

Artificial Intelligence and Environmental Decision-Making: Does AI Adoption Improve Corporate Resource Efficiency?

Tanvir Rahman^{1*}, Ashikuzzaman², Kazi Arham Ahmed³

¹ Stamford University, Dhaka, Bangladesh

² Daffodil International University, Dhaka, Bangladesh

³ Universiti Kebangsaan Malaysia, Selangor, Malaysia *Corresponding author: tanvir.bappi32@gmail.com

Received 13 August 2024; Revised 18 September 2024; Accepted 11 October 2024; Published: 28 November 2024

KEYWORDS

Artificial Intelligence,
Environmental Decision-Making,
Resource Efficiency,
Corporate Sustainability,
Rebound Effect,
Green AI,
Algorithm Adoption

Abstract

Companies increasingly use machine learning and other artificial intelligence (AI) tools to forecast demand, tune processes, choose suppliers, route goods, and monitor environmental performance, and these tools are often described as a route to lower use of energy, water, and materials. AI can improve the decisions that determine resource use, but it also consumes resources itself, it can increase the output that a firm produces, and its recommendations improve resource efficiency only if people and systems act on them. This study builds an evidence based framework for asking whether AI adoption improves corporate resource efficiency. It synthesizes peer reviewed studies published before 1 January 2024 on the environmental footprint of AI, the sustainability potential of AI, algorithm use in decision-making, firm level productivity effects of data driven decisions and AI, and rebound effects. The evidence does not support a universal efficiency gain: direct firm level evidence linking AI adoption to resource use is very thin, the footprint of large models is substantial and rarely netted against the savings they enable, and the productivity literature points to higher output as well as higher efficiency. The paper therefore proposes a resource ledger in which the net resource saving equals the efficiency gain from better decisions minus an output scale effect and minus the operating footprint of the AI system itself. It derives a break even condition for output growth, a footprint budget, and an offset ratio, models the efficiency gain as the product of decision error reduction and the rate at which recommendations are acted on, and develops research propositions, an empirical design, and an AI Environmental Decision Maturity Index (AEDMI). The central implication is that AI improves corporate resource efficiency only when the savings from better decisions exceed the resources used to produce those decisions and the additional production they enable.

1. Introduction

Corporate resource efficiency is decided in thousands of ordinary choices: how much to produce and stock, which supplier to use, how to set a process, which route a delivery takes, and which waste stream to tackle first. Most of these choices are made under uncertainty and with limited attention, so a firm’s resource use reflects the quality of its decisions as well as the quality of its equipment. This is the bounded rationality problem described by Simon (1955), and it is the reason why artificial intelligence (AI) is attractive in environmental management. Machine learning can forecast demand more accurately and so reduce overproduction and spoilage, anticipate equipment faults, optimize routes and schedules, screen suppliers on environmental attributes, and turn raw monitoring data into targets. Reviews of the climate and sustainability potential of AI list many such applications (Rolnick et al., 2022; Vinuesa et al., 2020).

Whether AI adoption actually improves corporate resource efficiency is nevertheless an open question, for three reasons. First, AI is a resource user. Training and running large models consumes electricity, cooling water, and the materials embedded in hardware (Strubell et al., 2019; Patterson et al., 2021; Wu et al., 2022; Luccioni et al., 2023; Li et al., 2023), and data center energy demand has been contained so far only because efficiency gains have offset very rapid growth in computing (Masanet et al., 2020). A firm that adopts AI therefore adds a new item to its resource ledger. Second, tools that improve decisions and productivity can raise output: firms that adopt data driven decision-making and AI report higher productivity and faster growth (Brynjolfsson and McElheran, 2016; Brynjolfsson et al., 2011; Czarnitzki et al., 2023), and the rebound literature has long warned that efficiency gains can be partly or fully offset when they expand activity (Khazzoom, 1980; Brookes, 1990; Greening et al., 2000; Sorrell and Dimitropoulos, 2008; Gillingham et al., 2016). Third, a recommendation changes resource use only when someone acts on it, and decision makers do not always follow algorithmic advice (Dietvorst et al., 2015; Logg et al., 2019). Assessments of AI and climate stress that direct, enabling, and systemic effects can point in different directions (Kaack et al., 2022; Freitag et al., 2021).

This paper therefore asks: *Does AI adoption improve corporate resource efficiency?* The question is framed at the firm and plant level and is broken into four research questions.

- RQ1** Does AI adoption raise resource productivity by improving the decisions that determine energy, water, and material use?
- RQ2** Does it raise output, and does the AI system itself add to resource use (scale and footprint)?
- RQ3** Under what conditions do the efficiency gains outweigh the scale and footprint effects?
- RQ4** How can the gains, the offsets, and the organizational processes behind them be measured at the firm level?

The paper makes five contributions. First, it treats AI as environmental decision infrastructure and links it to resource use through a five stage decision cycle. Second, it proposes a resource ledger that nets the efficiency gain against an output scale effect and the footprint of the AI system itself, with an exact accounting identity, a break even growth condition, a footprint budget, and an offset ratio. Third, it models the efficiency gain as the product of the potential gain, the reduction in decision error, and the rate at which recommendations are acted on. Fourth, it proposes measurable indicators and a composite index for firm level research. Fifth, it states testable propositions and an evidence gap map for future work.

The evidence boundary is 31 December 2023: publications from 2024 onward are not used as evidence. The paper proceeds as follows. Section 2 reviews the literature. Section 3 develops the theoretical framework. Section 4 describes the methodology and measurement. Section 5 synthesizes the evidence and illustrates the framework. Section 6 discusses implications, Section 7 sets out limitations and a research agenda, and Section 8 concludes. Figure 1 summarizes the decision cycle on which the framework is built.




Figure 1. The AI decision cycle for corporate environmental decisions.

2. Literature Review

2.1 The Sustainability Potential of AI

Assessments of AI and sustainability set out a wide set of applications. Rolnick et al. (2022) review how machine learning can contribute to climate change mitigation and adaptation across sectors such as electricity systems, transportation, buildings, industry, and agriculture, and they also stress that many of these opportunities depend on data, deployment, and complementary policy. Vinuesa et al. (2020) assess how AI relates to the Sustainable Development Goals and conclude that it can both enable and inhibit progress, depending on the target and on how it is governed. Kaack et al. (2022) distinguish the direct emissions of computing, the enabling effects of applications, and systemic effects of AI on consumption and production, and they argue that alignment with climate goals requires attention to all three. These sources establish that a plausible set of pathways exists. They do not measure how much resource use falls in a given firm after adoption, because they mostly describe potential.

2.2 The Resource Footprint of AI

AI systems use resources. Strubell et al. (2019) estimated the energy and carbon costs of training natural language processing models and showed that the cost of developing and tuning large models can be substantial. Patterson et al. (2021) show that the carbon emissions of training large neural networks depend on model architecture, hardware, data center efficiency, and the electricity mix, so that practices matter as well as scale. Wu et al. (2022) describe the environmental implications of AI across data, experimentation, training, inference, and hardware manufacturing, and Luccioni et al. (2023) estimate the footprint of a large language model across training, hardware manufacturing, and deployment. Water is a further resource: Li et al. (2023) describe the water consumed in cooling and electricity generation for AI model training and inference. At the system level, Masanet et al. (2020) show that global data center energy use rose only modestly over a period in which computing output grew several-fold, which illustrates the efficiency gains available in digital infrastructure and the need to track the aggregate. For a firm, the implication is that the operating footprint of an AI system, whether on premises or purchased as a cloud service, belongs on the same ledger as the resource savings it is meant to produce.

2.3 Decision-Making, Prediction, and Algorithm Use

The economic logic of AI is the logic of prediction. Agrawal et al. (2019) argue that AI lowers the cost of prediction and that its value depends on how prediction combines with human judgment and action in decisions. This matters for resource efficiency because a better forecast reduces waste only if the decision it informs actually changes. Behavioral research shows that people respond to algorithmic advice in uneven ways. Dietvorst et al. (2015) find that people can avoid algorithms after seeing them err, and Logg et al. (2019) find that people often prefer algorithmic to human judgment in some settings. A firm's environmental benefit from AI therefore depends on trust, accountability, and process design, and not on model accuracy alone.

2.4 Firm Productivity, Growth, and Complementarities

Studies of firm level data link data driven decision-making and AI with output. Brynjolfsson and McElheran (2016) document the rapid adoption of data driven decision-making in United States manufacturing, and Brynjolfsson et al. (2011) report that firms with greater use of data driven decision-making have higher output and productivity. Czarnitzki et al. (2023) study firms in Germany and report that firms using AI have higher productivity. Brynjolfsson et al. (2023) report field evidence that a generative AI assistant raised the productivity of customer support agents. Brynjolfsson et al. (2021) argue that general purpose technologies such as AI require intangible complements, so that measured returns lag investment. For resource efficiency these results matter in two ways. If AI adoption raises output, it creates the scale effect that works against resource savings, and if returns arrive with a lag, a short evaluation window can misstate both gains and offsets.

2.5 Management, Adoption, and Resource Efficiency

Evidence on industrial energy efficiency suggests that organizations often fail to capture apparently attractive savings. Gerarden et al. (2017) review explanations for the energy efficiency gap, and Anderson and Newell (2004) find that adoption of audit recommendations in plants depends on economic characteristics, not only on the information provided. Management capability appears repeatedly as a condition of performance: Boyd and Curtis (2014) link management practices with energy efficiency in United States manufacturing, and Bloom and Van Reenen (2007) show that management practices vary widely across firms and countries and relate to performance. These findings suggest that an AI system that produces better recommendations will improve resource efficiency only in organizations able to act on them.

2.6 Direct Evidence on AI and Environmental Outcomes

Direct empirical evidence on AI adoption and environmental outcomes in firms is very limited and mostly aggregate. Liu et al. (2022) study the industrial sector in China and report that artificial intelligence is associated with lower carbon intensity. The finding supports the intensity channel but concerns carbon intensity and not absolute resource use, and sector data cannot show whether output growth or the footprint of the AI system offset the gains inside individual firms. This is the main empirical gap that the present framework addresses.

2.7 Rebound, Decomposition, and Measurement

The rebound literature shows that efficiency improvements can be partly offset by additional use (Khazzoom, 1980; Brookes, 1990; Greening et al., 2000). Sorrell and Dimitropoulos (2008) distinguish direct and indirect mechanisms, and Gillingham et al. (2016) review the evidence on the size of rebound effects. Index decomposition methods allow changes to be attributed exactly among factors, and Ang (2004) compares methods and discusses log mean Divisia decomposition. The idea that win-win environmental improvements are possible for firms has been argued by Porter and van der Linde (1995). Measurement quality is a further constraint: Berg et al. (2022) show that environmental, social, and governance ratings diverge substantially across providers, which indicates that environmental data cannot be treated as settled. AI tools used for environmental reporting and decision support inherit these data problems.




Figure 2. Expected links between AI decision domains and resource streams.

2.8 Research Gap

Five gaps are identified. First, assessments of AI and sustainability describe potential more often than realized firm level outcomes. Second, research on the footprint of AI measures the resource use of models and rarely nets it against the savings those models enable. Third, studies of AI and productivity measure output and growth but not the resource use that accompanies it. Fourth, decision research on algorithm use is seldom connected to environmental outcomes. Fifth, a unified measurement framework connecting AI capability, decision integration, resource productivity, output, and AI footprint is still missing. This study addresses these gaps through a resource ledger framework. Table 1 summarizes the evidence base.

3. Theoretical Framework

3.1 AI as Environmental Decision Infrastructure

The paper treats AI as environmental decision infrastructure rather than as an inherently green technology. A forecasting model can reduce overproduction, but only if planners change production plans. A routing model can shorten delivery distances, but only if dispatchers accept its routes. A supplier screening model can flag high impact inputs, but only if procurement rules use the flag. Resource efficiency is therefore a joint outcome of the model, the organization using it, and the resources consumed by the model. The decision cycle in Figure 1 has five stages: sense, predict, decide, act, and learn.

The framework identifies five decision domains in which AI can influence resource use: demand and inventory planning, process and energy optimization, procurement and supplier choice, logistics and routing, and monitoring and reporting (Table 2). Each domain acts mainly on particular resource streams. Figure 2 shows the expected strength of these links, which is the authors' conceptual judgment and not an estimate.

3.2 Accounting Framework: The Resource Ledger

Let a firm use n resource streams (for example energy, water, materials, and waste disposal) and let w_i convert stream i to a common unit, such as cost or environmental weight. Resource use in stream i is the quantity needed to produce output plus the resource use of the AI system,

$$U = \sum_{i=1}^n w_i \left(\frac{Y}{P_i} + A_i \right), \quad (1)$$

where Y is output, P_i is resource productivity in stream i (output per unit of resource), and A_i is the operating footprint of AI in that stream, including electricity, cooling water, and amortized hardware, whether on premises or purchased as a service. Superscripts B and A denote the situation before and after AI adoption, with $A_i^B = 0$ for a firm without




Figure 3. Resource ledger for a hypothetical firm as a flow diagram.

the system. Define three effects:

$$EG = \sum_i w_i Y^A \left(\frac{1}{P_i^B} - \frac{1}{P_i^A} \right), \quad (2)$$

$$SC = \sum_i w_i \frac{Y^A - Y^B}{P_i^B}, \quad (3)$$

$$AF = \sum_i w_i (A_i^A - A_i^B). \quad (4)$$

The efficiency gain EG is the saving from higher resource productivity at the new level of output, the output scale effect SC is the extra resource use from producing more at baseline productivity, and the footprint effect AF is the change in resource use of the AI system. The net resource saving is

$$NRS = EG - SC - AF = U^B - U^A. \quad (5)$$

The last equality is an exact identity: adding and subtracting $\sum_i w_i Y^A / P_i^B$ in $U^B - U^A$ gives the three terms. The split between EG and SC depends on the order in which output and productivity are changed, although NRS does not, and log mean Divisia methods (Ang, 2004) give an order independent alternative. This is an accounting identity and not an empirical estimate. A positive NRS means that total resource use falls. Figure 3 shows the ledger for the hypothetical firm used in Section 5.

3.3 Break Even Condition, Footprint Budget, and Offset Ratio

Suppose productivity rises by a common proportion $\pi = P^A / P^B - 1$, output grows by $g = Y^A / Y^B - 1$, and the AI footprint equals a share a of the baseline resource burden. Then $U^A / U^B = (1 + g) / (1 + \pi) + a$, and the net saving is non-negative if and only if

$$g \leq g^* = \pi - a(1 + \pi). \quad (6)$$

The break even growth rate g^* falls one for one with the productivity gain that is lost to the AI footprint. Equivalently, for given π and g the footprint must stay within a footprint budget

$$a \leq a^* = \frac{\pi - g}{1 + \pi}. \quad (7)$$

The share of the gross saving consumed by the two offsets is the offset ratio

$$O = \frac{SC + AF}{EG}, \quad (8)$$

and the net saving is positive when $O < 1$. With $\pi = 7.35\%$, $g = 3\%$, and $a = 2\%$, which are the base values used below, $g^* = 5.2\%$, $a^* = 4.05\%$, and the net saving is 2.05% of the baseline burden. Setting $a = 0$ recovers the familiar rebound condition that output growth must stay below the productivity gain.

Table 1. Selected Pre 2024 Evidence Relevant to AI Adoption and Corporate Resource Efficiency

Study	Context	Main evidence as reported	Interpretation
Strubell et al. (2019); Patterson et al. (2021)	Training of large models	Energy and carbon cost of training depends on model size, hardware, data center, and energy mix	AI footprint
Wu et al. (2022); Luccioni et al. (2023)	AI life cycle; one large language model	Footprint spans data, training, hardware manufacturing, and deployment	AI footprint
Li et al. (2023)	AI model training and inference	Cooling and electricity generation consume water	AI footprint (water)
Masanet et al. (2020)	Global data centers	Energy use grew modestly while computing output grew several-fold	Efficiency of digital infrastructure
Rolnick et al. (2022); Vinuesa et al. (2020)	Reviews	AI can enable climate and sustainability goals; some effects may inhibit	Potential, not realized outcomes
Kaack et al. (2022); Freitag et al. (2021)	Assessments	Direct, enabling, and systemic effects of computing can differ in sign; estimates are uncertain	Rebound and systemic effects
Agrawal et al. (2019)	Theory	AI lowers the cost of prediction; value depends on judgment and action	Decision pathway
Dietvorst et al. (2015); Logg et al. (2019)	Experiments	People can avoid or prefer algorithmic advice depending on context	Adoption conditions
Brynjolfsson and McElheran (2016); Brynjolfsson et al. (2011)	United States firms	Data driven decision-making linked with higher output and productivity	Scale pathway
Czarnitzki et al. (2023)	German firms	AI use associated with higher firm productivity	Scale pathway
Brynjolfsson et al. (2023)	Customer support firm	Generative AI assistant raised worker productivity	Scale pathway
Brynjolfsson et al. (2021)	Aggregate analysis	Returns to general purpose technology lag because of intangible complements	Lag and complementarity
Boyd and Curtis (2014); Bloom and Van Reenen (2007)	Manufacturing firms	Management practices linked with energy efficiency and performance	Capability condition
Anderson and Newell (2004); Gerarden et al. (2017)	Plants; review	Adoption of efficiency projects depends on economic characteristics	Adoption economics
Liu et al. (2022)	China industrial sector	AI associated with lower carbon intensity	Intensity channel (aggregate)
Berg et al. (2022)	ESG rating providers	Environmental ratings diverge across providers	Data quality

Table 2. Five Decision Domains and Their Resource Pathways

Domain	AI task	Efficiency pathway	Scale or footprint pathway	Example indicator
Demand and inventory planning	Demand forecasts; replenishment and production plans	Less overproduction, spoilage, and surplus material	Faster fulfillment can raise volume	Forecast error; surplus stock; waste per unit
Process and energy optimization	Setpoint, schedule, and load optimization	Lower energy, water, and scrap per unit	Higher throughput and yield; compute energy	Energy per unit; yield; model energy use
Procurement and supplier choice	Supplier screening; material substitution	Lower embodied material use; lower impact inputs	Lower unit cost can raise purchasing	Material per unit; share of screened inputs
Logistics and routing	Route and load optimization	Fewer kilometers and empty trips	Cheaper delivery can raise shipments	Fuel per tonne-kilometer; load factor
Monitoring and reporting	Anomaly detection; automated environmental reporting	Earlier detection of leaks and waste; better targets	Data processing footprint; false confidence in poor data	Leak detection time; data coverage

3.4 Decision Mechanism: From Better Predictions to Realized Gains

The productivity gain is not a property of the model. It depends on the value of better decisions in the firm’s setting, on how much decision error the AI system removes, and on how often recommendations are acted on. A simple representation is

$$\pi = \pi_{\max} \theta \alpha, \quad (9)$$

where $\pi_{\max}$ is the productivity gain that error free decisions, fully implemented, would deliver, $\theta \in [0, 1]$ is the proportional reduction in decision error (such as forecast error) that the system achieves, and $\alpha \in [0, 1]$ is the share of recommendations that are acted on. The relation is multiplicative, so a high value of either factor cannot compensate for a very low value of the other. The decision chain behind α and θ has four stages: a relevant and accurate signal is produced (p_1), the recommendation is trusted (p_2), it is acted on (p_3), and the action delivers a net benefit after offsets (p_4). The share of deployments completing the chain is

$$\Pi = \prod_{j=1}^4 p_j. \quad (10)$$

Because the probabilities multiply, modest failure at every stage compounds into a small overall effect. The pathway is expected to be stronger where firms integrate recommendations into routine decisions, assign accountability for outcomes, and have reliable resource data.

3.5 Moderators and Boundary Conditions

The framework expects decision integration and governance to strengthen the pathway, and it expects the quality of resource data to condition the gain. It treats the carbon intensity of electricity and the location of computing, the choice between large and small models, product market conditions, and regulation as boundary conditions. The same system can have different consequences where computing is powered by low carbon electricity, where models are sized to the task, or where the firm is capacity constrained and so can use savings to grow.

4. Methodology

4.1 Research Design

The study is a structured secondary evidence synthesis. The evidence boundary is 31 December 2023, and sources published from 1 January 2024 onward are excluded. Peer reviewed journal articles, established conference proceedings and preprints from the AI research community, academic reviews, and authoritative institutional materials are prioritized. The unit of analysis is the firm or plant decision pathway. The study does not claim a new firm dataset or a new causal estimate; it develops a framework for future empirical testing.

Seven evidence groups are used: the footprint of AI; the sustainability potential of AI; rebound and systemic effects; decision science and algorithm use; firm productivity and growth; management and adoption conditions; and direct environmental outcomes and environmental data quality. Each study was coded by theme and by level of analysis (Section 5.1).

Table 3. Evidence Selection Framework

Criterion	Included	Excluded
Date	Published before 2024	Published in 2024 or later
Quality	Peer reviewed research, preprints from established groups, authoritative reports	Unsupported commentary
Topic	AI footprint, AI potential, decisions, productivity, rebound	Unrelated digital technologies
Outcome	Resource use, productivity, output, adoption, footprint	Outcomes without resource relevance
Unit	Firm, plant, industry, economy, model	Purely technical benchmark studies

4.2 Measurement Variables

The primary independent construct is AI adoption (AI). It can be measured with the share of decisions supported by models, the number of production use cases, spending on AI services, and adoption dates. The




Figure 4. Measurement pyramid for firm level research.




Figure 5. Rank changes of three hypothetical firms under four weighting schemes for the AEDMI.

main dependent constructs are resource productivity in each stream, output, and the AI footprint, which together determine the net saving. Environmental outcomes should include both intensity and absolute measures. Table 4 lists the constructs, and Figure 4 shows how raw records are turned into the net saving.

4.3 AI Environmental Decision Maturity Index

The study proposes an AI Environmental Decision Maturity Index (AEDMI) built from five components: data foundation (D), model capability (M), decision integration (I), governance and oversight (G), and footprint management (F). Each raw indicator z_{jh} for firm h is first rescaled to $[0, 100]$,

$$\tilde{z}_{jh} = 100 \frac{z_{jh} - z_j^{\min}}{z_j^{\max} - z_j^{\min}}, \quad (11)$$

and the index is the weighted sum

$$\text{AEDMI}_h = w_1 D_h + w_2 M_h + w_3 I_h + w_4 G_h + w_5 F_h, \quad \sum_{j=1}^5 w_j = 1. \quad (12)$$

Equal weights may be used in exploratory work, while factor analysis or expert weights may be used in more advanced studies. Weights matter because they can change rankings. Figure 5 shows three hypothetical firms scored under four weighting schemes. Firm Y (integration-led) ranks first under equal weights (64.0 against 63.0 for firm Z and 58.0 for firm X) and under weights that emphasize decision integration (70.3, 60.8, and 54.3), firm X (model-rich) ranks first when model capability is emphasized (71.0), and firm Z (footprint-careful) ranks first when footprint management is emphasized (68.8 against 62.8 for firm Y and 49.3 for firm X). Weight sensitivity should therefore be reported. The AEDMI describes capability; it does not replace the absolute outcome measure NRS, which is estimated separately.

4.4 Empirical Models for Future Testing

Because $U = (Y/P)(1 + \phi)$ with $\phi = AP/Y$ the AI footprint as a share of operating resource use, Equation (1) implies $\ln U = \ln Y - \ln P +$

$\ln(1 + \phi)$ for a single stream measured consistently. The net effect can therefore be recovered from three parallel firm panel regressions. For each outcome $y \in \{\ln P, \ln Y, \ln(1 + \phi)\}$,

$$y_{it} = \alpha_0^y + \alpha_1^y AI_{it} + \alpha_2^y DI_{it} + \alpha_3^y X_{it} + \mu_i + \lambda_t + \varepsilon_{it}, \quad (13)$$

where μ_i represents firm effects and λ_t represents time effects. The total proportional effect of AI adoption on resource use is the combination of the three coefficients,

$$c = \alpha_1^Y - \alpha_1^P + \alpha_1^\Phi, \quad (14)$$

where α_1^P is expected to be positive (efficiency), α_1^Y is expected to be positive (scale), and α_1^Φ is expected to be positive (footprint). Resource use falls when $c < 0$, that is, when $\alpha_1^P > \alpha_1^Y + \alpha_1^\Phi$, which mirrors NRS in Equation (5). This structure separates the three pathways while respecting the accounting identity. Firm fixed effects remove time invariant differences, but adoption of AI is not random. Credible identification therefore calls for event study designs around adoption dates that use estimators robust to staggered adoption (Callaway and Sant'Anna, 2021; Sun and Abraham, 2021), since two way fixed effects estimates can be biased under heterogeneous timing (Goodman-Bacon, 2021), together with instruments based on plausibly exogenous variation in the availability of vendor services or connectivity, and placebo tests on outcomes that should not respond. Because returns to general purpose technologies can arrive with a lag (Brynjolfsson et al., 2021), dynamic specifications with several post adoption periods are required.

4.5 Moderation Tests

Decision integration is expected to strengthen the relationship between AI adoption and resource productivity,

$$\begin{aligned} \ln P_{it} = & \beta_0 + \beta_1 AI_{it} + \beta_2 DI_{it} \\ & + \beta_3 (AI_{it} \times DI_{it}) \\ & + \beta_4 X_{it} + \mu_i + \lambda_t + \varepsilon_{it}, \end{aligned} \quad (15)$$

and a positive β_3 would indicate a larger productivity gain among firms that act on recommendations. The same specification with data quality (DQ) in place of DI tests the role of resource data, and with the AI footprint share as the outcome tests whether right-sized and well-governed models have lower footprints.

4.6 Data for a Bangladesh Application

A study in Bangladesh could combine a firm survey on AI adoption, utility records for electricity, gas, and water, procurement and waste records, and production data for resource intensive manufacturing such as textiles and garments or food processing. Candidate sources include the establishment surveys of the Bangladesh Bureau of Statistics, factory level resource audit records, billing records, and a purpose built survey that measures AI use cases, decision integration, and data quality. The footprint of AI services should be taken from cloud usage logs and vendor disclosures where available, with emission factors for the national electricity mix applied consistently. Access, sampling, confidentiality, and consent conditions must be confirmed before fieldwork.

4.7 Validation Plan

The index and the ledger should be validated before they are used for inference. For the AEDMI, content validity can be reviewed by experts, dimensionality tested with confirmatory factor analysis, measurement invariance checked across sectors and firm sizes, and reliability examined through repeated measurement. For the ledger, results should be reported under alternative orderings and with a log mean Divisia decomposition, under alternative stream weights w_i , and with and without the footprint of purchased AI services.

5. Analysis and Results

5.1 Evidence Synthesis

The pre 2024 literature provides no single direction of effect. Figure 6 counts the reviewed studies by theme, and four patterns emerge.

First, the best documented regularity concerns the footprint of AI: six studies document the energy, carbon, or water use of models and

Table 4. Constructs, Indicators, and Candidate Data for Firm Level Research

Construct	Symbol	Example indicators	Unit	Candidate data source
AI adoption	AI	Share of decisions model supported; production use cases; adoption date; AI spending	Share or count	Firm survey; vendor and cloud records
Resource productivity	P_i	Output per unit of energy, water, and material; waste per unit	Output per unit	Meters; ERP; procurement and waste records
Output	Y	Output volume or value added; capacity utilization	Units or currency	ERP and production records
AI footprint	A_i	Electricity, cooling water, and hardware attributable to AI workloads	kWh; m ³ ; kg	Cloud usage logs; utility bills; hardware inventory
Net resource burden	U	Weighted sum of streams including AI footprint	Weighted index	Resource ledger
Decision integration	DI	Share of recommendations acted on; override rate; accountability for outcomes	Share; score	Decision logs; firm survey
Data quality	DQ	Metering coverage; verification; update frequency	Score (0–100)	Firm survey; audit records
Controls	X	Firm size, age, sector, resource prices, capacity utilization	Various	Survey; administrative data

Table 5. AEDMI Components and Example Scoring

Component	Example indicators	Scoring
D: data foundation	Metering coverage; data verification; integration of ERP and sensor data	Min–max
M: model capability	Validated forecast error; model monitoring; retraining routine	Min–max
I: decision integration	Share of decisions model supported; share of recommendations acted on	Min–max
G: governance and oversight	Named owner of outcomes; audit of model use; override policy	Score (0–100)
F: footprint management	Measured AI footprint; model right-sizing; low carbon computing	Score (0–100)




Figure 6. Number of reviewed pre 2024 studies by theme. A study can appear under more than one theme.

data centers (Strubell et al., 2019; Patterson et al., 2021; Wu et al., 2022; Luccioni et al., 2023; Li et al., 2023; Masanet et al., 2020), and they show that the footprint depends on model size and on how and where computing is done. Second, evidence that data driven decision-making and AI raise output, productivity, or growth is consistent where it is measured (Brynjolfsson and McElheran, 2016; Brynjolfsson et al., 2011; Czarnitzki et al., 2023; Brynjolfsson et al., 2023), which establishes the conditions for a scale effect, and the benefits may arrive with a lag (Brynjolfsson et al., 2021). Third, assessments of the sustainability potential of AI list many applications (Rolnick et al., 2022; Vinuesa et al., 2020) but warn that direct, enabling, and systemic effects can differ in sign (Kaack et al., 2022; Freitag et al., 2021), and decision research shows that realized gains depend on whether advice is used (Agrawal et al., 2019; Dietvorst et al., 2015; Logg et al., 2019; Simon, 1955). Fourth, direct evidence on resource outcomes is very thin: one study reports lower carbon intensity in an industrial sector (Liu et al., 2022), and evidence on management and adoption indicates that capability conditions the capture of efficiency gains (Boyd and Curtis, 2014; Bloom and Van Reenen, 2007; Anderson and Newell, 2004; Gerarden et al., 2017). No reviewed study measures the efficiency gain, the output effect, and the AI footprint together in the same firms. The absence of net evidence may reflect what is measured as much as real differences in effects, and it supports measuring all three together.

5.2 The Resource Ledger: A Numerical Illustration

The ledger in Section 3.2 can be illustrated with a hypothetical firm. The numbers are assumed for exposition and are not estimates from data. The firm's baseline resource burden is 100 units in a common




Figure 7. Resource productivity by stream before and after charging the AI footprint to the stream.

weighted index, made up of energy (35), water (15), materials (40), and waste disposal (10). After AI adoption, output rises by $g = 3\%$, and resource productivity rises by 6% in energy, 5% in water, 9% in materials, and 12% in waste, because forecasting reduces overproduction and surplus material, process models trim energy and water use, and supplier screening shifts material choices. The AI system adds 1.2 units of energy, 0.5 of water, and 0.3 of materials (embodied in hardware), a total footprint of 2.0 units. Table 6 reports the ledger and Figure 3 shows it as a flow.

In this example the efficiency gain is 7.28 units, the scale effect 3.00, and the footprint 2.00, so $NRS = +2.28$: the burden falls from 100.00 to 97.72 and the offset ratio is 69%. The aggregate result hides the pattern by stream. Energy and water show negative net savings, because the AI system draws on these streams and the gains in them are modest, while materials and waste carry the benefit. A firm that evaluated AI only by its effect on materials and waste would overstate the benefit, and one that looked only at energy would see no gain. Figure 7 makes the same point in productivity terms: gross of the AI footprint, productivity rises by 7.6% across streams, but net of the footprint and the associated resource use it rises by 5.4%, and in energy it rises by 2.4% instead of 6.0%.

5.3 Sensitivity: Which Assumptions Matter Most

Under the base values $\pi_{\max} = 14\%$, $\theta = 0.75$, $\alpha = 0.70$, $g = 3\%$, and $a = 2\%$, the productivity gain is 7.35% and the net saving is 2.05% of the baseline burden. Figure 8 varies each parameter in turn across a plausible range. The result is most sensitive to the parameters that determine the efficiency gain and to output growth. Moving the action rate from 0.50 to 0.90 moves the net saving from 0.1% to 3.9%, as does moving the maximum achievable gain from 10% to 18%. Moving output growth from 5% to 1% moves the net saving from 0.2% to 3.9%, moving the AI footprint share from 4% to 0.5% moves it from 0.1% to 3.6%, and moving decision-error reduction from 0.60 to 0.90 moves it from 0.7% to 3.4%. Because the multiplicative structure in Equation (9) lets any one weak factor pull the gain down, the table of footprint budgets below helps to see how little slack the base case has.

Table 7 reports the break even growth rate and net savings for different AI footprint shares at the base productivity gain. The break even growth rate falls from 7.4% with no footprint to 5.2% at a footprint of

Table 6. Hypothetical Resource Ledger for One Firm ($g = 3\%$; weighted index units, baseline burden = 100; positive NRS means resource use falls)

Stream	Baseline use	π_i (%)	AI foot-print	Effic. gain	Scale effect	Foot-print effect	Net saving	Use after	Offset ratio (%)
Energy	35.0	6	1.20	2.04	1.05	1.20	-0.21	35.21	110
Water	15.0	5	0.50	0.74	0.45	0.50	-0.21	15.21	129
Materials	40.0	9	0.30	3.40	1.20	0.30	+1.90	38.10	44
Waste	10.0	12	0.00	1.10	0.30	0.00	+0.80	9.20	27
Total	100.0	—	2.00	7.28	3.00	2.00	+2.28	97.72	69




Figure 8. One at a time sensitivity of the net resource saving to the model parameters around the base case (solid line, 2.05%).

2% and 0.9% at a footprint of 6%, so a heavier model can exhaust the gains that the same decisions would otherwise have delivered.

Table 7. Footprint Budget at the Base Productivity Gain ($\pi = 7.35\%$)

AI footprint, a (%)	g^* (%)	NRS, $g=3\%$	NRS, $g=6\%$
0	7.35	+4.05	+1.26
1	6.28	+3.05	+0.26
2	5.20	+2.05	-0.74
4	3.06	+0.05	-2.74
6	0.91	-1.95	-4.74

NRS as a percentage of the baseline burden.

5.4 Uncertainty: A Simulation Across Deployment Designs

Firms differ in how they deploy AI, and parameter uncertainty is large. Figure 9 shows the distribution of the net saving from a simulation of 40,000 draws for each of three assumed deployment designs. In pilot analytics, productivity gain is drawn from a normal distribution with mean 3% and standard deviation 2%, output growth with mean 1% and standard deviation 1.5%, and the footprint share from a lognormal distribution with mean 0.4%. In embedded decision support, the corresponding means are 7.5%, 3%, and 1.5% (standard deviations 2.5% and 2%). In large model scale-up, the means are 8.5%, 5%, and 3.5% (standard deviations 3% and 2.5%). The probability that resource use falls is about 73% for pilot analytics, 81% for embedded decision support, and 48% for the large model scale-up, with medians of +1.5%, +2.7%, and -0.2%. The ordering shows that the larger productivity gain of a bigger deployment can be cancelled by its growth and footprint, so that scaling a system raises the stakes of measuring all three terms. The distributions are produced by assumed inputs and are not empirical.

5.5 Decision Integration as the Missing Link

AI adoption should not be expected to change resource use automatically. A signal must be produced, the recommendation must be trusted, it must be acted on, and the action must deliver a net benefit after offsets, as in Equation (10). The chain can fail at any stage: a forecast may be accurate but ignored by planners, a planner may override a recommendation after one visible error, or an implemented saving may be used to raise output. Figure 10 illustrates the compounding with assumed stage probabilities. For a standalone analytics dashboard, $p = (0.75, 0.40, 0.50, 0.60)$ gives 9.0% of deployments completing the chain. With decision support embedded in planning routines, $p = (0.85, 0.60, 0.70, 0.70)$ gives 25.0%, and with closed loop execution under governance, $p = (0.90, 0.75, 0.85, 0.80)$ gives 45.9%. The




Figure 9. Simulated distribution of the net resource saving for three assumed deployment designs (40,000 draws each).




Figure 10. Illustrative share of deployments completing all four stages of the decision chain (signal, trust, action, net benefit).

values are hypothetical, but the arithmetic shows why the integration and governance of a system can matter more than the accuracy of its models.

5.6 Decision Quality and Action Rate as Moderators

The efficiency gain depends jointly on the decision-error reduction and the action rate. Figure 11 plots the net saving implied by Equation (9) as both vary, with $\pi_{max} = 14\%$, $g = 3\%$, and $a = 2\%$. Net savings range from -4.1% when error reduction is low and recommendations are rarely acted on to +7.6% when both are at their maximum, and the base case sits at 2.05%. The surface is nearly a plane in the region shown, so improving either factor helps, but the zero plane cuts through the lower part of the surface: a firm with a low action rate gains little from a better model, and a firm with a good model gains little if it does not act. This is the pattern that the moderation test in Equation (15) is designed to detect. As decision-error reduction θ and action rate α vary, with $\pi_{max} = 14\%$, $g = 3\%$, and $a = 2\%$. The gray plane marks zero net saving and the dot marks the base case. Values are assumed, not estimated.

5.7 Cross Country and Sector Context

The environmental effect of AI adoption may vary across countries and sectors because electricity systems, cooling needs, resource prices, regulation, firm size, and digital infrastructure differ. The reviewed evidence on firm productivity comes mainly from the United States and Germany, and the sector level evidence on carbon intensity comes from China, so results should not be generalized automatically to Bangladesh, India, or other emerging economies. The footprint of the same model is lower where computing runs on low carbon electricity and higher where it does not, and the value of a saving depends on the resource price and the scarcity of the resource.




Figure 11. Net resource saving implied by the decision model.




Figure 12. Illustrative allocation of the saving potential of an AI portfolio across decision domains.

5.8 Research Propositions

The synthesis supports five propositions. Table 8 states each proposition with the model and expected sign used to test it, and Figure 12 shows how the expected saving potential is distributed across decision domains and resource streams for an illustrative firm.

6. Discussion

The answer to the research question depends on what is counted and how the system is used. If only the efficiency gain is counted, AI looks like an unambiguous improvement. If the output it enables and the resources it consumes are counted as well, the net result depends on their relative size, and the numerical illustration shows that a plausible case can produce a gain in materials and waste together with no gain in energy and water.

The decision perspective suggests where the benefit comes from. Demand and inventory planning and process optimization carry most of the expected saving potential in the illustration, in line with the idea that better prediction reduces the cost of acting under uncertainty (Agrawal et al., 2019). These are also the domains in which the link between a forecast and a physical action is shortest, so they are the natural places to start. Monitoring and reporting applications improve visibility but depend on data that may be unreliable (Berg et al., 2022), so a model trained on poor environmental data can increase confidence without increasing accuracy.

Resource efficiency should be measured with intensity and with

absolute use. A reduction in resource use per unit is not necessarily a reduction in total use, and the break even condition in Equation (6) makes this concrete: a productivity gain of a given size is beneficial only if output growth and the AI footprint together stay within it. Output growth is not a pure cost, since it can reflect competitiveness and employment, so environmental evaluation should be balanced against the economic value of the additional production (Porter and van der Linde, 1995).

6.1 Managerial Implications

Managers should put AI footprint on the resource ledger from the start, request footprint data from cloud and model providers, and compare the footprint with the savings the system is expected to generate. Model size should be chosen to fit the task, since a heavier model that adds little accuracy consumes part of the footprint budget in Equation (7) without raising θ . Decision integration deserves at least as much attention as accuracy: assigning owners for outcomes, building recommendations into routine planning, and reviewing overrides can raise the action rate α . Savings should be tracked by stream, so that a gain in materials is not mistaken for a gain in energy. Table 9 links actions to the mechanisms they target and the indicators that show whether they work.

6.2 Policy Implications

Policy makers can support resource efficient AI adoption through standards for reporting the footprint of AI services, guidance on measurement and verification of savings, and support for the digital and management capabilities that convert recommendations into action. Programs that subsidize AI adoption should be evaluated with resource indicators as well as productivity indicators, and a program that raises output but not efficiency may need complementary measures such as resource pricing. Public agencies can also support data infrastructure so that firm resource, output, and AI usage data can be linked for evaluation without compromising confidentiality.

6.3 Implications for Bangladesh and Emerging Economies

Emerging economies are an important setting for future research because AI services can be adopted quickly through cloud platforms while resource systems, infrastructure, and regulation change at the same time. In Bangladesh, resource intensive export manufacturing and food processing use substantial energy, water, and materials, and forecasting, process optimization, and procurement tools could lower waste and input use, but the effect should be tested with firm level data and not assumed from adoption alone. A purchased AI service moves the footprint to a provider, often abroad, so a firm level ledger should explicitly decide whether and how to count it. Given the variation in electricity mixes, the footprint and the saving in energy terms should be valued with local emission factors.

7. Limitations and Future Research

The study has several limitations. First, it is a structured evidence synthesis rather than a new causal firm level study, and it is not a systematic review with a registered protocol. Second, direct firm level evidence on AI adoption and resource use is very limited, so the framework relies partly on adjacent literatures on productivity, decision science, rebound, and the footprint of AI. Third, AI adoption is measured differently across studies, and the footprint of purchased services is often not observable. Fourth, resource effects may appear with a time lag. Fifth, firm level resource data linked to AI use are rarely available. Sixth, the numerical illustrations, sensitivity analysis, and simulations are hypothetical and use assumed parameters; they demonstrate the logic of the framework and are not estimates. The classification in Figure 6 summarizes the literature reviewed and does not weight studies by quality or sample size. Table 10 pairs each limitation with a response.

The proposed AEDMI is a measurement framework and not a validated universal scale. Future research should validate it with confirmatory factor analysis, panel data, and observed firm records. Priority questions include how large the footprint budget a^* is in practice for different firms, whether decision integration raises gains more than model accuracy does, how the effects differ between large firms and

Table 8. Research Propositions and Their Empirical Tests

No.	Proposition	Test	Expected sign	Key data
P1	AI adoption is associated with higher resource productivity when recommendations are embedded in decisions	Eq. (13), $y = \ln P$: α_1^P	+	Firm panel of AI and P
P2	Decision integration and governance strengthen the association between AI adoption and resource productivity	Eq. (15): β_3	+	AI, P , DI
P3	AI adoption can increase output, and so total resource demand, through productivity and growth	Eq. (13), $y = \ln Y$: α_1^Y	+	Output; capacity
P4	The AI system consumes resources, and the net effect depends on the relative size of efficiency, scale, and footprint effects	Eq. (13), $y = \ln(1 + \phi)$: α_1^ϕ ; Eq. (14): sign of c ; NRS	+, sign of c varies	Cloud logs; utility bills; ledger
P5	Better resource data strengthens the efficiency gain from AI adoption	Eq. (15) with DQ	+	AI, P , DQ

Table 9. Implications for Actors: Actions, Targeted Mechanisms, and Monitoring Indicators

Actor	Recommended action	Mechanism targeted	Indicator to monitor
Firms	Put AI footprint on the resource ledger; assign owners for outcomes; track savings by stream	Footprint; decision integration	Offset ratio O ; action rate α ; stream level NRS
Model and cloud providers	Report energy, water, and hardware footprint by workload and by region	Footprint	Footprint per unit of service
Lenders and investors	Ask for absolute resource use together with intensity, and for AI footprint disclosure	Scale; footprint	Absolute resource use; AI footprint share a
Regulators and standards bodies	Set standards for reporting AI footprint and for verification of savings claims	Footprint; data quality	Disclosure coverage; verified savings
Statistical agencies	Publish firm level resource and output data linkable to AI adoption	Measurement	Availability of linked data; update frequency
Researchers	Report efficiency, scale, and footprint effects together with break even growth and offset ratio	Measurement	NRS; g^* ; O ; weight sensitivity of AEDMI

Table 10. Limitations and Proposed Responses

Limitation	Proposed response
No new causal estimate	Test P1–P5 with firm panels, staggered adoption event studies, and instruments
Thin direct evidence on resource outcomes	Collect firm level resource, output, and AI usage data
Unobserved footprint of purchased AI	Request provider disclosures; use workload based estimates with sensitivity
Time lags	Use multi year panels and dynamic specifications
Limited linked data	Link meter, ERP, procurement, and cloud usage records
Illustrative numbers	Replace assumed values with estimated productivity, growth, and footprint shares

small and medium enterprises, how long rebound effects last, and how the footprint of purchased AI services should be allocated across customers.

8. Conclusion

This study examines whether AI adoption improves corporate resource efficiency. The evidence available before 2024 indicates that the relationship is neither automatically positive nor automatically negative. AI can improve the decisions that determine resource use by forecasting demand, optimizing processes, screening suppliers, routing goods, and monitoring performance. At the same time, the AI system consumes energy, water, and materials, it can raise output, and its recommendations change resource use only when they are acted on. Direct firm level evidence is very limited, and the main conclusion is therefore conditional. AI improves corporate resource efficiency when its efficiency gain exceeds the output it enables and the resources it consumes. The paper’s ledger gives a simple test: with a productivity gain π and an AI footprint equal to a share a of baseline use, resource use falls only if output growth stays below

$$g^* = \pi - a(1 + \pi). \quad (16)$$

The effects of AI adoption may also vary across firms and contexts. Firms that integrate recommendations into routine decisions, assign accountability, hold reliable resource data, and size their models to the task may be more likely to convert AI capability into lower resource use. In contrast, firms that deploy large models mainly to support growth may experience stronger scale and footprint effects. The consequences may differ across resource streams, as the illustration shows, and across domains of decision-making. Future studies should therefore avoid treating AI as a homogeneous technology and instead examine differences across use cases, resource streams, firm sizes, and computing arrangements.

The findings also have implications for firms, AI providers, financial institutions, and policy makers. Adopting AI alone may not be sufficient to achieve resource efficiency objectives. Firms can place the

AI footprint on the resource ledger and invest in the decision processes that turn predictions into action, providers can disclose the footprint of their services, and lenders and regulators can ask for absolute resource use together with intensity. Such approaches would allow the productivity benefits of AI to be retained while directing a larger share of the gains toward lower resource use.

An additional consideration is the rebound effect. Efficiency improvements enabled by AI may lower the cost of production and increase output, so that a reduction in resource use per unit does not necessarily translate into a reduction in total use. The offset ratio in Equation (8) shows how much of the gross saving is consumed: in the illustration, 69% of the gross saving is offset by output growth and the AI footprint, and in the energy and water streams the offsets exceed the gains. Evaluating AI therefore requires intensity based and absolute indicators, together with the resources used by the AI system itself.

There are also important limitations to the available evidence. Much of the existing literature relies on conceptual assessments, estimates of the footprint of models, or aggregate measures of productivity and environmental outcomes. These approaches can identify potential and associations but may not fully establish the causal mechanisms connecting AI adoption, decisions, and resource use. Differences in the measurement of AI adoption, resource productivity, and AI footprint also make comparison across studies difficult. Moreover, rapid changes in generative AI, model efficiency, and cloud computing mean that findings from earlier periods may not fully represent current environments.

Future empirical research should therefore test the framework with firm panel data, resource meter data, cloud usage logs, decision logs, and measures of data quality. Where possible, researchers should distinguish between AI used to optimize physical operations and AI used mainly to expand volume or marketing. Natural experiments, staggered adoption, and longitudinal firm surveys could help identify causal effects and reveal whether efficiency gains persist over time.

Overall, the environmental consequences of AI adoption depend not only on how much AI a firm uses, but also on how it shapes the decisions that determine resource use, the scale of production, and the resources consumed by the AI system itself. The central policy and research question is therefore not whether AI is inherently green, but under what organizational, technological, and energy-system conditions the resources it saves exceed the resources it uses and the output it enables. The paper’s framework provides a basis for making this distinction.

Data Availability

No new primary dataset was collected for this evidence synthesis. The paper is designed to support future empirical testing with firm level resource and production data, cloud usage records, and management surveys. The code used to produce the illustrative figures, the numerical

example, the sensitivity analysis, and the simulation (`make_figures.py`) is available from the authors on request.

References

- Agrawal, A., Gans, J. S., and Goldfarb, A. (2019). Prediction, judgment, and complexity: A theory of decision-making and artificial intelligence. In A. Agrawal, J. Gans, and A. Goldfarb (Eds.), *The Economics of Artificial Intelligence: An Agenda*, 89–110. University of Chicago Press.
- Anderson, S. T., and Newell, R. G. (2004). Information programs for technology adoption: The case of energy-efficiency audits. *Resource and Energy Economics*, 26(1), 27–50. <https://doi.org/10.1016/j.reseneeco.2003.07.001>
- Ang, B. W. (2004). Decomposition analysis for policymaking in energy: Which is the preferred method? *Energy Policy*, 32(9), 1131–1139. [https://doi.org/10.1016/S0301-4215\(03\)00076-4](https://doi.org/10.1016/S0301-4215(03)00076-4)
- Berg, F., Kölbel, J. F., and Rigobon, R. (2022). Aggregate confusion: The divergence of ESG ratings. *Review of Finance*, 26(6), 1315–1344. <https://doi.org/10.1093/rof/rfac033>
- Alam, M. R. U., and Nahid, M. A. H. (2023). A socio-technical analysis of employee resistance to AI automation during Industry 4.0 migrations. *Digital Transformation and Technology Dynamics*, 3(1).
- Hasan, S., Rahman, M. M., and Ahmed, M. S. (2022). Renewable energy transition and microgrid adoption in manufacturing SMEs: A structural equation and technoeconomic analysis. *Eco-Business and Environmental Progress Journal*, 2(1).
- Alvi, M. F. I., Haque, M. E., and Hasan, M. (2022). Decarbonizing urban freight logistics: Evaluating the environmental and economic impacts of electric fleet transition policies. *Eco-Business and Environmental Progress Journal*, 2(1).
- Gupta, M. S., Alam, M. R. U., and Islam, S. (2023). The impact of ESG disclosure quality on corporate cost of capital and financial performance: Panel evidence from Asian manufacturing firms. *Eco-Business and Environmental Progress Journal*, 3(1).
- Bloom, N., and Van Reenen, J. (2007). Measuring and explaining management practices across firms and countries. *Quarterly Journal of Economics*, 122(4), 1351–1408. <https://doi.org/10.1162/qjec.2007.122.4.1351>
- Boyd, G. A., and Curtis, E. M. (2014). Evidence of an “energy-management gap” in U.S. manufacturing: Spillovers from firm management practices to energy efficiency. *Journal of Environmental Economics and Management*, 68(3), 463–479.
- Brookes, L. (1990). The greenhouse effect: The fallacies in the energy efficiency solution. *Energy Policy*, 18(2), 199–201. [https://doi.org/10.1016/0301-4215\(90\)90145-T](https://doi.org/10.1016/0301-4215(90)90145-T)
- Brynjolfsson, E., Hitt, L. M., and Kim, H. H. (2011). Strength in numbers: How does data-driven decisionmaking affect firm performance? SSRN Working Paper No. 1819486. <https://doi.org/10.2139/ssrn.1819486>
- Brynjolfsson, E., Li, D., and Raymond, L. R. (2023). Generative AI at work. NBER Working Paper No. 31161. <https://doi.org/10.3386/w31161>
- Brynjolfsson, E., and McElheran, K. (2016). The rapid adoption of data-driven decision-making. *American Economic Review*, 106(5), 133–139. <https://doi.org/10.1257/aer.p20161016>
- Brynjolfsson, E., Rock, D., and Syverson, C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1), 333–372. <https://doi.org/10.1257/mac.20180386>
- Callaway, B., and Sant’Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2), 200–230. <https://doi.org/10.1016/j.jeconom.2020.12.001>
- Czarnitzki, D., Fernández, G. P., and Rammer, C. (2023). Artificial intelligence and firm-level productivity. *Journal of Economic Behavior & Organization*, 211, 188–205.
- Dietvorst, B. J., Simmons, J. P., and Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114–126. <https://doi.org/10.1037/xge0000033>
- Freitag, C., Berners-Lee, M., Widdicks, K., Knowles, B., Blair, G. S., and Friday, A. (2021). The real climate and transformative impact of ICT: A critique of estimates, trends, and regulations. *Patterns*, 2(9), 100340. <https://doi.org/10.1016/j.patter.2021.100340>
- Gerarden, T. D., Newell, R. G., and Stavins, R. N. (2017). Assessing the energy-efficiency gap. *Journal of Economic Literature*, 55(4), 1486–1525. <https://doi.org/10.1257/jel.20161360>
- Gillingham, K., Rapson, D., and Wagner, G. (2016). The rebound effect and energy efficiency policy. *Review of Environmental Economics and Policy*, 10(1), 68–88. <https://doi.org/10.1093/reep/rev017>
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2), 254–277. <https://doi.org/10.1016/j.jeconom.2021.03.014>
- Greening, L. A., Greene, D. L., and Difiglio, C. (2000). Energy efficiency and consumption: The rebound effect, a survey. *Energy Policy*, 28(6–7), 389–401. [https://doi.org/10.1016/S0301-4215\(00\)00021-5](https://doi.org/10.1016/S0301-4215(00)00021-5)
- Kaack, L. H., Donti, P. L., Strubell, E., Kamiya, G., Creutzig, F., and Rolnick, D. (2022). Aligning artificial intelligence with climate change mitigation. *Nature Climate Change*, 12, 518–527.
- Khazzoom, J. D. (1980). Economic implications of mandated efficiency standards for household appliances. *The Energy Journal*, 1(4), 21–40.
- Li, P., Yang, J., Islam, M. A., and Ren, S. (2023). Making AI less “thirsty”: Uncovering and addressing the secret water footprint of AI models. arXiv preprint arXiv:2304.03271. <https://doi.org/10.48550/arXiv.2304.03271>
- Liu, J., Liu, L., Qian, Y., and Song, S. (2022). The effect of artificial intelligence on carbon intensity: Evidence from China’s industrial sector. *Socio-Economic Planning Sciences*, 83, 101002.
- Logg, J. M., Minson, J. A., and Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, 151, 90–103. <https://doi.org/10.1016/j.obhdp.2018.12.005>
- Luccioni, A. S., Viguiet, S., and Ligozat, A.-L. (2023). Estimating the carbon footprint of BLOOM, a 176B parameter language model. *Journal of Machine Learning Research*, 24(253), 1–15.
- Masanet, E., Shehabi, A., Lei, N., Smith, S., and Koomey, J. (2020). Recalibrating global data center energy-use estimates. *Science*, 367(6481), 984–986. <https://doi.org/10.1126/science.aba3758>
- Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L.-M., Rothchild, D., So, D., Texier, M., and Dean, J. (2021). Carbon emissions and large neural network training. arXiv preprint arXiv:2104.10350.
- Porter, M. E., and van der Linde, C. (1995). Toward a new conception of the environment-competitiveness relationship. *Journal of Economic Perspectives*, 9(4), 97–118. <https://doi.org/10.1257/jep.9.4.97>
- Rolnick, D., Donti, P. L., Kaack, L. H., Kochanski, K., Lacoste, A., Sankaran, K., Ross, A. S., Milojevic-Dupont, N., Jaques, N., Waldman-Brown, A., Luccioni, A. S., Maharaj, T., Sherwin, E. D., Muckavilli, S. K., Kording, K. P., Gomes, C. P., Ng, A. Y., Hassabis, D., Platt, J. C., Creutzig, F., Chayes, J., and Bengio, Y. (2022). Tackling climate change with machine learning. *ACM Computing Surveys*, 55(2), Article 42. <https://doi.org/10.1145/3485128>
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69(1), 99–118. <https://doi.org/10.2307/1884852>
- Sorrell, S., and Dimitropoulos, J. (2008). The rebound effect: Microeconomic definitions, limitations and extensions. *Ecological Economics*, 65(3), 636–649. <https://doi.org/10.1016/j.ecolecon.2007.08.013>
- Strubell, E., Ganesh, A., and McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3645–3650. <https://doi.org/10.18653/v1/P19-1355>
- Sun, L., and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2), 175–199. <https://doi.org/10.1016/j.jeconom.2020.09.006>
- Vinuesa, R., Azizpour, H., Leite, I., Balaam, M., Dignum, V., Domisch, S., Felländer, A., Langhans, S. D., Tegmark, M., and Fuso Nerini, F. (2020). The role of artificial intelligence in achieving the Sustainable Development Goals. *Nature Communications*, 11, 233.
- Wu, C.-J., Raghavendra, R., Gupta, U., Acun, B., Ardalani, N., Maeng, K., Chang, G., Aga, F., Huang, J., Bai, C., and others (2022). Sustainable AI: Environmental implications, challenges and opportunities. In *Proceedings of Machine Learning and Systems*, 4, 795–813.