

<https://doi.org/10.1038/s44168-026-00411-0>

AI-driven productivity gains enable more CO₂ emissions than they avoid in a global energy–economy model



Will Alpine¹✉, Nathan Geldner², Holly Alpine¹ & Maksym G. Chepeliev³

The net climate impacts of artificial intelligence (AI) depend largely on how its applications propagate through competing energy pathways. Predominant analyses examine the relationship between datacenter energy demand, renewables optimization, and demand-side efficiencies, but insufficiently address how AI also reshapes fossil fuel supply economics. We instead model AI as a bidirectional productivity amplifier in a global computable general equilibrium model, quantifying both enabled emissions from fossil fuel productivity gains and avoided emissions from renewables productivity gains. Under parallel adoption scenarios, net annual CO₂ emissions increase by 0.47–1.8 gigatonnes (1.2–4.8% of 2024 global energy-related CO₂ emissions). Enabled emissions exceed avoided emissions whenever fossil-sector gains are nonzero; net emissions reductions require renewables gains 4–5× greater than fossil fuel gains. Absent policy steering, AI's modeled effects increase the carbon intensity of the global economy and reinforce fossil fuel incumbency—outcomes that current analytical and governance frameworks do not fully capture.

Climate change presents profound societal and economic risks¹, and limiting global warming to 1.5 °C requires rapid decarbonization centered on phasing out fossil fuels². Yet fossil fuels continue to supply over 80% of global primary energy³ and global consumption continues to rise⁴, despite renewables' increasing deployment⁵ and cost parity⁶. This reflects energy addition rather than substitution: renewables have historically expanded total energy supply while fossil fuels remain largely undisplaced⁷. Beyond electricity, fossil fuels also remain deeply embedded in industrial pathways throughout the broader global economy. Furthermore, path-dependent co-evolution of technological systems with governing institutions reinforces incumbent energy regimes⁸. Together, these dynamics complicate efforts to align technological progress with internationally agreed climate targets.

Artificial intelligence (AI) is a general-purpose productivity technology diffusing rapidly across the global economy, yet its interaction with these energy system dynamics remains underexamined. AI's climate impacts extend beyond the direct operational footprint of its computational infrastructure to the broader effects of its applications across the economy. As a bidirectional productivity amplifier, AI's applications enhance productivity across both high-carbon and low-carbon energy supply pathways simultaneously. Whether AI enables more emissions than it avoids depends on the global balance between these competing pathways.

Assessing this net effect requires understanding how technology is applied across both high- and low-carbon energy supply systems. Fossil fuel production faces a structural constraint that makes technological

reinforcement particularly consequential: in the absence of continuous reinvestment, oil and gas production would fall by approximately 8% annually—a dynamic the International Energy Agency (IEA) characterizes as requiring the industry to “run much faster just to stand still”⁹. The IEA noted as early as 2005 that technological progress has repeatedly delayed “peak oil” forecasts by unlocking resources previously considered unviable¹⁰. AI extends this dynamic, pushing back peak-supply expectations in at least some fossil sectors today¹¹. Operating across the entire value chain, its applications predominantly expand the pool of economically viable supply by increasing extraction productivity¹², lowering production costs¹³, and reducing operational risk¹⁴, thus extending the industry's economic viability (see Methods; Supplementary Information (SI) Notes S3.1.1 and S8). While these applications predominantly enable additional emissions resulting from increased fossil fuel throughput, the same capabilities can also avoid emissions within fossil operations or end-uses, for example, through methane leak detection¹⁴ or efficiency gains in power generation¹⁵. Applied to renewables (and other low-carbon generation), AI's applications predominantly avoid emissions: they offer potential to accelerate deployment and optimize generation efficiency, for example, through forecasting, predictive maintenance, and power generation optimization, while also extending the productive life of installed capacity and improving grid integration¹⁴ (see Methods; SI Note S3.1.2). This pathway-level balance between high- and low-carbon systems is amplified by the broader economic structure through which AI's productivity gains propagate.

¹Enabled Emissions Campaign, Seattle, WA, USA. ²Unaffiliated, Seattle, WA, USA. ³Purdue University, West Lafayette, IN, USA.

✉e-mail: will@enabledemissions.com

At a systemic level, this raises the question of whether AI's net effect reinforces or displaces incumbent energy systems. We characterize displacement as the following: eroding fossil supply economics, reducing total fossil consumption, and reducing economic activity within fossil-embedded systems; reinforcement is the inverse. The answer has significant implications for energy transition pathways, given that AI is enhancing productivity within an energy-economic system still anchored in fossil fuels. Answering this systemic question is complicated by two compounding mechanisms through which AI's applications generate indirect, system-wide emissions effects. First, productivity gains from digital optimization frequently generate rebound effects, defined as efficiency improvements that stimulate additional consumption¹⁶. Second, AI's impacts extend beyond these familiar dynamics to lesser-studied induction effects, in which new technological options reshape production economics by altering capital, labor, and resource productivity¹⁷. Together, rebound and induction effects suggest that AI's most consequential climate impacts may arise through indirect, economy-wide channels.

Although these indirect, system-wide impacts of its applications may substantially exceed AI's comparatively well-studied direct operational footprint from computational infrastructure^{14,18–20}, they remain largely unquantified^{21,22}. This gap is visible in three limitations of existing methodological approaches. First, industry and solution-oriented assessments predominantly emphasize the potential of emissions-reducing applications, often leaving emissions-increasing uses outside the defined analytical boundary^{23,24}. Second, prevailing quantitative frameworks are not designed to capture economy-wide, cross-sector dynamics. Partial-equilibrium approaches risk underrepresenting economy-wide feedbacks²⁵; avoided-emissions methodologies omit system-level rebound and induction effects within their stated boundaries²⁶; and even net carbon impact approaches that identify AI-enabled fossil expansion as a higher-order mechanism acknowledge that rigorous quantification remains beyond their present scope²⁷. Third, the climate modeling architectures most commonly used for long-term projections, including Integrated Assessment Models (IAMs), are less suited to resolving the near-term, cross-sector supply dynamics through which AI reshapes competing energy pathways²⁸. An extended discussion of these gaps is provided in SI Note S2.2.

Together, these gaps give rise to a broader pattern in analyses of AI's net climate effect, illustrated by two prominent examples. First, even analyses that explicitly call for quantitative frameworks capable of capturing the full range

of indirect effects, such as Luers et al.²², nonetheless quantitatively frame AI's net impact primarily as a trade-off between operational footprints and avoided-emissions potential²⁹, leaving emissions-increasing pathways outside the analytical boundary (see SI Note S2.1). Second, the IEA's 2025 Energy and AI report quantifies a limited set of emissions-increasing effects but does not provide a systematic treatment, and its methodology is not publicly available, precluding independent assessment or reproduction¹⁴ (see SI Note S2.1). Because governance frameworks reflect the boundaries of the analyses that inform them, this analytical gap propagates into a governance gap: policy, disclosure, and assessment frameworks built on a narrow framing centered on operational footprints and avoided-emissions potential cannot address emissions pathways that their underlying analyses do not recognize. In practice, these indirect system-level emissions-increasing effects currently remain outside the substantive scope of prominent AI and climate governance frameworks^{30,31}.

To address these gaps, we present a global, economy-wide quantitative assessment of AI as a bidirectional productivity amplifier operating across competing energy system pathways. We modeled two AI-driven energy supply pathways by applying productivity gains in both fossil fuels and renewables, distinguishing their predominant emissions consequences as enabled and avoided emissions, respectively (Box 1). Drawing on empirical evidence to parameterize sectoral productivity gains, we employed a computable general equilibrium (CGE) model that quantified enabled and avoided emissions within a unified equilibrium structure. We additionally assessed whether fuel-neutral pathways (such as grid infrastructure improvements and demand-side efficiency gains) moderated or amplified the supply-side balance. We then tested whether AI-driven productivity gains caused enabled emissions to exceed avoided emissions within a predominantly fossil-based energy system, and whether they increased the carbon intensity of economic growth. Next, we assessed whether the modeled equilibrium reinforced or displaced incumbent carbon-intensive systems using the criteria defined above. Finally, we translated these findings into governance priorities for AI and climate policy.

Results

Overview

Using a CGE model, we quantified the CO₂ emissions consequences of AI-driven productivity gains across two competing supply-side pathways (fossil fuels and renewables) alongside fuel-neutral applications (grid

Box 1 | AI as a bidirectional productivity amplifier

This study treats AI as a general-purpose productivity-amplifying technology within a competing-pathways framework. The schematic (Fig. 1) illustrates the model's analytical boundary, including supply-demand interactions, deployment bottlenecks, and system-level productivity effects. We adopted a consequential analysis: emissions outcomes are determined by economy-wide market responses, rather than by attributional methods that trace emissions within fixed system boundaries.

Model approach

We used a computable general equilibrium (CGE) model to represent AI adoption as productivity shocks across competing energy sectors. "Productivity" describes AI's effect on sectoral output per unit input (capital, labor, energy, etc.); "efficiency" is reserved for established technical terms (e.g., energy, grid, or extraction efficiency). Unlike isolated efficiency metrics, productivity shocks capture system-wide effects: output expansion, cross-sector substitution, factor reallocation, and price feedbacks. For example, AI in upstream oil extraction can reduce marginal costs and increase extraction efficiency, inducing a global expansion in economically viable supply, whereas AI in solar forecasting can raise generation efficiency from installed capacity.

Impact taxonomy

Following the European Green Digital Coalition (EGDC) taxonomy²⁷, we distinguish across the following impacts: first-order effects (direct operational impacts of computing infrastructure, such as datacenter energy use), second-order effects (the indirect consequences of AI applications), and higher-order effects (systemic dynamics, chiefly rebound and induction effects, which propagate economy-wide). This study addresses second- and higher-order effects.

Consequence categories

Within these effects, we distinguish enabled and avoided emissions as consequence categories, not pathway labels: the same technology can produce either outcome (e.g., enhanced extraction enables emissions; methane leak detection avoids them). Enabled and avoided emissions are, respectively, additional and reduced CO₂ attributable to tech-driven gains. Fuel-neutral pathways (infrastructure, demand-side applications) were modeled symmetrically and can yield either outcome. Higher-order dynamics are detailed in Methods (SI Notes S4 and S5); SI Note S2.5 gives terminological context.




Fig. 1 | Conceptual systems diagram of AI-energy market dynamics across first-, second-, and higher-order effects. Modified Causal Loop Diagram⁴⁹ representing energy market dynamics across first-order (operational), second-order (avoided/enabled), and higher-order (amplification) effects. Technology is modeled as productivity shocks applied to competing energy pathways: Path 1 (fossil productivity

gains) and Path 2 (renewables optimization). The diagram includes both modeled and unmodeled mechanisms to illustrate whole-system dynamics. Notable deployment barriers are indicated for system comprehension but not explicitly modeled. Element sizes are representative, not quantitative. Designer credit: Rachel Woods-Robinson

infrastructure and select demand-side pathways). We first calculated the net emissions effect of AI-driven fossil fuel and renewables supply-side productivity shocks under parallel adoption assumptions, and related these estimates to first-order datacenter emissions for reference. We then applied uniform productivity shocks across fossil and renewables pathways to identify the structural asymmetry and breakeven threshold to achieve net emissions reductions. Next, we evaluated empirically derived scenarios in which adoption rates vary independently across pathways, testing robustness through elasticity variation, baseline comparison, and carbon pricing. We then assessed the impact of fuel-neutral applications (grid infrastructure and select demand-side pathways) on the supply-side asymmetry. Finally, we assessed macroeconomic implications, examining how AI productivity gains affect the growth of emissions relative to GDP. These results characterize comparative-static equilibria, isolating the effect of the modeled productivity shocks while holding other structural change constant (*ceteris paribus*); they describe the direction and approximate scale of economic incentives under current conditions, not time-resolved forecasts. Real-world outcomes are shaped by co-evolving economic dynamics that the static modeling framework used here does not capture, and are likely to differ in magnitude while maintaining directionality (see Discussion; Methods).

Net global emissions increase under parallel adoption

Under parallel adoption assumptions—a neutral benchmark representing all pathways at the same adoption tier (a methodologically conservative simplification; see “Methods”)—AI-driven productivity increased net global annual CO₂ emissions by 0.47–1.8 Gt, equivalent to 1.2–4.8% of total global energy-related CO₂ emissions in 2024⁵. Isolating individual pathways revealed that fossil fuel productivity gains were the primary contributor to the net increase, while renewables optimization offset only part of that effect, as illustrated in Fig. 2.

Enabled emissions (0.6–2.4 Gt CO₂ annually) exceed the IEA’s 2025 datacenter emissions estimates (0.18 Gt CO₂ annually)¹⁴ by 3.3–13.3×, and its projected 2035 datacenter emissions (0.3–0.5 Gt CO₂ annually) by 1.2–8× (Fig. 2). This contextualizes magnitude only; it is not an aggregation across analytical frameworks. Modeled net second-order effects under parallel adoption exceed the IEA’s 2025 datacenter estimates across all modeled scenarios and their projected 2035 estimates across most (see SI Note S2.3).

Renewables gains must exceed fossil gains by 4–5× to reach breakeven

With uniform productivity shocks applied simultaneously across fossil fuel and renewables pathways, renewables productivity gains must exceed fossil

Fig. 2 | AI emissions (Modeled). Annual CO₂ emissions changes attributable to AI-mediated productivity gains are shown across modeled pathways. Gray bars show unmodeled first-order data-center estimates from the IEA¹⁴ for reference only; these are not directly additive to the modeled second-order effects (see Discussion; SI Note S2.3). Blue bars show avoided emissions from renewables-only productivity shocks applied at low, medium, and high adoption levels, with no fossil shocks. Red bars show enabled emissions from fossil-only productivity shocks applied at low, medium, and high adoption levels, with no renewables shocks. Purple bars show the modeled net effect under parallel adoption, in which fossil and renewables supply pathways receive productivity shocks at the same adoption tier (see Methods) indicating an annual CO₂ emissions increase of 0.47–1.8 Gt. Because the net effect is modeled under simultaneous system-wide shocks, it is not equal to the sum of the isolated pathway effects. Infrastructure and demand-side efficiencies are not visualized; variations are explored in SI Figs. S3–S13. Additional scenario-based outcomes are shown in Fig. 4. Designer credit: Camille Hu





Fig. 3 | Change in annual CO₂ emissions: fossil fuels and renewables. Estimated net annual CO₂ emissions (gigatonnes, Gt) under uniform productivity shocks to fossil fuels (extraction, refining, and generation) and renewables (generation) sectors, respectively. Fossil productivity increases linearly from bottom to top, while renewables productivity increases linearly from left to right. Productivity gains are shown as percentage changes relative to baseline; upper ranges are illustrative rather than realistic given available evidence. Dark red cells indicate larger increases in CO₂ emissions while dark green cells indicate larger reductions in CO₂ emissions. Gray dashed line indicates the net-zero emissions isoquant (breakeven), with average slope of approximately 1:5 (corresponding to 4–5× modeled asymmetry mentioned throughout). Shock types differ by pathway (see “Methods”). Fuel-neutral productivity gains are excluded.



gains by approximately 4–5× for net emissions to break even: each 1% marginal productivity gain in fossil fuels requires 4–5% gains in renewables (Fig. 3). Uniform shocks consistently increased net emissions across the full modeled shock range—from no AI productivity gains to a doubling of productivity (this upper range is illustrative rather than realistic given available evidence)—indicating that the asymmetry is not limited to specific scenario parameterizations (see “Methods”).

Upstream fossil fuel extraction productivity is the primary driver of this asymmetry: small fossil extraction gains are sufficient to outweigh much larger improvements in renewables generation. This is revealed by isolating fossil fuel shocks across both extraction and generation relative to the same renewables baseline: extraction-only shocks (Fig. S1) produced outcomes comparable to combined shocks, while generation-only shocks (Fig. S2) showed slight emissions reductions across all scenarios, confirming that the asymmetry originates in upstream fossil fuel productivity gains.

Directional asymmetry is robust across scenarios, baselines, and parameter variation

The 4–5× modeled asymmetry between fossil fuels and renewables persisted across 64 empirically derived scenario combinations, and was robust to adjustment across baseline assumptions, parameter variation, and energy

mix (parameterization in SI Note S3). In parallel adoption scenarios, net emissions increased in every case (Fig. 4). Net emissions reductions occurred only when fossil-sector productivity gains were zero; emissions increased even when pairing upper-bound renewables assumptions (15–20% generation gains for mainstream technologies; higher for nascent “other” categories) with limited fossil productivity improvements (2–5% upstream gains).

To test whether the asymmetry reflects the current electricity generation mix rather than deeper economic structure, we compared results from our primary 2024-adjusted baseline against the original 2017 Global Trade Analysis Project (GTAP) Data Base baseline, which has a lower renewables share (see “Methods”). Despite this difference, the enabled-to-avoided ratio changed only marginally (Fig. S22), indicating that the core asymmetry is robust to short-term changes in the electricity generation mix.

Directional robustness was further confirmed by varying two key elasticities by $\pm 50\%$ (capital-energy substitution and cross-technology substitutability in power generation), which shifted the enabled-to-avoided ratio from approximately 2–3× (net increase of 1.0–1.3 Gt CO₂, under lower capital-energy substitution and higher cross-technology substitutability) to approximately 7× (net increase of 2.4 Gt CO₂, under the reverse conditions) as shown in Figs. S18–S21.




Fig. 4 | Net annual CO₂ emissions changes under scenario-based productivity shocks across fossil, renewables, and fuel-neutral pathways. Changes in net annual CO₂ emissions attributable to fossil fuel combustion (Gt) under scenario-based productivity shocks are shown across 64 combinations spanning four adoption tiers per pathway (4 fossil × 4 renewables × 4 fuel-neutral). Panels indicate differing fuel-neutral adoption levels: **a**, none; **b**, low; **c**, medium; **d**, high. Within each panel, fossil

productivity increases from bottom to top and renewables productivity increases from left to right. Dashed diagonals indicate parallel adoption scenarios. Dark red cells indicate larger increases in CO₂ emissions while dark green cells indicate larger reductions. Negative net emissions occur only in scenarios with zero fossil productivity gains. Extraction-only and generation-only fossil shocks are shown in Figs. S12–S13.

Fuel-neutral efficiency gains do not reverse the asymmetry

To assess whether fuel-neutral productivity gains can reduce emissions enough to offset the emissions-increasing supply-side dynamics, we modeled improvements in power infrastructure efficiency (grid transmission and distribution) and a subset of demand-side end-use efficiency alongside supply-side pathways. The demand-side subset covers energy-intensive industry and maritime shipping—together about 25% of the IEA's projected end-use efficiency savings from widespread AI adoption by 2035^{4,14} (see Methods; SI Note S3.3).

Across the modeled scenarios, fuel-neutral improvements moderated net emissions but did not reverse the supply-side effects. Under parallel adoption assumptions, a high fuel-neutral adoption scenario (fossil and renewables shocks excluded) reduced net emissions by only approximately 0.1 Gt CO₂ annually (Fig. 4). The 4–5× asymmetry persisted across sensitivity analyses incorporating fuel-neutral shocks alongside supply-side pathways (Figs. S3–S11), confirming the asymmetry's supply-side origin.

Carbon pricing narrows—but does not reverse—the asymmetry

Under carbon pricing at \$80/tCO₂, enabled emissions (1.0 Gt) still exceeded avoided emissions (0.3 Gt) by approximately 3×, producing a net increase of 0.7 Gt CO₂. At \$308/tCO₂, the gap narrowed further: enabled emissions (0.3 Gt) exceeded avoided emissions (0.2 Gt) by approximately 1.5×, producing a net increase of 0.1 Gt CO₂ (Figs. S14–S17).

Across all evaluated carbon price levels, the asymmetry was narrowed, but not eliminated. Net emissions decreased only when renewables gains substantially outpaced fossil gains: at \$80/tCO₂, this requires zero fossil productivity gains; at \$308/tCO₂, the threshold eased, with low fossil gains paired with medium or higher renewables gains producing net reductions.

Carbon intensity of economic growth increases under parallel adoption

To assess the macroeconomic implications of these findings, we compared emissions growth to GDP growth across scenarios. Under parallel adoption scenarios, AI productivity increased CO₂ emissions faster than economic growth, indicating negative decoupling within this static framework. A reversal of this pattern, where emissions fall or stabilize as the economy grows (absolute decoupling, a prerequisite for the conditions associated with environmentally sustainable growth), occurred only when supply-side AI productivity gains were zero. Even the more modest goal (relative decoupling, where emissions grow more slowly than GDP) requires renewables and fuel-neutral adoption to outpace fossil fuel productivity gains; parallel adoption is insufficient. AI-driven productivity consistently increased the emissions intensity of GDP across tested conditions. Consistent with the pathway and value-chain decompositions above, fossil fuel productivity gains primarily determined these decoupling outcomes; shock-type decomposition further attributed the effect predominantly to upstream fossil fuel productivity propagating through the broader economy, consistent with induction effects (see SI Note S2.4 and Figs. S23–S26).

Discussion

Our findings challenge the predominant framing of AI's net climate impact as a tradeoff between the energy it consumes and the emissions it helps avoid. Instead, we characterize AI as a bidirectional productivity amplifier: the same capabilities that optimize renewables generation and end-use efficiency also sustain and expand the economic viability of fossil fuels, with asymmetric emissions consequences.

Within modeled scenarios, enabled emissions, predominantly driven by fossil fuel supply pathways, consistently exceed avoided emissions from renewables optimization. Parallel adoption of AI across fossil fuel and renewable energy pathways increases net global annual emissions by 0.47–1.8 Gt CO₂ (1.2–4.8% of global 2024 energy-related CO₂ emissions)—a conservative benchmark, as discussed below. Avoided emissions exceed enabled emissions only when renewables productivity gains outpace fossil gains by 4–5×. Furthermore, fuel-neutral gains (including grid efficiency and select demand-side end-use efficiency) moderate but do not reverse the

asymmetry. These findings indicate that AI productivity gains, as modeled, help sustain fossil supply economics and increase total fossil consumption, leading to an increase in net global emissions.

At the macroeconomic level, AI-driven productivity shocks increase the carbon intensity of gross domestic product (GDP) under parallel adoption assumptions, indicating how AI's productivity gains amplify economic activity within fossil-embedded systems. Within our static framework, this pattern is inconsistent with the conditions associated with environmentally sustainable economic growth; even in economies where emissions intensities have been falling, our modeled equilibrium shows AI-driven productivity would partially counteract rather than reinforce this pattern (see SI Note S2.4).

This outcome reflects the underlying conditions of the economy that AI productivity gains amplify. With approximately 80% of global primary energy still fossil-based, proportionate AI-driven productivity improvements applied evenly across fossil and low-carbon systems under parallel adoption induce disproportionate emissions—a scale effect of the asymmetric baseline tied to economic activity within carbon-intensive systems. This effect is driven predominantly by upstream fossil fuel extraction productivity gains, whose impacts propagate economy-wide through rebound and induction effects. The 80% fossil base provides the starting point for understanding the asymmetry, but its persistence reflects more than the energy mix alone. The asymmetry persists at a similar magnitude when renewables productivity is adjusted to match the 2024 electricity generation mix (see Methods; Fig. S22), indicating that fossil fuels are embedded in industrial, chemical, and material systems that renewables penetration alone does not readily reach in the short term. AI's productivity gains propagate across these embedded systems through mechanisms such as cross-sector substitution, factor reallocation, and price feedback.

In aggregate, under market forces alone, AI's modeled equilibrium meets our criteria for reinforcement rather than displacement of incumbent carbon-intensive pathways: it sustains fossil supply economics, increases total fossil consumption, and amplifies economic activity within these fossil-embedded systems. To interrupt this dynamic, frameworks must close the analytical-governance gap by distinguishing enabled from avoided emissions and tracing how productivity gains propagate through a predominantly fossil-based global economy. On the basis of these results, we offer five governance priorities.

First, governance regimes must recognize enabled emissions as a distinct policy-relevant category and constrain AI-enabled increases in fossil fuel productivity. Analytical recognition is foundational: given that enabled emissions dominate net outcomes across modeled scenarios, any analysis that leaves emissions-increasing pathways outside the analytical boundary—as prevailing assessments of AI's net climate impact often do²⁴—will remain analytically incomplete. Supply-side production constraints are the governance complement: enabled emissions constitute the largest modeled contribution to net impact, and such constraints are well-established as a necessary complement to demand-side interventions in achieving net emissions reductions^{32–34}. Extending this logic to AI governance means treating limits on AI-enabled fossil fuel productivity as a distinct policy lever.

Second, research and governance need to address both first-order impacts (the direct operational impacts of computing infrastructure) and second- and higher-order effects (the broader indirect economic consequences of AI's applications) in tandem. Our modeled enabled emissions exceed the IEA's datacenter estimates, by 3.3–13.3×, suggesting that the indirect economy-wide consequences of AI-driven productivity gains substantially exceed AI's direct operational impacts. Frameworks limited to operational footprints therefore risk omitting the larger emissions consequences arising from AI's applications. Moreover, datacenter electricity demand and AI-driven supply-side productivity gains are coupled and mutually reinforcing: datacenter load draws on the same fossil-intensive energy system whose economic viability AI simultaneously enhances globally, underscoring why AI's operational footprint and the economy-wide impacts of its applications must be addressed in tandem (see SI

Note S2.3). Accordingly, an integrated research and governance agenda that treats operational and indirect impacts as connected rather than separate concerns is needed to capture AI's full climate footprint.

Third, where governance regimes pursue AI-driven renewables productivity gains, it should pair them with supply-side constraints on AI-enabled fossil productivity. AI-driven renewables productivity gains alone, even at scale, do not produce net emissions reductions absent constraints on fossil productivity. As modeled, net emissions reductions occur only when fossil productivity gains are zero; available evidence indicates that realized fossil productivity gains are already substantial, with further gains expected (see SI Note S3.1.1). The breakeven asymmetry of 4–5 $\times$ reveals the scale of the requisite renewables productivity gains, but realizing gains of this magnitude depends on resolving compound deployment barriers: renewables-side constraints (permitting delays, interconnection backlogs, geographic variability) and related grid-side AI adoption barriers that the IEA characterizes as predominantly institutional rather than technical¹⁴ (see SI Notes S3.1.2, S3.2). AI optimization may be able to partially address these barriers, but cannot fully resolve either set. Our findings, together with these barriers, demonstrate the insufficiency of exclusively pursuing AI-driven renewables acceleration. More fundamentally, this approach risks reinforcing a broader pattern that energy transition scholarship cautions against: exclusive stimulation of green innovation can paradoxically protect incumbent regimes by directing attention away from fossil fuel incumbency³⁵. Our findings echo this caution: focusing governance effort on AI-driven renewables acceleration, without addressing supply-side fossil dynamics, risks under-deploying the more effective levers our results identify.

Fourth, AI-driven efficiency gains—within fuel-neutral applications and fossil operations alike—should be evaluated for their whole-system effect rather than credited as inherently decarbonizing. In our modeled scenarios, fuel-neutral improvements moderate but do not reverse the direction of net outcomes; net emissions reductions from fuel-neutral gains are achieved only when upstream fossil fuel productivity gains are constrained (Fig. 4 and SI Figs. S3–S11). This pattern is consistent with broader scholarship showing that demand-side measures complement—rather than substitute for—supply-side intervention³². Efficiency improvements within fossil operations are similarly bounded: our shock decompositions show that downstream AI-driven generation efficiency gains reduce emissions only modestly relative to upstream emissions-increasing effects (Figs. S1–S2, S26). This distinction is obscured when efficiency gains that lower carbon intensity within fossil operations are reported as environmental benefits in isolation from the throughput increases that the same systems enable³⁶. The translation of efficiency gains into system-wide emissions reductions is further constrained by rebound effects^{36,37,38}, limiting the decarbonizing potential of fossil-side efficiency. Efficiency gains should therefore not be assumed to reduce net emissions, even where they lower carbon intensity, without accounting for the whole-system interactions that determine their actual net effect.

Fifth, because price signals alone moderate but do not reverse the directional outcome, carbon pricing should operate as part of a broader portfolio that includes supply-side constraints on AI-enabled fossil productivity gains. Economy-wide price signals, as modeled, reach beyond the electricity generation system to the embedded sectors where fossil fuels are most deeply entrenched, moderating the structural advantage of fossil productivity gains. However, net emissions increase persist under parallel adoption across all modeled price levels, including at \$308/tCO₂, a level substantially above emissions-weighted carbon tax rates currently in effect in any country³⁹. Of the levers modeled, carbon pricing produced the largest reduction in net emissions. This suggests that structural economic interventions are a necessary component of an effective portfolio and that reliance on AI's applications alone (such as AI-driven renewables or end-use efficiency gains) is insufficient to drive net emissions reductions.

These five recommendations rest on findings that are directionally robust under sensitivity analysis and conservatively bounded by the modeling assumptions. Two modeling assumptions bias estimates toward

understating net emissions increases (see SI Notes S3 and S5). First, headline net impact calculations use symmetric adoption rates despite documented incumbent advantages for fossil applications (see Methods; SI Note S3). Second, renewables productivity gains are parameterized at the upper bounds of documented technical potential—a deliberately optimistic choice that, given weaker realized-deployment evidence than for fossil applications, likely understates the fossil-renewables gap (see SI Note S3.1.2). Excluded higher-order dynamics, including infrastructure lock-in, political-economy dynamics, and path-dependent investment, are expected to further amplify the understated net effect (see SI Notes S4 and S5). Independently, sensitivity analysis confirms directional robustness: elasticity variation shifts the enabled-to-avoided ratio without reversing its direction. In sum, across the scenarios, conservative bounds, and sensitivity tests examined here, any residual uncertainty is expected to affect magnitude but not directionality.

Achieving net emissions reductions from AI-driven productivity gains is possible, but doing so requires a whole-system governance framework that coordinates intervention across all five of these priorities simultaneously. Without such coordinated frameworks, interventions risk targeting only a fraction of AI's climate impact, while the larger share remains ungoverned.

AI does not inherently decarbonize or intensify emissions; its productivity gains reinforce the economic and institutional structures of the system in which it operates, making it the latest general-purpose technology operating within an energy system still structurally anchored in fossil fuels. This situates AI within a long-observed pattern: technological progress applied to incumbent energy systems often extends their persistence—including by delaying anticipated peak supply in some fossil sectors^{10,11}—contributing to prolonged coexistence with emerging alternatives rather than rapid displacement^{8,40}. A 2026 patent-based study reaches a convergent conclusion for AI specifically: absent policy steering, AI has so far entrenched existing energy systems rather than redirected them⁴¹.

Whether broader structural shifts in the underlying economy could redirect these dynamics at decarbonization-relevant timescales, given AI productivity gains are already operating on globally scaled fossil fuel infrastructure, remains an open question for future dynamic modeling work (see “Methods”). The directional finding from this analysis, however, is already unambiguous.

Whether AI repeats this historical pattern or sets a new trajectory depends on coordinated interventions across multiple pathways, closing the analytical-governance gap identified above. Until then, AI's potential to accelerate the energy transition must be reconciled with our finding that, under modeled conditions, AI-driven productivity gains enable more emissions than they avoid—reinforcing fossil fuel incumbency rather than displacing it.

Methods

Modeling framework

This study adopted a consequential (rather than attributional) framework: emissions outcomes are determined endogenously through economy-wide market responses, including cross-sector substitution, price adjustments, and demand feedbacks. AI adoption is represented as productivity shocks across competing energy pathways; this formulation captures system-level effects beyond isolated technical efficiency changes.

We employed a global computable general equilibrium (CGE) model, GTAP-E-Power⁴², calibrated to the GTAP-Power database version 11⁴³. GTAP-E-Power extends GTAP-E by incorporating 11 generation technologies (including grid transmission and distribution), representing substitution possibilities among different energy sources, while retaining GTAP-E's capital–energy substitution structure.

The model is calibrated to a 2017 base year, corresponding to the latest publicly available version of the GTAP-Power database⁴⁴. For simulation purposes, the database was aggregated to 20 economic sectors (including coal, oil, natural gas, and disaggregated fossil and renewables power generation) across five global regions (see SI Note S6). This approach resolves

Table 1 | Shock parameter ranges by pathway (detailed in SI Note S3)

Modeling pathway	Value chain category	Description	Shock type	Range
Enabled	Supply	Fossil upstream (commodity extraction)	Factor augmenting: Capital, labor, natural resources ^a	0–20%
Enabled	Supply	Fossil midstream (refining)	Factor- and energy- augmenting: Capital, labor, energy	2–15%
Enabled	Supply	Fossil downstream (power generation)	Factor augmenting: Capital, labor	2–10%
Avoided	Supply	Renewables power generation	Output-augmenting TFP ^b	2–30%
Fuel-neutral	Infrastructure	Grid transmission and distribution	Output-augmenting TFP	2–10%
Fuel-neutral	Demand	Maritime shipping and logistics	Output-augmenting TFP	5–15%
Fuel-neutral	Demand	Energy-intensive industry	Output-augmenting TFP	2–6%

Ranges span none–low–medium–high adoption tiers (see “Methods”). Pathway labels indicate the predominant emissions consequence of each shock category; see Box 1 for formal definitions.

^aNatural resource shocks apply to oil and gas extraction only; coal receives capital and labor shocks only.

^bRenewables magnitudes vary by technology (solar, wind, nuclear, hydro, other).

cross-sector feedback, rebound effects, and market adjustments that alternative frameworks are not designed to capture (see SI Note S2.2). We focused exclusively on CO₂ from fossil fuel combustion, calculated using fuel-specific coefficients applied to energy consumption volumes derived from IPCC Tier 1 methodology⁴⁵; non-CO₂ gases and industrial processes are excluded for analytical tractability.

The model was solved as a static, comparative-equilibrium system in which prices adjust to clear goods and factor markets under standard GTAP-E-Power closure. Production is represented using nested constant-elasticity-of-substitution (CES) functions; household demand uses the constant-differences-in-elasticities (CDE) utility function; and trade follows the Armington specification, differentiating goods by region of origin. Three factor categories are represented in the model: labor (skilled and unskilled), capital, and natural resources/land. The model assumes full employment, with capital reallocated across sectors to equalize rates of return within each region. Additional model structure and closure details are provided in SI Note S1.

Because share of non-fossil energy in global primary energy consumption has grown since 2017, we adjusted the primary energy baseline by increasing total factor productivity for low-carbon generation technologies until the electricity generation mix approximated 2024 levels⁴⁶. This ensures the smaller 2017 market share does not lead the model to underestimate the impacts of AI-driven renewables improvements. The robustness to this adjustment is evaluated in the reported sensitivity analysis. Other structural changes since 2017 are not resolved; this limitation warrants future research as updated GTAP databases become available. An extended discussion of model structure, trade-offs, validation, and extension opportunities is provided in SI Note S5. AI-driven productivity shocks are forward-looking, so direct replicative validation against historical outcomes is not feasible. We relied on structural validation via empirically grounded elasticities and sensitivity analysis over key parameters reported in Results. Prior replication and validation efforts for the GTAP-E framework are discussed in SI Note S1.

Pathway identification and empirical basis

Given meaningful model uncertainty, we adopted an exploratory scenario approach rather than deterministic forecasting^{47,48}. Shock parameterization requires identifying the mechanisms through which AI influences energy-sector emissions (documented across pathways below, and discussed in detail in SI Notes S3 and S8).

For fossil fuel supply-side pathways, AI improves fossil fuel productivity across the entire value chain: upstream (e.g., exploration and production), midstream (e.g., processing), and downstream (e.g., transport and generation). By improving productivity across capital, labor, and natural resource inputs simultaneously, AI reduces per-unit production costs, shifting sectoral supply curves outward. The reported scale is substantial: AI-powered analysis estimates that deployment of best-practice recovery techniques (including AI-driven methods) across existing fields could yield 470 billion to over 1 trillion additional barrels of technically recoverable oil

and alleviate peak oil supply expectations^{49,50}. Goldman Sachs analysts suggest these AI-driven gains could ultimately expand recoverable shale resources and delay peak U.S. shale supply¹¹—consistent with the supply expansion dynamics this study models. Goldman also estimates¹³ that AI-driven improvements could expand U.S. shale reserves by 8–20%, reduce drilling costs by approximately 30%, and reduce marginal production costs by roughly \$5 per barrel, potentially improving the commercial viability of marginal projects. Operator reporting corroborates this marginal-viability dynamic in two ways: by bringing previously uneconomic resources into production (the extensive margin) and by raising output and lowering costs at already-producing fields (the intensive margin). On the extensive margin, ExxonMobil’s vice president of exploration has stated that AI and new technologies enable the company to revisit oil prospects previously considered uneconomic⁵¹. On the intensive margin, a parallel dynamic is visible in established plays: break-even costs in the Permian Basin fell from over \$90 to approximately \$40 per barrel over the past decade, with AI cited as a contributor among multiple productivity-driving factors and expected to contribute further¹². Beyond upstream extraction, AI also reshapes midstream operations such as process optimization in refining^{52–54} and downstream operations such as combustion-efficiency gains in power generation^{14,55,56}, which can lower emissions per unit of output. Adoption is advancing; according to industry studies, 44% of upstream organizations already use AI in exploration, with 45% planning adoption within three years⁵⁷. Evidence used in the parameterization range (0–20% gains) primarily reflects realized or projected values; the upper bound is consistent with industry, operator, and financial-analysis estimates. Extensive calibration evidence and parameter mapping are provided in SI Notes S3.1.1, S8.

For renewables supply-side pathways, AI demonstrates nascent potential across diverse generation types. For brevity, we use the term “renewables” throughout this paper to refer to the aggregate of non-fossil electricity and heat generation technologies represented in the modeling framework, comprising nuclear, hydropower, wind, solar, biomass, geothermal, and other renewable technologies. For solar, research indicates 10–20% potential gains through intelligent tracking and maximum power point algorithms^{58,59}; for wind, the National Renewable Energy Laboratory (NREL) indicates a potential 6% increase in annual energy production⁶⁰, while Google’s DeepMind announced a 20% increase in value through improved forecasting⁶¹. For nuclear, the U.S. Department of Energy (DOE) reported pilots at nuclear plants that have improved operational efficiencies⁶² and provided cost reductions of up to 15% for predictive maintenance applications⁶³. However, our review found limited evidence of industry-wide realized gains beyond academic studies or pilots, suggesting that these estimates reflect technical potential rather than realized sectoral deployment. Furthermore, AI’s impact on renewables adoption rates is expected to be limited due to preexisting cost competitiveness⁶⁴; moreover, renewables adoption faces persistent institutional constraints¹⁴ (including interconnection backlogs, permitting delays, and curtailment) that AI optimization alone cannot resolve (see SI Note S3.1.2). We nonetheless

parameterized renewables generation gains at the upper bounds of documented technical potential (2–30%), with high scenarios assuming these deployment barriers are resolved. Additional calibration evidence and parameter mapping are provided in SI Note S3.1.2.

Infrastructure and demand-side pathways were both modeled as fuel-neutral, meaning shocks are applied symmetrically regardless of energy source (Box 1). For grid infrastructure, AI enhances grid performance primarily through increasing transmission capacity, improving operational reliability, and optimizing dispatch and forecasting¹⁴. The IEA estimates these applications could unlock 175 GW of transmission capacity without new construction, and that reducing renewable curtailment by one percentage point globally could avoid approximately 120 Mt CO₂ annually by 2035¹⁴. However, component-level technical potential consistently exceeds realized system-wide gains due to system-level infrastructure constraints, conservative operator adoption, and selective deployment on critical bottlenecks rather than fleet-wide gains (see SI Note S3.2). The IEA notes that grid AI deployment barriers are predominantly institutional rather than technical¹⁴. Industry adoption evidence is mixed: an industry survey finds that while 73% of utility executives report using AI, only 18% report that the technology has met expectations⁶⁵. Though modeled symmetrically, some applications (particularly curtailment reduction and renewables forecasting) could disproportionately benefit variable renewables generation; because the magnitude of any such effect is likewise not yet evidenced at realized, system-wide scale, we did not model it explicitly. We therefore conservatively parameterized grid efficiency gains at 2–10%, reflecting these observed gaps and the documented evidentiary basis. Additional calibration evidence and parameter mapping are provided in SI Note S3.2.

For demand-side pathways, AI improves end-use efficiency through applications such as real-time process control, predictive maintenance, and operational optimization. We modeled two applications based on IEA estimates of AI-driven efficiencies¹⁴: energy-intensive manufacturing (2–6% efficiency gains, constrained by thermodynamic limits in processes already near theoretical maxima) and maritime shipping (5–15%, where information-intensive logistics optimization faces fewer physical constraints). Together, these represent approximately 3.45 EJ—energy-intensive industry (~3 EJ) and maritime shipping (~0.45 EJ)—or about 25% of the 13.5 EJ of potential energy savings the IEA projects from widespread AI adoption across industry and transport by 2035^{4,14} (see SI Note S3.3). Remaining projected savings in additional sectors could increasingly couple with electrification pathways whose emissions consequences would be influenced by the supply-side dynamics already modeled (see Discussion). These efficiency shocks were applied symmetrically regardless of energy source, but the modeled end-use sectors themselves remain predominantly fossil-fueled, meaning realized savings primarily reduce fossil consumption. Additional calibration evidence and parameter mapping are provided in SI Note S3.3.

Shock parameterization

Attributing productivity gains specifically to digital technology, with AI as a subset, remains an open methodological challenge^{21–24}. Recognizing this, we calibrated shock magnitudes from sources that explicitly attribute gains to AI/ML, predominantly treat reported gains as ceiling estimates, and frame the results as conditional scenarios spanning plausible ranges rather than as predictions (see SI Note S3). The plausibility of these magnitudes is corroborated by historical evidence: comparable productivity gains were realized through earlier generations of digital technology in fossil fuel operations, establishing a precedent for near-term, rolling improvements of similar scale (see SI Notes S3.1.1 and S8).

The translation of these magnitudes into model shocks requires matching shock types to the physical and economic structure of each pathway. Fossil fuel pathways employed factor-augmenting shocks to the relevant inputs (capital, labor, energy, and, in upstream extraction, natural resources) because AI improvements in these sectors act on specific input

categories (such as capital and labor in drilling and refining, natural resources through improved reservoir characterization and recovery rates), and a composite Total Factor Productivity (TFP) shock would obscure the input substitution responses that drive supply expansion. Renewables pathways employed output-augmenting TFP shocks, reflecting improvements in output or service delivered per composite unit of input; AI in renewables generation primarily increases energy delivered from installed capacity (rather than accelerating new deployment; see SI Note S3.1.2) without differentially reducing one input factor over another through applications such as forecasting, tracking, and dispatch optimization. In fuel-neutral applications, AI improvements affect system-level performance (grid transmission and distribution efficiency, shipping and logistics optimization, energy-intensive industrial process efficiency) rather than fuel-specific extraction or generation.

For each pathway, we constructed corresponding adoption rates: none (no adoption), low, medium, and high (widespread adoption), calibrated to IEA 2035 estimates¹⁴. Throughout, parallel adoption refers to scenarios in which all pathways operate at the same adoption tier simultaneously; this symmetric configuration anchors headline calculations and is conservative, given documented incumbent advantages for fossil applications (see Discussion; SI Note S3). Because each pathway's productivity gains are calibrated to its own empirical evidence, parallel adoption specifies symmetric adoption intensity, not equal productivity magnitudes; the productivity gain ranges differ across pathways accordingly (see Table 1; SI Note S3).

Parameterization draws on three evidentiary tiers: authoritative institutional projections (principally IEA 2025¹⁴), peer-reviewed academic studies, and industry operator reports, in descending order of calibration confidence; consultancy estimates, financial-analyst notes (e.g., Goldman Sachs^{11,13}), and industry disclosures bound high adoption scenarios only and are not used for central values. Where sources converge across tiers, calibration confidence is highest; where evidence is sparse or conflicting, we adopted conservative parameterization and noted interpolations explicitly (see SI Notes S3.1–S3.4).

Two shock parameterization choices bias estimates toward understating the modeled fossil-renewables gap. First, renewables productivity gains were parameterized at the upper bounds of documented technical potential despite weaker realized-deployment evidence than for fossil applications (see SI Note S3.1.2). Second, evidentiary sourcing is asymmetric: fossil fuel calibrations derive primarily from realized, reported, or projected sectoral gains, while renewables calibrations derive predominantly from academic studies and single-site pilots (see SI Notes S3.1–S3.2). The differential shock specification accordingly does not favor fossil magnitudes: the resulting outcome therefore originates in the structure of the modeled economy, not in shock formulation.

We represented AI-driven productivity gains as shocks that shift supply curves downward by adjusting input-output coefficients: a 10% total factor productivity shock means inputs that previously generated 1 unit of output now generate 1.1 units; input-specific shocks work analogously (e.g., a 10% capital productivity shock in coal extraction means 1 unit of capital produces the equivalent of 1.1 units). Throughout, “productivity” denotes sectoral output gains, while “efficiency” is reserved for established technical terms. The full set of shock types and parameter ranges by pathway is summarized in Table 1.

Experimental design

We evaluated uniform shocks (0–100%) representing productivity gains applied simultaneously across fossil and renewables pathways to identify breakeven thresholds and the underlying asymmetry (Fig. 3); the upper end is illustrative rather than empirically realistic. To estimate net emissions under documented adoption patterns, we evaluated 64 empirically derived scenario combinations spanning 4 fossil × 4 renewables × 4 fuel-neutral levels (none, low, medium, high; Fig. 4; SI Note S3). To assess parametric sensitivity, we varied two key elasticities (capital–energy substitution and cross-technology substitutability in power generation) by ±50% from central values, a range consistent with observed differences between short- and

long-run elasticities in the empirical literature^{37,66} (SI Figs. S18–S21). We additionally tested sensitivity to the energy mix by comparing the 2024-adjusted baseline against the original 2017 GTAP-Power baseline (SI Fig. S22). We next evaluated carbon pricing scenarios (\$80–\$308/tCO₂), corresponding to Social Cost of Carbon estimates at 3% and 1.5% discount rates, respectively⁶⁷ (SI Fig. S14–S17). To assess macroeconomic implications, we compared modeled changes in CO₂ emissions against GDP, evaluating absolute and relative decoupling and the emissions intensity of GDP.

Limitations, excluded dynamics, and conservative assumptions

Figure 1 illustrates the analytical boundary of the model, distinguishing modeled, partially modeled, and unmodeled mechanisms. Below, we highlight the most relevant modeling constraints and excluded dynamics; extended discussion is provided in SI Notes S4 and S5.

The framework's structural constraints are standard in the CGE literature, including full employment, sectoral capital reallocation, and substitution elasticities (see SI Note S1); these shape output and emissions magnitudes but do not reverse the directional asymmetry, as the elasticity sensitivity analysis confirms. No model resolves all relevant dynamics; ours is designed to capture the mechanisms most consequential for the research question: cross-sector feedbacks, endogenous price responses, and economy-wide rebound and induction effects that alternative approaches are not designed to capture (see SI Note S2.2). These constraints describe limits of the framework as constructed and point to directions for future methodological extension.

A further constraint is sectoral and regional aggregation. The model aggregates heterogeneous actors across five regions and 20 sectors; it resolves aggregate trade and capital flows but not firm- or project-level heterogeneity (see SI Note S5). The model is calibrated to a 2017 base year with renewables generation adjusted to approximate 2024 levels; broader structural changes since 2017, including shifts in relative costs, policy environments, and market dynamics, are not resolved (see SI Note S6). While our framework resolves the global economy across 20 sectors with particular focus on electricity generation technologies, future research would benefit from more detailed parameterization of selected sectoral shocks. For instance, an explicit disaggregation of the key energy-intensive activities (such as steel, aluminum, cement, or fertilizers) which are all grouped in a single set of “energy-intensive industries” in the current study, combined with activity-specific parameterization of the AI-related shocks, would be an important extension of this analysis.

As a static framework, the model quantifies the direction and approximate scale of economic incentives under comparative-static conditions; it cannot capture temporal dynamics such as infrastructure lock-in that could amplify or attenuate modeled effects over time. The findings characterize structural incentives under current economic conditions, not a time-specific forecast; dynamic models in which adoption rates evolve over time (rather than being fixed as in our static framework) would be needed to resolve these dynamics as noted in the Discussion (see SI Notes S4 and S5).

On market dynamics, the model endogenously resolves price-mediated consumption and production responses, including output expansion induced by productivity gains, but cannot capture the full extent of documented rebound effects, new market creation, or asymmetries in how fossil and renewables productivity gains propagate through different market channels (see SI Notes S4 and S5).

Several system-level pathways are likewise unrepresented: distributed energy resources, integrated generation systems, active demand-side management (e.g., load shifting, demand response), and supply-chain dynamics are not represented in the framework. The directional impact on net emissions of including these is ambiguous; representing them at the granularity needed for finer-grained analysis represents a direction for future research. The model also abstracts from physical constraints, leaving operational and geological limits unresolved (grid physics, thermodynamic limits, and natural decline rates in fossil production); these bound the translation of theoretical productivity gains into realized emissions changes

in either direction (see SI Note S5). Some emergent low-carbon technologies are captured only within aggregated GTAP categories: geothermal, wave/tidal, and advanced storage fall under an “other renewables/other baseload” category, parameterized at a higher range (10–30%) than mature generation to reflect their low optimization baselines and lack of legacy constraints; with limited deployment history, this rests on structural inference rather than evidence and carries wide uncertainty (see SI Note S3.1.2). Others have no GTAP sector and cannot yet be represented (e.g., green hydrogen, small modular nuclear reactors, direct air capture); their exclusion understates avoided emissions potential, but as early-stage technologies at a negligible shares of supply, with many still at demonstration or prototype phase⁶⁸, their near-term effect on the modeled balance is minimal.

Several further dynamics below are excluded or only partially captured; addressing these represents a substantial agenda for future research (see SI Note S5). In most cases, available evidence suggests that inclusion would tend to increase rather than decrease estimated net emissions; accordingly, the reported results likely understate net effects.

Higher-order dynamics are only partially captured: the model represents policy channels through carbon pricing scenarios but does not resolve political-economy feedbacks (e.g., lobbying, subsidy persistence), infrastructure lock-in, or social, informational, and institutional dynamics (see SI Note S4 and Fig. 1).

Adoption asymmetries are also absent: productivity shocks are applied uniformly, implicitly assuming symmetric adoption rates despite documented deployment barriers for renewables and an evidentiary differential favoring fossil fuel applications (see Discussion; SI Notes S3–S5).

First-order datacenter emissions are excluded from equilibrium calculations for analytical tractability. These represent coupled dynamics: datacenter electricity demand adds load to the same fossil-intensive energy system whose productivity AI simultaneously enhances. Their inclusion would likely add to, rather than offset, the modeled net emissions effects (see SI Note S2.3) due to the mutually reinforcing effect identified (see Discussion).

Finally, the model tracks CO₂ only; non-CO₂ gases (including methane and other process emissions across fossil fuel operations, industrial processes, and downstream sectors) are excluded (see SI Note S5). These effects are bidirectional: AI applications such as methane leak detection can reduce fugitive emissions, while expanded fossil and industrial operations increase non-CO₂ emissions. Our directional finding is for CO₂ specifically; the net non-CO₂ contribution would need to be quantified separately to assess the full GHG impact, and unlike the other excluded dynamics, its effect on net emissions is not directionally determined.

Data availability

The datasets generated and analyzed during the current study are available in the accompanying GitHub repository (https://github.com/geldner/enabled_emissions_CGE), archived at Zenodo (<https://doi.org/10.5281/zenodo.21540204>), including shock parameters, experiment configurations, and output data. Empirical justifications for sector-specific shock parameterizations are documented in SI Notes S3 and S8. The GTAP-Power 5 × 20 aggregation, included in the repository, can be run using either a free installation of the GEMPACK modeling software or via the RUNGTAP software, publicly available through the GTAP website (<https://www.gtap.agecon.purdue.edu/>). The underlying GTAP-Power database is distributed by the Global Trade Analysis Project under its own access terms.

Code availability

All model code, simulation scripts, experiment configurations, and output data are available in the accompanying GitHub repository https://www.github.com/geldner/enabled_emissions_CGE, archived at Zenodo (<https://doi.org/10.5281/zenodo.21540204>).

Received: 10 November 2025; Accepted: 21 July 2026;

Published online: 04 August 2026

References

1. IPCC. *Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change* (IPCC, 2023).
2. Energy Transitions Commission. *Fossil Fuels in Transition: Committing to the Phase-down of All Fossil Fuels* (Energy Transitions Commission, 2023).
3. Ritchie, H. & Rosado, P. *Energy mix. Our World in Data* <https://ourworldindata.org/energy-mix> (2020).
4. IEA. *World Energy Outlook 2025* (IEA, 2025).
5. IEA. *Global Energy Review 2025* (IEA, 2025).
6. IRENA. *Renewable Power Generation Costs in 2024* (IRENA, 2025).
7. York, R. & Bell, S. E. Energy transitions or additions? Why a transition from fossil fuels requires more than the growth of renewable energy. *Energy Res. Soc. Sci.* **51**, 40–43 (2019).
8. Unruh, G. C. Understanding carbon lock-in. *Energy Policy* **28**, 817–830 (2000).
9. IEA. *The Implications of Oil and Gas Field Decline Rates* (IEA, 2025).
10. IEA. *Resources to Reserves* (IEA, 2005).
11. Cunningham, S. AI holds potential to extend shale era. Argus Media <https://www.argusmedia.com/en/news-and-insights/latest-market-news/2610913-ai-holds-potential-to-extend-shale-era> (2024).
12. Salzman, A. AI could lift the oil industry's profits. It's happening in the Permian Basin. Barron's <https://www.barrons.com/articles/ai-oil-industry-profit-permian-basin-8adcacdd> (2024).
13. Reuters. AI likely to weigh on oil prices over the next decade, Goldman says. <https://www.reuters.com/technology/artificial-intelligence/ai-likely-weigh-oil-prices-over-next-decade-goldman-says-2024-09-03/> (2024).
14. IEA. *Energy and AI* (IEA, 2025).
15. Ashraf, W. M., Dua, V. & Debnath, R. Domain consistent industrial decarbonisation of global coal power plants. *Commun. Sustain.* **1**, 10 (2026).
16. Lange, S., Pohl, J. & Santarius, T. Digitalization and energy consumption. Does ICT reduce energy demand? *Ecol. Econ.* **176**, 106760 (2020).
17. Lange, S. et al. The induction effect: why the rebound effect is only half the story of technology's failure to achieve sustainability. *Front. Sustain.* **4**, 1178089 (2023).
18. Xiao, T., Nerini, F. F., Matthews, H. D., Tavoni, M. & You, F. Environmental impact and net-zero pathways for sustainable artificial intelligence servers in the USA. *Nat. Sustain.* **8**, 1541–1553 (2025).
19. Dodge, J. et al. Measuring the carbon intensity of AI in cloud instances. In *Proc. 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22)*, 1877–1894 (Association for Computing Machinery, 2022).
20. de Vries, A. The growing energy footprint of artificial intelligence. *Joule* **7**, 2191–2194 (2023).
21. Kaack, L. H. et al. Aligning artificial intelligence with climate change mitigation. *Nat. Clim. Change* **12**, 518–527 (2022).
22. Luers, A. et al. Will AI accelerate or delay the race to net-zero emissions? *Nature* **628**, 718–720 (2024).
23. Bieser, J. C. T., Coroamă, V. C., Bergmark, P. & Stürmer, M. The greenhouse gas (GHG) reduction potential of ICT: a critical review of telecommunication companies' GHG enablement assessments. *J. Ind. Ecol.* **28**, 1132–1146 (2024).
24. Bieser, J. AI and climate protection: research gaps and needs to align machine learning with greenhouse gas reductions. in *Proc. 2024 10th International Conference on ICT for Sustainability (ICT4S)* 1–9 (IEEE, 2024).
25. Lanz, B. & Rausch, S. General equilibrium, electricity generation technologies and the cost of carbon abatement: a structural sensitivity analysis. *Energy Econ.* **33**, 1035–1047 (2011).
26. Russell, S. *Estimating and Reporting the Comparative Emissions Impacts of Products* (GHG Protocol, 2018).
27. European Green Digital Coalition. *Net Carbon Impact Assessment Methodology for ICT Solutions* (European Green Digital Coalition, 2024).
28. Ackerman, F., DeCanio, S. J., Howarth, R. B. & Sheeran, K. Limitations of integrated assessment models of climate change. *Clim. Change* **95**, 297–315 (2009).
29. Luers, A. Net zero needs AI — five actions to realize its promise. *Nature* **644**, 871–873 (2025).
30. European Parliament & Council. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Off. J. Eur. Union L, 2024/1689 <http://data.europa.eu/eli/reg/2024/1689/oj> (2024).
31. US Congress. Artificial Intelligence Environmental Impacts Act of 2024, S. 3732, 118th Congress (introduced 2024; not enacted). <https://www.congress.gov/bills/118th-congress/senate-bill/3732>.
32. Pellegrini, L. & Arsel, M. The supply side of climate policies: keeping unburnable fossil fuels in the ground. *Glob. Environ. Polit.* **22**, 1–14 (2022).
33. Newell, P. & Daley, F. Supply-side climate policy: a new frontier in climate governance. *WIREs Clim. Change* **15**, e909 (2024).
34. Erickson, P., Lazarus, M. & Piggot, G. Limiting fossil fuel production as the next big step in climate policy. *Nat. Clim. Change* **8**, 1037–1043 (2018).
35. Geels, F. W. Regime resistance against low-carbon transitions: introducing politics and power into the multi-level perspective. *Theory Cult. Soc.* **31**, 21–40 (2014).
36. Li, M., Trencher, G. & Asuka, J. The clean energy claims of BP, Chevron, ExxonMobil and Shell: a mismatch between discourse, actions and investments. *PLOS ONE* **17**, e0263596 (2022).
37. Kulmer, V. & Seebauer, S. How robust are estimates of the rebound effect of energy efficiency improvements? A sensitivity analysis of consumer heterogeneity and elasticities. *Energy Policy* **132**, 1–14 (2019).
38. Jarvis, A. & King, C. W. Economy-wide rebound and the returns on investment in energy efficiency. *Energy Effic.* **17**, 60 (2024).
39. Ritchie, H. & Rosado, P. Which countries have put a price on carbon? Our World in Data <https://ourworldindata.org/carbon-pricing> (2022).
40. Smil, V. *Energy and Civilization: A History* (MIT Press, 2018).
41. Herman, K. S., Amortegui, T. G. & Griffiths, S. Are AI and environmental technology innovations converging? *Technol. Soc.* **85**, 103222 (2026).
42. Peters, J. C. GTAP-E-Power: an electricity-detailed economy-wide model. *J. Glob. Econ. Anal.* **1**, 156–187 (2016).
43. Chepeliev, M. GTAP-Power data base: version 11. *J. Glob. Econ. Anal.* **8**, 2 (2023).
44. Chepeliev, M. *A Revised CO₂ Emissions Database for GTAP*. GTAP Research Memorandum (Purdue University, 2022).
45. Chepeliev, M. Chapter 13A: CO₂ emissions from fossil fuels combustion. Center for Global Trade Analysis http://www.gtap.agecon.purdue.edu/resources/res_display.asp?RecordID=7130 (2024).
46. Ritchie, H., Rosado, P. & Roser, M. Energy production and consumption. Our World in Data <https://ourworldindata.org/energy-production-consumption> (2020).
47. Kwakkel, J. H. & Pruyt, E. Exploratory modeling and analysis, an approach for model-based foresight under deep uncertainty. *Technol. Forecast. Soc. Change* **80**, 419–431 (2013).
48. Johnson, D. R. & Geldner, N. B. Contemporary decision methods for agricultural, environmental, and resource management and policy. *Annu. Rev. Resour. Econ.* **11**, 19–41 (2019).
49. Latham, A., Marnell, O. & Dixon, J. Every last drop: using AI-powered analysis to find oil-field upside potential. Wood Mackenzie <https://www.woodmac.com/horizons/ai-powered-analysis-oil-field-potential/> (2025).

50. Flowers, S. How AI can unlock an extra trillion barrels of oil. Wood Mackenzie <https://www.woodmac.com/blogs/the-edge/how-ai-can-unlock-an-extra-trillion-barrels-of-oil/> (2025).
51. Williams, C. Exxon using AI for faster analysis of Guyana's oil fields, VP of exploration says. *Reuters* <https://www.reuters.com/business/energy/exxon-using-ai-faster-analysis-guyanas-oil-fields-vp-exploration-says-2026-05-05/> (2026).
52. ExxonMobil. Smart technologies for a smart refinery. <https://corporate.exxonmobil.com/who-we-are/technology-and-collaborations/smart-technologies-intelligent-operations> (2021).
53. Ba, L., Tangour, F., El Abbassi, I. & Absi, R. Analysis of digital twin applications in energy efficiency: a systematic review. *Sustainability* **17**, 3560 (2025).
54. McKinsey. Technology transformation for oil and gas companies. <https://www.mckinsey.com/capabilities/operations/our-insights/harnessing-volatility-technology-transformation-in-oil-and-gas> (2022).
55. Lee, K.-J., Choi, S.-W. & Lee, E.-B. Artificial intelligence-driven approach to optimizing boiler power generation efficiency: the advanced boiler combustion control model. *Energies* **18**, 820 (2025).
56. Boswell, B., Buckley, S., Elliot, B., Melero, M. & Smith, M. An AI power play: fueling the next wave of innovation in the energy sector. McKinsey <https://www.mckinsey.com/capabilities/tech-and-ai/how-we-help-clients/an-ai-power-play-fueling-the-next-wave-of-innovation-in-the-energy-sector> (2022).
57. IBM. *Oil and Gas in the AI Era* (IBM Institute for Business Value, 2025).
58. Mamodiya, U. et al. Numerical modeling and neural network optimization for advanced solar panel efficiency. *Sci. Rep.* **15**, 23492 (2025).
59. Azad, A. M. A. S. et al. Harnessing the sun: framework for development and performance evaluation of AI-driven solar tracker for optimal energy harvesting. *Energy Convers. Manag.* **X** **26**, 100990 (2025).
60. Vijayakumar, G. & King, R. *Enabling Innovation in Wind Turbine Design Using Artificial Intelligence* (NREL, 2022).
61. Elkin, C. & Witherspoon, S. Machine learning can boost the value of wind energy. Google DeepMind Blog <https://deepmind.google/blog/machine-learning-can-boost-the-value-of-wind-energy/> (2019).
62. US Department of Energy. New AI tools could save Constellation reactor fleet millions. <https://www.energy.gov/ne/articles/new-ai-tools-could-save-constellation-reactor-fleet-millions> (2024).
63. Walker, C. et al. *Explainable Artificial Intelligence Technology for Predictive Maintenance* (US Department of Energy, 2023).
64. Stern, N. et al. Green and intelligent: the role of AI in the climate transition. *Npj Clim. Action* **4**, 56 (2025).
65. Juchno, M. AI can help utilities predict grid outages. EY Insights https://www.ey.com/en_us/insights/power-utilities/ai-can-help-utilities-predict-grid-outages (2025).
66. Gallaway, M. P., McDaniel, C. A. & Rivera, S. A. Short-run and long-run industry-level estimates of U.S. Armington elasticities. *North Am. J. Econ. Financ.* **14**, 49–68 (2003).
67. Rennert, K. et al. Comprehensive evidence implies a higher social cost of CO₂. *Nature* **610**, 687–692 (2022).
68. IEA. *Net Zero by 2050: A Roadmap for the Global Energy Sector* (IEA, 2021).
69. Ekchajzer, D., Bornes, L., Combaz, J., Letondal, C. & Vingerhoeds, R. Decision-making under environmental complexity: the need for moving from avoided impacts of ICT solutions to systems thinking approaches. In *Proc. 2024 10th International Conference on ICT for Sustainability (ICT4S)* 29–40 (IEEE, 2024).

Acknowledgements

We thank Rachel Woods-Robinson for contributions to the design of Fig. 1 and for editorial support, and Camille Hu for the design of Fig. 2. We further

thank the numerous others whose insights improved this work: Amory Lovins, Michael Lazarus, Vlad C. Coroamă, Bill Wehl, Jo Lindsay Walton, Parker Caswell, Kate Schnipper, Kenneth P. Pucker, Tamara Kneese, Kathy Mulvey, Delta Merner, Lynn Kaack, Gauthier Roussilhe, John Kos-tyack, Chris Adams, Sean Higgins, Bruce Fryer, Mark Butcher, Eugene Kirpichov, Sasha Luccioni, Auden Schendler, Lou Welgryn, Theo Alves, Molly Webb, Anna Lerner Nesbitt, Drew Wilkinson, Brian Bills, Elliott Fox, Lucas Pratesi, and Nathaniel Burola, in addition to several others who declined to be named due to professional affiliations. Any remaining errors are our own. This work received no dedicated grant funding and was conducted on a volunteer basis by the authors. W. Alpine and H. Alpine gratefully acknowledge organizational support (~\$3,000 USD) from Lush Cosmetics, which had no role in the study.

Author contributions

W. Alpine proposed the analytical framing of enabled versus avoided emissions and higher-order feedbacks, performed the literature review, developed the scenarios, and was responsible for the overall manuscript structure, framing, and writing. Dr. N. Geldner performed the simulations and analysis and contributed substantially to methodological framing, the interpretation of results, and substantial writing. H. Alpine contributed to the conceptual framing of the study and to substantive revision of the manuscript. Dr. M. G. Chepeliev provided the model code, GTAP-Power 5 × 20 database aggregation, feedback on the analytical methods and framing, and contributions to the paper's structure and writing.

Competing interests

W. Alpine and H. Alpine are affiliated with the Enabled Emissions Campaign, a nonprofit research and advocacy project focused on AI's climate impacts and technology governance, and declare this affiliation as a non-financial competing interest. The authors declare no financial competing interests, and no external party influenced the study design, analysis, interpretation, or decision to publish.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s44168-026-00411-0>.

Correspondence and requests for materials should be addressed to Will Alpine.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2026