Commodity Special Feature: Market Developments and the Impact of AI on Energy Demand

Primary commodity prices increased 1.9 percent between August 2024 and March 2025, with the rise driven by natural gas, precious metals, and beverage prices. In oil markets, prices fell amid concerns that a trade war could dampen global demand, adding to downward pressure from robust oil production growth outside OPEC+ (Organization of the Petroleum Exporting Countries plus selected nonmember countries, including Russia) and the unwinding of OPEC+ supply cuts. With the notable exception of gold prices, which continued to soar owing to geopolitical uncertainty, and prices of some staples like wheat, most commodity prices have dropped since the announcement of additional tariffs by the US administration on April 2. This Special Feature also analyzes the impact of artificial intelligence (AI) on energy demand.

Commodity Market Developments

Oil prices declined 9.7 percent between August 2024 and March 2025 as trade war fears, strong non-OPEC+ supply growth, and the unwinding of OPEC+ cuts more than offset lingering supply risks. Oil prices then plummeted in early April amid escalating trade tensions, adding to an already-bearish outlook. This latest catalyst compounded weak fundamentals, with supply growth expected to likely outpace tepid global demand growth through 2025 and 2026. Demand concerns were exacerbated by sluggish Chinese demand, partly dented by the rising penetration of electric vehicles (EVs).

In this context, OPEC+ policy will be pivotal: Facing pressure to roll back its deep and sustained cuts, OPEC+ has decided to start gradually unwinding them despite a broader environment of falling prices. The harshest sanctions on Russia to date (imposed on January 10, 2025) have not materially disrupted oil flows. Russian oil, exported primarily to China and India, has traded at a $5–$15 discount to Brent. Futures markets indicate that oil prices will average $66.9 per barrel in 2025, a 15.5 percent decline, before falling to $62.4 in 2026 (Figure 1.SF.1, panel 2). Risks to this outlook are balanced. Upside price risks from potential disruptions in oil supply from countries subject to sanctions or a de-escalation of trade barriers are offset by the possibility of a further escalation in the trade war and additional increases in OPEC+'s production schedule.

The contributors to this Special Feature are Christian Bogmans, Patricia Gomez-Gonzalez, Giovanni Melina (team co-lead), Jorge Miranda-Pinto, Andrea Paloschi, Andrea Pescatori (team lead), and Sneha Thube, with research assistance from Ganchimeg Ganpurev, Maximiliano Jerez Osses, and Joseph Moussa. This Special Feature is based on Bogmans and others (2025).

Figure 1.SF.1. Commodity Market Developments

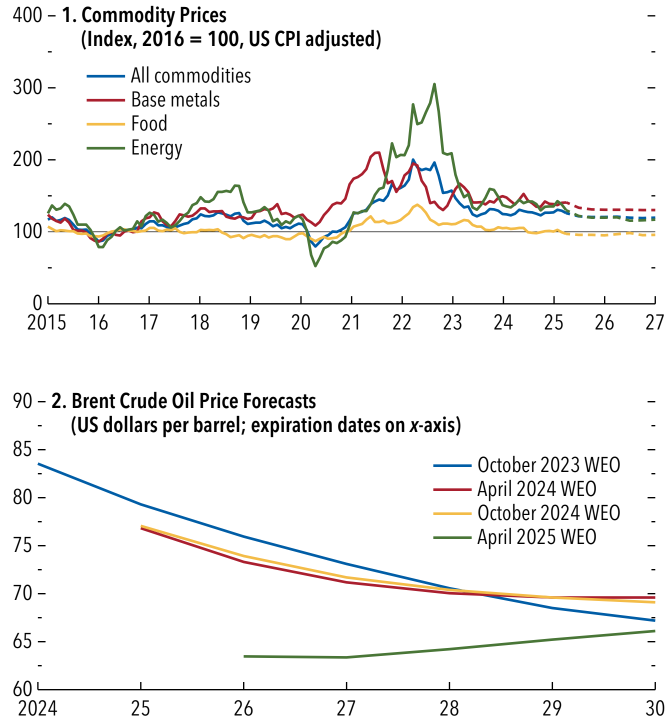

Sources: Bloomberg, L.P.; Haver Analytics; IMF, Primary Commodity Price System; International Energy Agency; and IMF staff calculations.

Note: In panel 1, latest actual CPI value is applied to forecasts, represented by the dashed portions of the graph lines. CPI = consumer price index; WEO = *World Economic Outlook*.

Natural gas prices reversed course in the first week of April, beginning to decline alongside oil prices after a six-month period of gains. Title Transfer Facility (TTF) trading hub prices in Europe rose 7.7 percent between August 2024 and March 2025 to $13.1 a million British thermal units (MMBtu). This was above the historical average but well below the 2022 peak. Among other factors, a cold snap and various supply disruptions, including a halt of Russian gas to Europe through Ukraine at the beginning of January 2025, explained the upward trend. Similarly, harsh weather and a surge in demand for gas exports led to a doubling in Henry Hub prices.

Weak demand from China, in contrast, kept Asian liquefied natural gas prices almost constant over the same period. Following the April 2 tariff announcement, gas prices reversed course, with concerns about future energy demand pushing gas prices down across the board. As of April 4, futures markets suggested that TTF prices will average $12.5 a MMBtu in 2025, steadily decreasing to $7.8 a MMBtu in 2030. Henry Hub prices are expected to decline from $4.0 a MMBtu in 2025 to $3.3 a MMBtu in 2030. Risks to this outlook are balanced.

Metals prices rose amid safe-haven demand and supply disruptions until the end of March, but things changed abruptly on April 2. The IMF's metals price index increased by 11.2 percent between August 2024 and March 2025 (Figure 1.SF.1, panel 1), with the rise driven mainly by gold, aluminum, and copper prices. Among base metals, aluminum (12.7 percent) and copper prices (8.4 percent) increased the most because of supply concerns. Both metals also faced demand pressures from front-loading ahead of tariffs. Like those for energy, industrial metals prices dropped abruptly in the first week of April as trade tensions escalated. Futures markets now predict a downturn in prices for base metals, with price declines of 5.7, 4.5 and 14.3 percent for aluminum, copper, and iron ore, respectively, by the end of 2026. This stands in contrast to what has taken place regarding prices for precious metals: Gold prices have repeatedly set new records amid policy and geopolitical uncertainty, recently surpassing their historical high at $3,000 per ounce.

Agricultural commodity prices increased as a result of adverse weather. Between August 2024 and March 2025, the IMF's food and beverages price index increased by 3.6 percent, with the rise driven by higher beverage prices. Cereal prices increased modestly, by 0.6 percent, as concerns over crop conditions for wheat and corn subsided. Coffee prices jumped 33.8 percent, with the IMF coffee index reaching historic highs in February because of weather-related supply concerns in Brazil. Meanwhile, rice prices fell 26.0 percent as crop conditions improved in India and other parts of Asia. New trade barriers imposed in April had heterogeneous effects on agricultural prices. The price of income-elastic (coffee) and trade-sensitive (soybeans) crops have declined sharply, whereas prices for staples like corn and wheat are so far less affected. Upside risks stem from trade disruptions and adverse weather; larger-than-expected harvests, trade war intensification, and broader uncertainty are the main downside risks.

Power Hungry: How AI Will Drive Energy Demand

The rapid development and adoption of generative AI models, including large language models, require building more data centers that consume vast amounts of electricity. Large language models' costs have two main components: a large fixed cost for training the models and variable costs for operating and responding to user prompts.[1] Because substantial computational resources are required during both stages, electricity consumption represents a critical input for companies delivering AI services. In northern Virginia, which features the largest concentration of data centers in the world, the square footage of server-filled warehouses is now roughly equivalent to the floor space of eight Empire State Buildings (Cushman & Wakefield 2024).

Using a multicountry computable general equilibrium (CGE) model, IMF-ENV (Chateau and others 2025), this Special Feature seeks to answer the following questions: (1) How fast have sectors involved in the development and delivery of AI-related services grown in recent years, and what has happened to their electricity consumption? (2) How does the projected electricity demand from AI by 2030 compare with other drivers of demand, such as EVs? (3) What is the impact on energy prices and the mix of electricity sources under alternative policy scenarios? (4) What will be the impact of data centers' growth on carbon emissions?

The Growing Macroeconomic Relevance of AI-Producing Sectors

In the US, AI-producing sectors' value added quadrupled from $278 billion (in constant 2017 dollars) to $1.13 trillion between 2010 and 2023, a rate much faster than those for private nonfarm and manufacturing value added. As a result, these sectors' share in total US GDP increased from 2.4 percent in 2013 to 3.5 percent in 2023, with the data-processing sector nearly doubling its share in the same period. Meanwhile, the share of manufacturing declined by 1.5 percentage points (Figure 1.SF.2, panel 1). This fast growth of AI-producing sectors was driven by remarkable gains in labor productivity, with value added per employee in the data-processing sector

[1]Large fixed costs create economies of scale that concentrate AI development among a few large players (Korinek and Vipra 2024), although this pool has expanded recently as more variation in the cost structure of large language models has emerged.

Figure 1.SF.2. The Growing Macroeconomic Relevance of AI-Producing Sectors

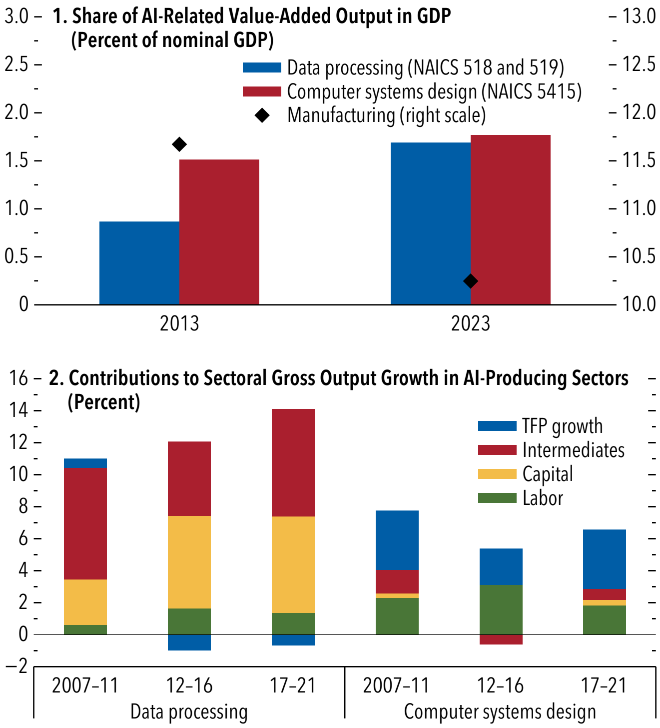

Sources: Haver Analytics; BEA-BLS Integrated Industry-Level Production Accounts (KLEMS); and IMF staff calculations.
Note: AI = artificial intelligence; NAICS = North American Industry Classification System; TFP = total factor productivity.

growing about four times faster than that in the whole economy over the past 10 years (see Online Annex Figure 1.1.2, panel 1 in Online Annex 1.1).[2] This productivity growth was largely the result of elevated investment in physical capital and the complementarity of intermediate inputs, contrary to what was the case in computer systems design, in which labor and total factor productivity (TFP) contributed significantly to output growth (Figure 1.SF.2, panel 2). Hence, the high output per employee in data centers, compared with that in other sectors, is the result of rapid capital accumulation, which has increased energy consumption as an intermediate input.

AI's Demand for Electricity

Electricity costs make up 13–15 percent of total costs for data center companies, whereas they account for only 0.8–1.5 percent for semiconductor firms and AI service companies. However, the latter have almost doubled the share of electricity costs in their total costs in less than five years (see Online Annex Figure 1.1.3 in Online Annex 1.1). As these companies integrate vertically by building, operating, and leasing their own data centers, that share will likely continue to grow.

The broader implications for global electricity consumption are substantial. Worldwide electricity consumption from data centers and AI is estimated to have reached 400–500 terawatt-hours (TWh) in 2023, more than double the level in 2015 (OPEC 2024). For the United States, where growth is the fastest, electricity demand from data centers is expected to increase from 178 TWh in 2024 to 606 TWh in 2030 under a medium-demand scenario (McKinsey & Company 2024a). By 2030, AI-driven global electricity consumption could hit 1,500 TWh, conceivably making its level comparable to that of India's current total electricity consumption, the third highest in the world. This projected electricity demand from AI by 2030 is about 1.5 times higher than expected demand from EVs, another emerging source of electricity demand (Figure 1.SF.3).

Recent developments in the AI industry have increased uncertainty about its future compute and energy demands. Companies such as DeepSeek are

Figure 1.SF.3. AI's Demand for Electricity
(Thousands of terawatt-hours; electricity demand for data centers compared with that in top electricity-consuming countries in 2023)

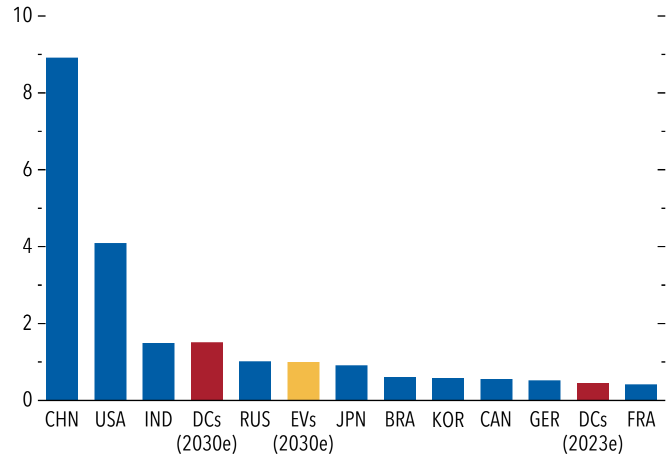

Sources: International Energy Agency (IEA); Organization of the Petroleum Exporting Countries (OPEC); and IMF staff calculations.
Note: Estimates for data centers (DCs) and electric vehicles (EVs) are for the world and come from OPEC and the IEA, respectively. Data labels in the figure use International Organization for Standardization (ISO) country codes. e = estimate.

[2]All online annexes are available at www.imf.org/en/Publications/WEO.

achieving breakthroughs in algorithmic efficiency that may lower the computational costs of AI models faster than previously anticipated. However, these efficiency gains may be counterbalanced by greater use of compute by companies pursuing better-performing models (Hoffmann and others 2022). Adding to this complexity is the recent emergence of reasoning models—which require more compute in their deployment—and possibly greater AI use driven by lower costs and availability of open-source models.

The Effects of Increased Demand for Electricity

In the IMF-ENV model, the impact of AI is captured by an increase in information technology (IT) sectors' TFP in China, the United States, and Europe to match the expected increase in data center power demand between 2025 and 2030 (see Online Annex Table 1.1.1. in Online Annex 1.1). This growth is projected at constant annual rates of 22, 13, and 10 percent, respectively (JP Morgan 2024; McKinsey & Company 2024a, 2024b).

Three scenarios are simulated here: (1) a *baseline* scenario, which excludes the AI-related TFP shock but reflects energy and emissions projections consistent with policies introduced through 2024; (2) an *AI scenario under current energy policies*, which models the AI-related TFP shock, assuming that the composition of electricity generation remains identical to that in the *baseline* scenario; and (3) an *AI scenario under alternative energy policies*, under which the share of renewables in total electricity generation is aligned with regions' long-term strategies using feed-in tariffs for renewables, though in practice policy choices will be guided by countries' preferences.[3] Results for both AI scenarios are reported as deviations from the *baseline* scenario, unless stated otherwise.

The AI shock increases electricity consumption by the IT sector, and power producers are expected to expand generation. The composition of electricity generation by technologies varies across countries and is based on their relative production costs and current policies. By 2030, in the *AI scenario under current energy policies*, total electricity supply increases by 8 percent in the United States (525 TWh), 3 percent in Europe (145 TWh), and 2 percent in China (237 TWh) relative to the baseline scenario. In the *AI scenario under alternative energy policies*, the increase in total electricity supply is kept the same, but its composition shifts in favor of renewables. In China, the United States, and Europe, generation from solar and wind sources offsets about 166 TWh, 58 TWh, and 35 TWh of generation, respectively, from other sources, including largely coal power in China and natural gas in the US (Figure 1.SF.4, panel 1).

In both scenarios, the rising marginal costs of electricity supply mean that the increase in generation is less than proportional to economy-wide demand growth, which drives electricity prices up. At the same time, strong commitment of major AI players to resolving medium-term power supply rigidities[4] could lead to a smaller increase in electricity prices. In this case, the surge would be 0.9 percent in the United States, 0.45 percent in Europe, and 0.35 percent in China under current energy policies (Figure 1.SF.4, panel 2). However, material pressure on prices would be added if the renewables scale-up slows from recent trends and if further investments are not made in transmission and distribution capacities (relative to those in the *baseline*). The price increase in the *AI scenario under current energy policies* could escalate up to 5.3 percent in China, 8.6 percent in the United States, and 3.6 percent in Europe by 2030 (Figure 1.SF.4, panel 2), adding to price pressures coming from many other sources.[5]

In addition, without further investments in transmission and distribution, support for the expansion of the AI sector would require redirecting electricity from other economic activities. Such a shift would pose significant challenges, especially for energy-intensive manufacturing sectors. In the United States, for example, annual growth in the value added of these sectors would fall by an average of 0.3 percentage point compared with that in the *baseline* scenario, reducing

[3]AI expansion relies on electricity growth, so countries' energy policies should focus on supply. Different supply-side policies affect prices, GDP, and revenue (Chateau, Jaumotte, and Schwerhoff 2024). Feed-in tariffs for solar photovoltaic (PV) and wind are simulated owing to their historical inclusion in policy packages and because these renewables are cost competitive with fossil fuels in these regions (IRENA 2024).

[4]Public investments are being made in the United States for upgrading transmission and distribution infrastructure to meet rising electricity demand. Innovative solutions like power coupling (Engel, Posner, and Varadarajan 2025) and small modular nuclear reactors could offer flexibility, making constraints less restrictive than expected. Most new nuclear capacity in the United States is expected online no earlier than the early 2030s.

[5]Chandramowli and others (2024) estimate a 19 percent rise in US wholesale electricity prices from 2025 to 2028 because of increased demand driven not only by data centers, but also by electrification of buildings and transportation, battery and fuel cell manufacturing, AI, and cryptocurrency mining.

Figure 1.SF.4. The Effects of Increased Demand for Electricity

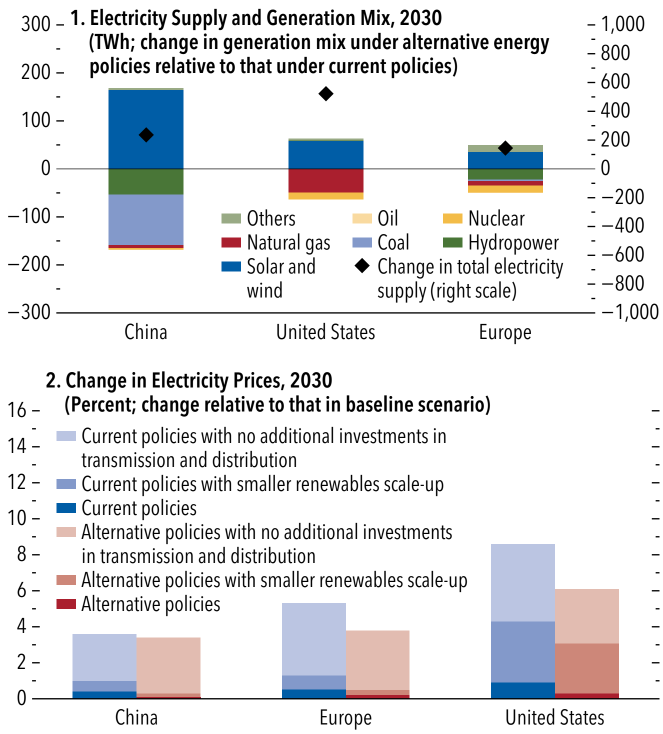

Sources: IMF, IMF-ENV model; and IMF staff calculations.

Note: In panel 1, the left axis shows the change in generation mix under alternative energy policies relative to current policies in terawatt-hours (TWh). Feed-in tariffs increase generation from solar and wind sources. The right axis shows the total increase in electricity supply relative to the baseline scenario in TWh, which is identical under both current energy policies and alternative energy policies.

Figure 1.SF.5. Emission Impacts of Expansion in IT Sector

(MtCO$_2$e; cumulative greenhouse gas emissions; Percent change relative to that in baseline, right scale)

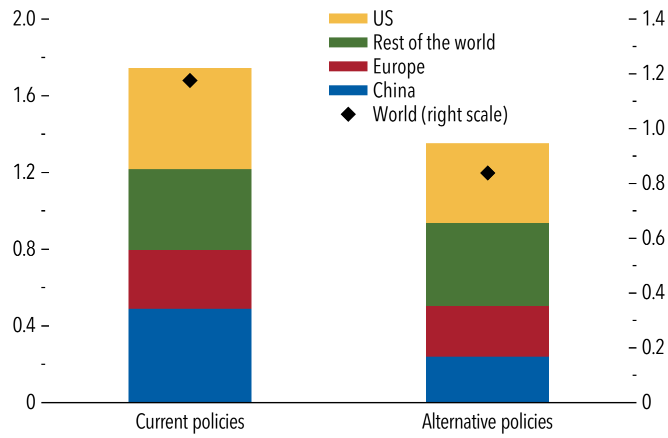

Sources: IMF, IMF-ENV model; and IMF staff calculations.

Note: The left axis shows the total greenhouse gas emissions increase in metric tons of carbon dioxide equivalent (MtCO$_2$e) between 2025 and 2030 resulting from information technology (IT) sector expansion in selected regions. The right axis shows the total increase in global emissions in 2030 relative to the baseline emissions as a result of this expansion.

annual GDP growth by 0.1 percentage point. The electricity price increase is more muted in the *AI scenario under alternative energy policies* owing to feed-in tariffs on solar and wind. The tariffs reduce the generation price of these technologies, which have relatively low production costs and a higher share in total electricity generation compared with those in the *AI scenario under current energy policies*.

In both AI scenarios, global and regional greenhouse gas (GHG) emissions increase because of the increased energy demand resulting from the expanded IT sector and its spillovers to the economy. In the *AI scenario under current energy policies*, the 2030 increase is 5.5, 3.7, and 1.2 percent in the US, Europe, and China, respectively, with a global average increase of 1.2 percent (Figure 1.SF.5). In cumulative terms, this translates into a global GHG emissions increase of 1.7 gigatons (Gt) between 2025 and 2030, which is similar to Italy's energy-related GHG emissions over a five-year period. Notably, in the *AI scenario under alternative energy policies*, even a modest decarbonization of the power sector limits the total cumulative global GHG emissions increase to 1.3 Gt by 2030, which is 24 percent less than in the *AI scenario under current energy policies*.[6]

In the *AI scenario under current energy policies*, the AI shock raises the average annual growth rate of global GDP by 0.5 percentage point between 2025 and 2030, in line with previous IMF estimates ranging between 0.1 percentage point and 0.8 percentage point (April 2024 *World Economic Outlook*). The impact is greater in countries where the projected growth rate of the IT sector and its relative importance in the economy are higher. In the *AI scenario under alternative energy policies*, these gains are slightly reduced because of the feed-in tariff polices. The total fiscal costs of these tariffs range from 0.3 percent to 0.6 percent of GDP across countries and are financed through increased lump-sum taxes, which slightly reduce household consumption. However, the growth benefits from AI expansion far outweigh these costs, resulting in similar average annual GDP growth across both scenarios.

In summary, although the AI-induced expansion of the IT sector is expected to raise global GDP, the development also comes at the cost of higher carbon emissions. Drawing on a median social cost of carbon

[6]This estimate is conservative compared with that of Stern and Romani (2025), who project that AI's energy demand could contribute between 0.4 and 1.6 Gt of carbon dioxide equivalent annually by 2035.

estimate of $39 per ton—based on 147 published studies with more than 1,800 estimates (Moore and others 2024)—the additional social cost of 1.3 to 1.7 Gt of carbon-dioxide-equivalent emissions is about $50.7 billion to $66.3 billion, or 1.3 percent to 1.7 percent of the AI-driven increase in real world GDP between 2025 and 2030.

Conclusions and Policy Implications

As AI technologies continue to evolve and proliferate, demand for computational power and electricity is poised for a significant surge. Despite challenges related to higher electricity prices and GHG emissions, the gains to global GDP from AI are likely to outweigh the costs of the additional emissions. The economic benefits, however, may not be evenly distributed across countries and among different groups within societies, potentially exacerbating existing inequalities.

Increasing electricity demand from the IT sector will stimulate overall supply, which—if sufficiently responsive—will lead to a small increase in electricity prices. More sluggish supply responses will lead to much stronger price surges. In the United States, the country with the largest expected surge in electricity demand, AI expansion alone could increase electricity prices by up to 9 percent, adding to price pressures coming from many other sources.

In addition, under current energy policies, the AI-driven rise in electricity demand could add 1.7 Gt in global greenhouse gas emissions between 2025 and 2030, an amount similar to Italy's energy-related GHG emissions over a five-year period. The social cost of these extra emissions is minor compared with the expected economic gains from AI, yet it still adds to the worrying buildup of worldwide emissions.

Demand for computing and electricity from AI service producers is subject to wide uncertainty, which may delay energy investments, causing underinvestment and higher prices. Policymakers and businesses must work together to ensure AI achieves its full potential, while minimizing societal costs.