since $\bar{r} < 1$ lowers effective γ)	$k_{AJ}^{**} > k^{**}$: further amplified
FCL rate-base effect	n/a	If FCL hardware in rate base: κ effectively reduced to $\kappa(1 - \bar{r})$, lowering $\bar{\kappa}^{ROR}$ threshold
Triple win threshold	n/a	$\bar{\kappa}^{ROR} > \bar{\kappa}$ (easier under ROR if hardware is rate-base eligible)
Carbon	AJ over-investment raises k_{AJ}, increases fossil utilization	FCL can partially offset if K also over-invested

Notes: This table summarizes optimal allocations and comparisons under rate-of-return (ROR) regulation following the Averch–Johnson (AJ) framework (Section 6). The allowed rate of return $\bar{r}$ satisfies $\bar{r} \geq \rho$ where ρ is the utility's cost of capital; when $\bar{r} > 0$, the effective net capital cost is $\gamma(1 - \bar{r}) < \gamma$, inducing over-investment in fossil capacity ($k_{AJ} > k^*$). If FCL hardware qualifies as rate base, the hardware cost parameter κ is effectively reduced to $\kappa(1 - \bar{r})$, widening the triple-win threshold to $\bar{\kappa}^{ROR} > \bar{\kappa}$. Compare with the unregulated baseline (Table 8).

Table 13: Deregulated Competitive Market: Equilibrium Conditions with and without FCL

Condition	**Without FCL**	**With FCL**
Fossil zero-profit	$(P_{\text{cap}} - c)\Pr(\tilde{x} > k) = \gamma$	Same (FCL curtails at scarcity, reduces blackout severity)
FCL zero-profit	n/a	$R\Pr(\tilde{x} \leq 0) + (R-c)\Pr(0 < \tilde{x} < k) = \kappa$
Blackout probability	$\dfrac{\gamma}{P_{\text{cap}} - c}$	$\dfrac{c + \gamma + \kappa - R}{P_{\text{cap}} - R}$
Triple win threshold	n/a	$\bar{\kappa}^{\text{comp}} = \dfrac{(R-c)(P_{\text{cap}} - c - \gamma)}{P_{\text{cap}} - c}$
"Missing money"	Scarcity rents $P_{\text{cap}} - c$ cover fixed cost γ	FCL increases fossil utilization, raising merchant revenue

Notes: This table presents equilibrium conditions in a deregulated competitive electricity market (Section 7). Fossil generators earn zero profit when scarcity rents $(P_{\text{cap}} - c)\Pr(\tilde{x} > k)$ exactly cover fixed cost γ; FCL operators earn zero profit when average computing revenue equals hardware cost κ. P_{cap} is the scarcity price cap (or value of lost load). Results are structurally isomorphic to the regulated baseline (Table 8) with P_{cap} replacing $\beta + p$; in particular, the triple-win threshold $\bar{\kappa}^{\text{comp}}$ exceeds $\bar{\kappa}$ when $P_{\text{cap}} > \beta + p$.

Table 14: Grid Scope and FCL Program Design: Local versus Global Hash Market

Scenario	**Assumption on R**	**Key implication**
Single local grid	R exogenous (price-taker)	Sections 4–7 apply unchanged
Multi-grid coordination	R endogenous: falls with $\sum_i \delta_i$	Nash equilibrium $\delta_i^* < \delta^{**}$ (under-provision vs. social optimum)
Social planner	R endogenous, internalized	Higher aggregate δ than Nash; carbon reduction stronger
Single global grid	One grid controls all mining, R endogenous; $R^G = \bar{R}/\mathbb{E}[d]$	$\delta^G = 0$ (FCL not offered): total mining income $\bar{R}$ is constant, marginal mining revenue vanishes; capacity reverts to no-FCL benchmark (Prop. 3)

Notes: This table contrasts the assumption on mining revenue R and its implication for optimal FCL program design across four scenarios analyzed in Section B. In the single-grid (local) case, the utility serves a small fraction of global miners so R is pinned by global Bitcoin markets (exogenous). In the multi-grid Nash case, N independent grids each choose δ_i, and mining revenue R falls endogenously in aggregate FCL capacity $\sum_i \delta_i$; the Nash equilibrium is inefficiently low relative to the social planner (tragedy of the commons). In the global monopolist case, one grid controls all mining: total mining income $\bar{R} = R \cdot \mathbb{E}[d]$ is constant, so the marginal mining-revenue incentive vanishes and $\delta^G = 0$ (Proposition 3).

Table 15: Master Summary: Blackout Probabilities and Triple-Win Threshold across Scenarios

Scenario	**Blackout (no FCL)**	**Blackout (with FCL)**	**Triple-win threshold $\bar{\kappa}$**
Baseline regulated, exog. p	$\gamma/(\beta+p-c)$	$(c+\gamma+\kappa-R)/(\beta+p-R)$	$(R-c)(\beta+p-c-\gamma)/(\beta+p-c)$
Case b (small FCL)	$\gamma/(\beta+p-c)$	$(c+\gamma+\kappa-R-cF(-\delta^*;K^{**}))/(\beta+p-R)$	$\bar{\kappa}^b > \bar{\kappa}$
Capacity case A (k fixed)	$1-F(k_0;K^*)$	$1-F(k_0;K^{**})$ (lower K effect)	$(R-c)F(k^*;K_0)$
Capacity case B (K fixed)	$\gamma/(\beta+p-c)$ at K_0	Lower iff $\kappa < \bar{\kappa}(K_0)$	$\bar{\kappa}(K_0) = (R-c)F(k^*;K_0)$
Capacity case C (both fixed)	$1-F(k_0;K_0)$	Unchanged	n/a (no capacity channel)
ROR regulation	$\gamma(1-\bar{r})/(\beta+p-c-\gamma\bar{r})$	Modified by AJ over-invest.	$\bar{\kappa}^{ROR} \geq \bar{\kappa}$ if hardware rate-based
Competitive market	$\gamma/(P_{\text{cap}}-c)$	$(c+\gamma+\kappa-R)/(P_{\text{cap}}-R)$	$(R-c)(P_{\text{cap}}-c-\gamma)/(P_{\text{cap}}-c)$
Corner regime	$\gamma/(\beta+p-c)$	0	Always satisfied: $\kappa \leq R-c-\gamma$
Global monopolist (N=1, m=0)	$\gamma/(\beta+p-c)$	Same as no-FCL ($\delta^G=0$)	n/a (FCL not adopted; Prop. 3)
Multi-grid Nash (N grids)	$\gamma/(\beta+p-c)$	$(c+\gamma+\kappa-R^{NE})/(\beta+p-R^{NE})$ at $R^{NE}<R$	$\bar{\kappa}^{NE} < \bar{\kappa}$ (harder to achieve)

Notes: This table provides a unified comparison of equilibrium blackout probabilities (with and without FCL program) and the triple-win threshold $\bar{\kappa}$ across all main scenarios. All expressions assume Case c ($\delta^{**} > K^{**}$) and the interior regime ($R < c+\gamma+\kappa$) at the simultaneous optimum; for Case b see Table 10, for the corner regime see Table 9. All blackout expressions are equilibrium values at the respective optimum. $\bar{\kappa}(K_0) \equiv (R-c)\cdot F(k^*;K_0)$ uses the no-FCL optimal k^* evaluated at fixed K_0; it equals $\bar{\kappa}$ when $K_0 = K^*$. Under the competitive market, P_{cap} replaces $\beta+p$ throughout. For the multi-grid Nash case, $R^{NE}<R$ is the endogenous revenue at the Nash equilibrium; see Appendix C. In the global monopolist case, total mining income $\bar{R}$ is constant in δ so the FCL program is not offered (Proposition 3).