Organisation for Economic Co-operation and Development

DAF/COMP/M(2025)3/ANN1/FINAL

Unclassified

English - Or. English

18 May 2026

DIRECTORATE FOR FINANCIAL AND ENTERPRISE AFFAIRS
COMPETITION COMMITTEE

Summary of Discussion of the Roundtable on Competition in AI infrastructure

Annex to the Summary Record of the 147th Meeting of the Competition Committee

4 December 2025

This document prepared by the OECD Secretariat is a detailed Summary of Discussion of the Roundtable on Competition in AI infrastructure, held by the Competition Committee on 4 December 2025. It presents a factual summary of the views expressed by speakers and delegations that intervened during the discussion.

More documents related to this discussion can be found at:
https://www.oecd.org/en/events/2025/12/competition-in-artificial-intelligence-infrastructure.html

Mr Ori SCHWARTZ
Email: ori.schwartz@oecd.org

JT03587125

Summary of Discussion of the Roundtable on Competition in AI infrastructure

On 4 December, the Competition Committee held a roundtable on Competition in Artificial Intelligence Infrastructure, chaired by **Benoît Cœuré**.

1. Introduction and Chairperson's Opening Remarks

The Chair opened the session by recalling that the roundtable formed part of the Committee's broader programme of work on competition issues arising along the artificial intelligence value chain. He noted that this work had progressed in stages over recent years, beginning with discussions on competition and data, followed by a dedicated roundtable on cloud computing services, and more recently by discussions at the Global Forum on Competition on downstream competitive effects of AI. The Chair also referred to the Committee's earlier discussions held the same day on the use of AI tools by competition authorities themselves.

The Chair explained that the objective of the current session was to move upstream in the AI stack, focusing on the physical infrastructure underpinning AI systems, including chips, compute, data centres, cloud services, networking and energy inputs. He emphasised that these layers are increasingly central to competitive outcomes in AI markets and raise complex questions for merger control, antitrust enforcement, market studies and advocacy.

The Chair introduced the speakers for the session. He noted that the background paper had been prepared by the OECD Secretariat Competition Division and they would briefly present on this alongside 3 expert speakers including:

- **Jan-Peter Kleinhans** from the OECD Directorate for Science, Technology and Innovation (STI), an expert in semiconductors and the compute supply chain.
- **Professor Helena Perrone**, Professor of Economics at Toulouse Business School, and
- **Professor Marc Bourreau**, Professor of Economics at Télécom Paris.

Finally, the Chair outlined the structure of the session. The discussion would begin with presentations by the Secretariat and the STI Directorate to set out the background and supply-chain context. This would be followed by a presentation from Prof Perrone and first round of discussion focusing on mergers and merger control in AI infrastructure, including country experiences. After a short break, the second part of the session would begin with a presentation from Prof Bourreau and broaden the focus to antitrust enforcement, market studies, advocacy and industrial policy, drawing on further expert input and delegations' interventions.

2. Background to the topic – Secretariate Competition division and Directorate for Science, Technology and Innovation

The Secretariat introduced the background paper prepared for the roundtable.

The Secretariat explained that the background note was organised around three pillars: an overview of the AI infrastructure supply chain; a set of common market features with

implications for competition; and a discussion of potential antitrust tools, together with the extent of current competition authority activity.

The secretariat noted that while many end users of AI do not directly observe the underlying physical supply chain, AI infrastructure is becoming increasingly important both for financial markets and for real economic growth across jurisdictions.

The secretariat clarified the session's definition of AI infrastructure. For the purposes of the roundtable, the focus was on the physical infrastructure required for generative AI, explicitly putting inputs such as data and skills to one side. Within this scope, he identified relevant infrastructure components as cloud service providers and data centres; the energy and networking that power and connect them; and the chips and other key technologies used within data centres and across the semiconductor supply chain.

The secretariat then passed over Jan-Peter Kleinhans to provide a technology and supply-chain overview to ground the competition discussion in the practical reality of the AI infrastructure stack.

Jan-Peter Kleinhans introduced the OECD Directorate for Science, Technology and Innovation's work on semiconductors through the Semiconductor Informal Exchange Network (SIEN). He described SIEN's objective as providing policymakers with analysis of the global semiconductor value chain, in order to better understand resilience challenges and inform the design of semiconductor-related policies. He referred to a recently released "chips landscape" providing a global overview of semiconductor production capacity by economy, and noted complementary STI work on the interaction between semiconductors and artificial intelligence.

He presented AI infrastructure as a layered compute stack, in which chips sit in servers, servers sit in data centres, and software orchestrates the system as a whole. He emphasised that AI performance increasingly depends on system-level integration rather than on individual components, highlighting the importance of networking infrastructure and internal data-centre orchestration. He also underlined that AI infrastructure depends heavily on energy, water, and specialised skills.

A central element of his presentation concerned the tight coupling between software layers and specific accelerators. He explained that widely used AI models are written using frameworks such as PyTorch and TensorFlow, which rely on lower-level software abstraction layers provided by accelerator manufacturers. These software environments are manufacturer-specific and not interchangeable: training on NVIDIA accelerators requires CUDA, while training on AMD accelerators requires ROC-m. This dependence was identified as a significant competitive factor.

Turning to accelerators, he described NVIDIA and AMD as the leading suppliers of AI accelerators, with hyperscalers such as Google, Amazon, and Microsoft among their largest customers. He noted that much of the market is concentrated among a small number of firms. He then explained the growing trend of hyperscalers developing their own AI accelerators to reduce dependence on NVIDIA and to design chips more specialised for machine-learning workloads, citing Google's TPU as a prominent example. He clarified that such accelerators are specialised for AI workloads but are not tied to a single AI model.

He noted that hyperscalers generally do not manufacture chips themselves, but instead collaborate with chip design houses such as Marvell, Broadcom, and MediaTek. He also highlighted a critical bottleneck in the supply chain: all AI accelerators for training in the cloud depend on high-bandwidth memory (HBM), which is supplied by only three companies globally—SK Hynix, Samsung, and Micron.

He then described the post-design stages of the semiconductor supply chain, including wafer fabrication and assembly, testing, and packaging. He explained that for cutting-edge AI accelerators, only two companies globally—TSMC and Samsung— currently have the capability to produce at the necessary scale. Each process step relies on highly specialised inputs, including advanced manufacturing equipment, specialised chemicals, and silicon wafers, with critical inputs often supplied by only one to three firms globally. As a result, effective multi-sourcing at the technological frontier is extremely limited.

He concluded by highlighting several competition-relevant observations: constraints on multi-sourcing are structural; hyperscaler self-supply reflects efforts to reduce dependence on dominant suppliers; and recent moves by Google to sell its latest accelerator to external customers represent a new development among U.S. hyperscalers. He highlighted CUDA as a particularly important proprietary software component in NVIDIA's accelerator ecosystem. He noted that access to, and dependence on, such proprietary software can be relevant for competition analysis, as it may affect interoperability, switching costs, and the ability of rival hardware suppliers to compete effectively.

Mr Kleinhans also stressed that competition conditions differ between AI training and AI inference. He explained that training typically requires very large-scale compute capacity and specialised accelerators, while inference places different performance and cost demands on hardware. As a result, market structure and competitive dynamics may differ across these two stages, and he suggested that they should be analysed distinctly for competition purposes.

Following the supply-chain presentation, the Secretariat, represented by Greg Jackson, outlined a set of common market features identified in the background note across many layers of AI infrastructure. These included high levels of innovation, high concentration and barriers to entry, increasing vertical integration and partnerships, evolving supply–demand imbalances, and significant levels of state intervention.

The Secretariat first highlighted the high level of innovation observed across AI infrastructure markets. Noting that innovation is generally a positive outcome, but that rapid technological change can complicate competition analysis, particularly in merger control, as authorities are required to assess likely future market developments under conditions of uncertainty. The Secretariat also explained that innovation and intellectual property may shift competitive interaction towards competition for the market rather than competition within the market, where technological breakthroughs can lead to step changes in provision.

The Secretariat then turned to market concentration and barriers to entry. Noting that many AI infrastructure markets are highly concentrated, reflecting not only successful innovation but also persistently high barriers to entry. These barriers include very large capital requirements, strong economies of scale, long lead times to market, and significant technological uncertainty. Taken together, these factors mean that only a limited number of well-funded firms are able to credibly enter or expand at scale. In such environments, the Secretariat cautioned that competition authorities may need to remain attentive to the risk that dominant firms seek to maintain market positions through anti-competitive conduct as markets evolve.

The Secretariat next addressed vertical integration, partnerships and investment trends. It noted a growing use of long-term purchasing commitments, strategic partnerships, equity investments and acquisitions across the AI infrastructure supply chain, with substantial transaction activity observed in recent months. While acknowledging that such arrangements may have efficiency justifications, including securing supply and supporting investment, it highlighted potential risks in concentrated markets. These risks include

exclusionary behaviour, preferential treatment of affiliated entities, and reduced access for independent competitors.

The secretariat then talked to supply–demand conditions and shortages. The Secretariat noted that several parts of the AI infrastructure supply chain have recently experienced excess demand, including acute shortages in certain server memory markets. It explained that these conditions may have ambiguous competitive effects. On the one hand, shortages can lead to price increases. On the other hand, they may temporarily lower barriers to entry by encouraging customers to consider alternative suppliers and by improving access to financing. The Secretariat cautioned, however, that these dynamics could change rapidly if demand expectations are not realised, potentially leading to sharp adjustments in market conditions.

Finally, the Secretariat addressed state intervention and industrial policy. It observed that public intervention has long been a feature of semiconductor manufacturing, but that its scale and intensity have increased in recent years due to the growing economic and geopolitical importance of AI infrastructure. It emphasised the role of competition authorities in advocating for the design of such interventions in ways that support effective competition where possible and limit unnecessary distortions.

In concluding, the Secretariat summarised the competition policy toolkit discussed in the background note. It referred to merger control, given the volume of transactions observed in the sector; antitrust enforcement, noting that while few AI-specific cases have been completed to date there are relevant historical precedents; market studies, as a means of building understanding in complex and rapidly developing markets; international co-operation, reflecting the global nature of supply chains; and advocacy, including the potential use of guidance or warning letters to influence conduct without waiting for lengthy enforcement proceedings.

3. Vertical links, merger control theories of harm, and AI-specific mechanisms

The Chair invited Professor Helena Perrone, Professor of Economics at Toulouse Business School, to open the first round of discussion with an expert perspective on mergers and vertical integration in AI infrastructure.

Professor Helena Perrone focused on competition issues upstream in the AI stack, particularly chips and compute, and framed her intervention around vertical integration, vertical agreements, and minority shareholdings that could affect competition. She observed that the AI stack is currently more concentrated upstream, including in chip design, manufacturing, and compute, than in downstream applications. She linked market power increasingly to access to critical inputs rather than access to users or data, contrasting this with earlier digital markets such as search or social networks. She highlighted the scarcity of GPUs and compute as a structural bottleneck, and noted that high fixed costs and economies of scale reinforce incumbency and create barriers to entry. She also referred to evidence that many AI-related acquisitions by major firms occur in upstream layers of the stack.

Prof. Perrone described different forms of vertical links across the AI stack. These included full vertical integration through mergers, citing Microsoft/Nuance, Intel/Habana Labs, and NVIDIA/Mellanox, as well as strategic partnerships and agreements short of control. She noted that such arrangements may generate competitive effects similar to mergers where they align incentives or enable foreclosure, and observed that they frequently involve preferential access to compute, special pricing, or early access to model capabilities.

Turning to theories of harm, she noted that many are standard in vertical settings but may arise in AI-specific ways. She identified risks of input foreclosure, including restricting or delaying access to GPUs or compute, or raising rivals' costs through discriminatory pricing for downstream developers. She also highlighted preferential access mechanisms, such as priority GPU access or discounted compute, which could affect innovation speed. In addition, she referred to possible conglomerate effects through bundling of compute, models, and application programming interfaces, which could increase switching costs and lock users into particular ecosystems.

Prof. Perrone further raised concerns about cross-ownership softening competition through the internalisation of rivals' profits, and about information flows within partnerships or integrated structures. She noted that the sharing of sensitive information, such as non-public technical details or demand forecasts, could create competitive advantages and may discourage unintegrated firms from sharing information with suppliers integrated with their rivals.

On potential pro-competitive effects, she explained that vertical integration can strengthen incentives to invest in innovation by internalising benefits across layers of the AI stack. She also recalled traditional efficiencies, including reduced double marginalisation, supply-chain stabilisation, and scale efficiencies. She highlighted that vertical integration or agreements may also be used to reduce dependence on dominant suppliers, for example where cloud providers invest in chip design or where NVIDIA partners with smaller cloud providers.

She emphasised that the balance between pro- and anti-competitive effects depends on adoption dynamics, multi-homing, and switching costs. She noted that downstream segments currently appear competitive, but cautioned that this could change as ecosystems develop and switching costs increase. She concluded that economic power currently lies upstream in chips, GPUs, and compute; that multiple instruments beyond mergers are used to create vertical linkages; and that authorities may need to consider a more unified approach to merger control, including the treatment of non-controlling stakes. She reiterated that access to chips and compute is, at present, the main source of market power in AI markets.

The Chair thanked Professor Perrone for her presentation and noted that it provided a clear analytical framework for understanding competition issues upstream in the AI stack. He observed that her intervention resonated with broader discussions within the Committee on mergers, including in relation to vertical effects, non-controlling interests, and supply-chain resilience. He also noted that the presentation highlighted the need for authorities to consider both potential competitive risks and efficiencies when assessing vertical arrangements in fast-moving and technologically complex markets such as AI infrastructure. The Chair then invited delegations to present their recent experience in applying merger control to AI infrastructure markets.

The **Korea** delegation presented the Korea Fair Trade Commission's (KFTC) recent merger reviews in the AI semiconductor sector. It highlighted several transactions reviewed since 2021, including SK Hynix's acquisition of Intel's SSD business, AMD's merger with Xilinx, and AMD's acquisition of shares in ZT Systems. These transactions were cleared unconditionally, as the KFTC identified a low risk of foreclosure or coordinated effects.

The delegation contrasted these clearances with the Synopsys/Ansys transaction reviewed in 2024–2025, which the KFTC approved in March 2025 subject to divestitures. It described this case as a significant conditional approval in the AI infrastructure space. The delegation explained that the review reflected the global nature of semiconductor design tools and involved extensive stakeholder outreach, including engagement with domestic

and foreign firms, the use of technical expertise, and close coordination with competition authorities in the European Union, the United Kingdom and the United States.

The delegation explained that both Synopsys and Ansys supply key software tools used in early-stage semiconductor design, and that the importance of such tools has increased as AI semiconductor development has accelerated. It identified unilateral effects concerns in three markets: register-transfer-level power consumption analysis software, optical design software, and photonics software. These concerns were grounded in high levels of concentration, with the merged entity exceeding 50% market share in each market, the elimination of direct competition between the two leading suppliers, and substantial barriers to entry linked to intensive R&D investment, specialised expertise, sunk costs and limited customer willingness to switch.

The KFTC also examined potential foreclosure strategies, including bundling of semiconductor design intellectual property with design tools, portfolio bundling, and the degradation of interoperability with rivals' products. The delegation reported that conglomerate foreclosure concerns were ultimately dismissed, noting in particular that bundling was constrained by incompatible business models, that customers typically mix and match best-performing products, and that interoperability is shaped by international standards.

Nevertheless, the delegation explained that the KFTC imposed structural remedies to address horizontal concerns in the affected markets. These included divestitures of relevant assets in register-transfer-level power analysis software and in optical and photonics software, covering the activities of the merging parties and certain affiliated entities.

The delegation highlighted the procedural significance of this case, noting that it was the first major application of a correction scheme submission system introduced in 2024. Under this mechanism, the parties submitted proposed divestitures, which the KFTC refined following market testing with rivals and customers.

Finally, the delegation described the KFTC's broader approach to assessing mergers in innovation-intensive markets. It noted that where sales-based market shares are not informative, the KFTC may rely on indicators such as R&D expenditure, specialised R&D assets and capabilities, patent portfolios and citations, and the number of meaningful innovators. In assessing innovation effects, the KFTC examines whether the merging parties are key innovators, the closeness of their rivalry, the competitive landscape post-merger, and relative technological strengths.

The **European Commission** described its recent experience in applying merger control to AI-relevant infrastructure markets. It emphasised that while AI technologies can support economic growth and competitiveness, merger control remains necessary to ensure that consolidation does not restrict access to critical computing resources in ways that harm competition and innovation in downstream markets.

The Commission first summarised its review of Broadcom/VMware in 2023. It explained that the transaction raised conglomerate foreclosure concerns due to Broadcom's strong position in Fibre Channel Host Bus Adapters and VMware's role as a leading provider of virtualisation software. The Commission noted that virtualisation software is an important component of AI infrastructure, as it enables efficient sharing of hardware and optimisation of resource allocation. Its investigation found that the merged entity could have both the ability and incentive to restrict interoperability between VMware's software and rival hardware components.

The Commission explained that potential foreclosure could arise through a range of mechanisms, including the exchange of information on application programming

interfaces, access to development tools and details of future product versions, as well as through VMware's certification processes for interoperability. In assessing incentives, the Commission described its use of critical switching analysis. It explained that even where a large number of customers would need to switch for foreclosure to be profitable, incentive could still exist given purchasing structures in which original equipment manufacturers select hardware components, while end customers choose virtualisation software and make significant investments in software infrastructure and staff training, thereby requiring continued compatibility between hardware and software.

To address these concerns, the Commission accepted a package of behavioural commitments. These included long-term commitments to ensure interoperability, guaranteed access to application programming interfaces, tools and technical support for development and certification, parity in the timing of access relative to Broadcom, open-source licensing of driver source code for current and future products, and organisational separation between hardware and software certification teams. The Commission noted that these remedies illustrate the importance of interoperability and component-level choice in safeguarding competition in AI-relevant infrastructure markets.

The Commission then turned to NVIDIA/Run:ai, which was cleared unconditionally in 2024. It explained that the case involved a conglomerate relationship between NVIDIA's data-centre GPUs and Run:ai's GPU orchestration software. While the products did not overlap, orchestration software must be compatible with GPUs, leading the Commission to assess whether the transaction could result in denial or degradation of interoperability.

The Commission described its assessment of a bi-directional conglomerate foreclosure theory, under which foreclosure would be profitable only if it strengthened NVIDIA's dominant position in GPUs. It reported that it found neither the technical ability nor the incentive for NVIDIA to foreclose competing orchestration software providers. In particular, it noted that key tools required for interoperability are open source or otherwise openly available and remain accessible once in the public domain. The Commission also observed that NVIDIA lacked visibility as to which third-party software relied on which tools, limiting any ability to selectively restrict access. Market participants indicated that restriction of such tools was unlikely, and the Commission considered that foreclosure would be unprofitable, given the relatively small size of the orchestration market compared with NVIDIA's much larger and profitable GPU business.

In concluding, the Commission reiterated that consolidation in AI-relevant infrastructure markets may influence access to critical computing inputs and affect downstream competition and innovation. It stated that it would continue to closely scrutinise transactions involving AI-critical inputs to prevent foreclosure of independent suppliers and to protect customer choice.

4. Delegate questions and further interventions

The Chair thanked the delegations for their merger case presentations and opened the floor for questions and additional interventions. He invited delegations to raise issues or share perspectives arising from the earlier presentations.

Japan asked the OECD Science, Technology and Innovation Secretariat to clarify which elements of the AI compute stack are protected by intellectual property, noting that semiconductor fabrication know-how often appears to rely on trade secrets rather than formal IP rights.

In response, **Jan-Peter Kleinhans** explained that intellectual property plays an important role primarily at the chip design stage. He described semiconductor IP blocks as pre-designed functional modules that are licensed from third-party suppliers and integrated into chip designs, alongside components developed internally. As an example, he explained that the interface between a GPU compute core and high-bandwidth memory may itself be licensed as an IP block. He further noted that electronic design automation tool providers such as Cadence, Synopsys and Siemens also develop semiconductor IP that is pre-qualified at specific foundries, which can reduce development time for chip designers. Responding to a follow-up question on pricing, he explained that there is no standard price for such IP, and that terms are typically determined through confidential bilateral negotiations depending on complexity and use.

The **Chinese Taipei** delegation highlighted the central role of the semiconductor industry in its economy, noting that TSMC alone accounts for a significant share of national output, and that firms based in Chinese Taipei are also active in chip testing and packaging, servers, and modular data centres. The delegation described growing integration and collaboration across the semiconductor supply chain, referring in particular to a recently cleared joint venture between Foxconn and SoftBank relating to servers and modular data centres, as well as a joint purchasing initiative involving firms from multiple sectors to secure green energy in response to AI-driven demand and sustainability requirements.

The delegation also reported that its competition authority had commissioned research on the semiconductor sector and completed a public consultation on competition issues related to generative AI. It noted mixed feedback, with some stakeholders expressing concern that dominant AI firms could leverage market power into adjacent markets, while others pointed to the pace of technological change as a potential constraint on such behaviour. Finally, the delegation highlighted certain non-technological features of TSMC's business model that may be relevant for competition analysis, including its pure-play foundry strategy and a conservative approach to expansion, which could inform assessments of incentives to foreclose or to engage in exclusionary strategies.

5. Broader antitrust concerns and policy interventions beyond enforcement

The Chair introduced the next part of the session, explaining that the discussion would broaden beyond merger control to consider wider antitrust and policy issues related to AI infrastructure. He invited Professor Marc Bourreau, Professor of Economics at Télécom Paris, to share his perspective on competition dynamics and possible policy approaches in AI infrastructure markets.

Professor Marc Bourreau broadened the discussion to consider competition and policy issues related to AI infrastructure. He recalled the high expectations associated with AI in terms of productivity, competitiveness and innovation, while noting that adoption by firms remains uneven. He emphasised that access to key inputs, in particular advanced compute, plays an important role in shaping competitive outcomes in AI markets.

He focused on AI accelerators and accelerated compute as critical inputs. He described the AI chip market as highly concentrated, referring to market analysis reported in the European Commission's NVIDIA/Run:ai decision. According to figures cited for 2024, NVIDIA holds a very large share of the data-centre GPU market by both volume and value. He noted that these shares have remained relatively stable despite rapid market growth, including a significant expansion in demand between 2022 and 2023.

Professor Bourreau identified two main factors contributing to this concentration. On the supply side, he referred to strong economies of scale and scope, high research and

development costs, and accumulated know-how. On the demand side, he highlighted complementarities and indirect network effects arising from the close linkage between NVIDIA's GPUs and its proprietary CUDA software platform, which reinforce adoption on both sides.

At the same time, he observed that the market may remain contestable. He pointed to efforts by other chip manufacturers, hyperscalers and large AI developers to develop or procure alternative solutions, including through internal development. He also noted the presence of more specialised suppliers offering differentiated products, for example targeting particular use cases or energy efficiency. Taken together, he described the market as highly concentrated but potentially contestable under current conditions.

He then referred to potential competition concerns associated with market power, including excessive pricing, restrictive contractual terms and exclusive dealing. He also noted that vertical expansion into adjacent activities, such as cloud services, could give rise to incentives for self-preferencing or discriminatory conduct in the allocation of compute resources.

Professor Bourreau emphasised that competition analysis should not focus solely on AI chip markets. He noted that AI developers often choose between building their own compute infrastructure and accessing accelerated compute through cloud providers. He observed that cloud-based access offers variable cost scaling and reduced exposure to technological obsolescence, which may be particularly relevant for smaller developers.

He characterised cloud computing markets as generally less concentrated than AI chip markets, with several large providers alongside a range of other global, regional and specialised firms. He identified switching costs as a potential concern, distinguishing between technical barriers related to system complexity and learning effects, and commercial barriers such as data transfer fees and discount or credit schemes that may contribute to customer lock-in. He concluded by noting that measures supporting interoperability or portability could help reduce switching costs and strengthen contestability, while allowing innovation to continue.

6. Findings from market studies: training vs inference, cloud concentration, and lock-in risks

The **Chair** noted that the discussion would now turn to findings from recent market studies conducted by competition authorities. He invited delegations to share their results, beginning with Japan.

Japan summarised the Japan Fair Trade Commission's discussion paper on generative AI published in October 2024 and its subsequent report issued in June 2025. It noted that while the analysis focused on domestic markets, many of the relevant firms operating across the AI value chain are global.

With respect to the infrastructure layer, the delegation highlighted findings concerning computing resources. It noted that GPUs are particularly well suited to training high-performance generative AI models and referred to a global GPU market in which a single manufacturer accounts for a very high share. At the same time, it emphasised that competitive conditions differ between the training and inference phases. The delegation explained that training typically requires very large computational capacity, making GPUs the dominant technology for parallel processing and supporting strong positions for leading suppliers with well-developed ecosystems. By contrast, inference generally requires less

compute and places greater emphasis on speed and energy efficiency, where the delegation observed more intense competition among chip suppliers.

The delegation also discussed inter-layer dependencies across the AI stack. It noted that only a limited number of firms, including large technology companies and some leading start-ups, possess their own computing infrastructure, while most AI developers rely on specialised cloud service providers to access GPU capacity. It reported indications of competition in cloud services, including through entry by domestic providers and differentiation in response to rising demand. At the same time, it observed that competition remains centred on a small number of leading global providers, which retain high shares in cloud-based GPU services due to their investment capacity and access to cutting-edge accelerators.

A central theme of the analysis concerned switching costs and potential lock-in. The delegation explained that dependence on specific semiconductor architectures and development environments can make it technically difficult and costly to switch between chip suppliers, often requiring software redesign and time. It noted that where a chip choice is closely linked to a particular cloud environment, switching chips may also require switching cloud providers, further increasing switching costs. The delegation added that stakeholder views were mixed, with some reporting that switching remains feasible and that certain cloud providers compete by facilitating migration, while others emphasised rising barriers.

Finally, the delegation noted that these structural features could give rise to competition concerns under Japan's Anti-Monopoly Act if firms with strong positions engage in conduct such as access restrictions, exclusionary practices or tying, potentially leading to foreclosure in downstream AI markets. It concluded by stating that the JFTC will continue to monitor developments closely given the rapid evolution of generative AI markets.

Portugal described a series of short issues papers intended to keep pace with rapid developments in AI markets, with the most recent work focusing on access to advanced chips, in particular GPUs. It reiterated concerns about high capital intensity, economies of scale and concentration across upstream stages of the semiconductor supply chain, including chip and software design, memory components, manufacturing and advanced packaging.

The delegation highlighted the role of network effects and ecosystems, noting that GPU suppliers often bundle software components enabling parallel computing. It referred in particular to CUDA as a widely adopted software environment that is compatible only with NVIDIA hardware. While acknowledging the benefits this ecosystem offers to developers, the delegation noted that such arrangements may contribute to lock-in effects and reinforce market power.

Finally, the delegation flagged the increasing prevalence of partnerships and long-term agreements involving chip producers as a potential competition risk warranting continued monitoring.

Korea summarised findings from the Korea Fair Trade Commission's study Generative AI and Competition, with emphasis on the cloud layer of AI infrastructure. It explained that the study analyses the generative AI ecosystem across three layers: infrastructure, foundation models and AI services.

With respect to cloud services, the delegation described competition in Korea between global cloud service providers, including Amazon, Microsoft and Google, and domestic providers such as Naver Cloud, KT Cloud and NHN Cloud. It observed that global providers benefit from extensive data-centre networks, strong investment capacity and

early access to cutting-edge accelerators, while domestic providers compete through localised technical support, compliance with Korean regulatory and security requirements, and strength in sectors such as public administration and financial services.

The delegation highlighted potential risks associated with restricted access to GPUs or large-scale compute, noting that market power at the infrastructure layer could be transferred downstream, particularly in the presence of network effects and high switching costs. At the same time, it pointed to mitigating factors, including the use of multi-cloud strategies and domestic efforts to diversify AI infrastructure capacity through investments in AI chips and data centres.

The delegation stated that the study did not identify clear evidence of illegal conduct at this stage. It concluded by noting that the KFTC will continue to monitor access conditions to ensure they remain fair and non-discriminatory, and may provide policy recommendations aimed at reducing excessive dependence on individual providers and improving interoperability, particularly in the context of public-sector cloud adoption.

7. Antitrust enforcement precedents and interim measures

The Chair noted that, in addition to mergers, advocacy and market studies, antitrust enforcement remains an important tool in AI infrastructure markets, and he invited the European Commission to present relevant enforcement experience.

The **European Commission** discussed antitrust enforcement cases in semiconductor markets that pre-date the emergence of generative artificial intelligence, presenting them as relevant precedents illustrating vulnerabilities of fast-moving and strategically important input markets to exclusionary conduct by dominant firms.

The Commission first referred to the Intel case concerning central processing units. It recalled that Intel had granted conditional rebates and made direct payments to computer manufacturers with the objective of halting or delaying the launch of products using competitors' CPUs. While the elements of the decision relating to conditional rebates were annulled by the courts, the findings relating to direct payments, characterised as naked restrictions, were upheld, leading to a renewed Commission decision adopted in 2023. The Commission noted that the case demonstrates how exclusionary practices affecting critical inputs can have significant competitive impact.

The Commission then referred to the Broadcom case concerning chips for television set-top boxes and modems. It explained that Broadcom had entered into exclusivity-inducing agreements with customers, prompting the Commission to adopt interim measures to prevent market foreclosure, followed by the acceptance of commitments. The Commission emphasised that interim measures can play an important role in preventing serious and irreparable harm to competition in fast-moving technology markets, while longer-term solutions are developed.

The Commission also recalled the Qualcomm case relating to 3G baseband chips, in which Qualcomm was found to have engaged in predatory pricing aimed at excluding a competitor. It noted that the decision was upheld by the General Court, reinforcing the application of exclusionary pricing principles in technology-intensive markets.

While these cases arose prior to the widespread deployment of AI technologies, the Commission noted that they highlight structural risks that may also arise in AI-related infrastructure markets, particularly where access to chips or large-scale compute acts as a bottleneck for downstream competition. The Commission referred to its policy work on competition in AI and virtual environments, which identifies chips as a potential constraint

on entry and expansion by AI competitors, and indicated that insights from past enforcement experience are being reflected in ongoing work on guidance concerning dominant-firm conduct, including in the context of draft Article 102 guidelines.

8. Industrial policy and public infrastructure

The **Chair** observed that competition authorities may need to take account of industrial policy and regulatory interventions affecting AI infrastructure. He then invited delegations to present their experience with such measures.

France described recent experience with public-sector interventions relevant to AI infrastructure, emphasising the potential role of public policy in supporting competition and market contestability. It noted that AI infrastructure raises particular challenges due to high fixed costs, capital intensity and scale requirements, which can make access difficult for smaller firms.

The delegation explained that public intervention can, under certain conditions, have pro-competitive effects, particularly where it facilitates entry or reduces reliance on dominant private providers. In this context, it highlighted the role of public infrastructure and shared facilities in enabling smaller firms to access essential AI inputs, including compute and data resources, on more equal terms.

At the same time, the delegation stressed that such interventions involve trade-offs and must be carefully designed. It noted that public funding and infrastructure provision can distort market incentives if they crowd out private investment or entrench incumbents, and therefore require careful calibration. In its view, competition authorities have an important role to play in ensuring that industrial-policy measures are implemented in a manner consistent with competitive principles.

The delegation concluded by emphasising that public intervention in AI infrastructure should aim to strike a balance between supporting innovation and maintaining effective competition, and that competition analysis can contribute to assessing when and how such interventions are likely to be beneficial.

Israel described its participation in an inter-governmental task force examining the use of artificial intelligence in the financial sector, and summarised insights from an interim report published in late 2024. It noted that AI is expected to play an increasingly central role in Israeli financial services, including credit decisions, insurance underwriting, customer interaction, fraud detection and risk management.

The delegation emphasised that the benefits of AI adoption in these areas depend on the preservation of effective and dynamic competition, particularly in sectors that already exhibit high levels of concentration, such as banking and parts of insurance. It noted regulatory concern that market developments in AI are progressing rapidly, and that access to key inputs may become concentrated before effective safeguards are in place.

In this context, the delegation described recommendations aimed at ensuring that smaller competitors, including non-bank institutions, are able to access essential data and infrastructure. These included expanded access to aggregated and anonymised data from the Bank of Israel's credit database and broader use of open-banking frameworks. The delegation indicated that these measures seek to reduce barriers to developing AI applications in financial services and to limit excessive dependence on incumbents or large technology providers.

The **European Commission** outlined the EU state-aid framework applicable to semiconductors and AI-related infrastructure, describing three main instruments.

First, it referred to state aid for research and development, which generally raises limited competition concerns. Second, it described the use of Important Projects of Common European Interest (IPCEIs) to support innovative projects up to first industrial deployment, subject to conditions including cross-border collaboration and knowledge dissemination. It noted that two microelectronics IPCEIs were approved in 2018 and 2023, with a total state budget of approximately EUR 10 billion, and that additional IPCEIs relating to advanced semiconductor technologies, computing infrastructure and AI are under preparation.

Third, the Commission discussed manufacturing aid assessed directly under the Treaty following the 2023 EU Chips Act. It emphasised that such aid is considered highly distortive and therefore exceptional, justified in the semiconductor sector by large investment requirements, high barriers to entry, the strategic importance of chips for multiple industries, and the availability of subsidies outside the Union. In assessing these measures, the Commission examines whether facilities are first of a kind, whether aid is limited to the funding gap, and whether positive effects such as security of supply, innovation, workforce development and cross-border co-operation outweigh potential competition distortions.

The Commission noted that since 2022 it has approved nine manufacturing-related measures, with state budgets totalling approximately EUR 13 billion and expected to unlock around EUR 30 billion in investment across several Member States. These include projects involving firms of different sizes, including joint ventures such as ESMC. It also observed that the EU currently lacks manufacturing capacity for leading-edge AI chips, and that attracting such investment is a key objective of the ongoing review of the Chips Act. The Commission noted that recent investment decisions illustrate that state aid is only one of several factors influencing firms' location choices.

The **Business at OECD (BIAC)** delegation expressed agreement with the Secretariat's background paper and described it as providing a balanced overview of competition issues in AI-related infrastructure. BIAC emphasised that competition authorities have an important role to play in advocating for pro-competitive policies that preserve innovation and efficiency, particularly in the context of actual or potential state intervention.

BIAC described AI infrastructure markets as characterised by dynamic competition, high levels of research and development, and significant reliance on intellectual property. In this context, it cautioned against pre-emptive enforcement based on concerns that could risk undermining innovation and investment incentives. BIAC therefore emphasised the importance of rigorous, evidence-based competitive assessment, alongside pro-competitive advocacy, to support the scale of investment required in AI infrastructure.

BIAC also underscored the heterogeneity across infrastructure layers, noting that different segments exhibit distinct competitive characteristics. It contrasted advanced chip markets, which are highly differentiated and IP-intensive, with data-centre construction, where entry may be more feasible and IP intensity lower. It further observed that energy supply and digital networking infrastructure often serve multiple uses and are subject to sector-specific regulation. On this basis, BIAC stressed that there is no single competitive paradigm applicable across all layers of AI infrastructure and that analysis should be tailored accordingly.

With respect to vertical mergers and partnerships, BIAC concurred that such arrangements may generate efficiencies in terms of innovation, coordination and investment, and should be assessed on a case-by-case basis. At the same time, it reiterated concerns that state

intervention can distort competition if not carefully designed, and referred to competitive neutrality principles as a useful reference point.

The **United States** delegation described recent U.S. initiatives relating to AI infrastructure, referring in particular to an "America's AI Action Plan" organised around three pillars: accelerating AI innovation, building AI infrastructure, and strengthening international co-operation on AI diplomacy and security. It explained that the infrastructure pillar includes measures aimed at streamlining permitting for data centres, semiconductor manufacturing facilities and energy infrastructure, alongside efforts to ensure security and expand domestic semiconductor production capacity.

The delegation noted that competition authorities can contribute to these objectives by supporting innovation, promoting entry and ensuring competitive AI markets. It also referred to recent competition-policy work touching on AI, including issues raised earlier in the week in relation to search markets.

The delegation further referenced a joint comment submitted with the U.S. Patent and Trademark Office in proceedings before the International Trade Commission concerning patented AI-related technologies, emphasising the importance of protecting innovation in AI markets.

The United States added that the FTC and the Department of Justice are supporting efforts by the Department of Energy to increase energy production for AI data centres, highlighting access to affordable and reliable power as a key input for AI infrastructure. The delegation referred to recent initiatives aimed at reinvigorating the nuclear energy industrial base and to co-operation with competition authorities under defined legal frameworks to enable certain forms of coordination in line with the Defence Production Act.

The United States delegation concluded by emphasising that the benefits of AI should not be viewed as a zero-sum outcome, and that fair and open regulatory frameworks, combined with effective competition, are essential to ensuring that AI delivers broad economic benefits.

9. Closing reflections

The **Chair** invited the expert speakers to offer any final reflections in light of the discussion.

In the final exchange, **Professor Marc Bourreau** responded to Professor Perrone's earlier observation that market power in AI markets currently derives more from access to inputs than from access to users. He suggested that for providers of accelerated compute, access to particular user groups, notably start-ups, may still be commercially significant. In this context, he referred to practices such as discounted cloud access, which may contribute to customer lock-in over time.

Professor Helena Perrone responded by emphasising the high degree of uncertainty surrounding the evolution of AI markets. She explained that her earlier assessment reflected a snapshot of market conditions at the present stage, and that compared with earlier digital markets such as search, access to users appears less decisive today, while remaining potentially important. She stressed that markets remain in formation, and that adoption dynamics, multi-homing behaviour and switching costs could evolve in ways that increase the relevance of user access. She referred in this context to observations made earlier in the discussion, including by Japan, regarding increasing switching costs.

The **Chair** thanked the experts, delegations and the Secretariat for their contributions. He noted that the discussion had highlighted the need for continued work to deepen understanding of the AI infrastructure stack while keeping pace with rapidly evolving market developments. He indicated that the Committee would continue to consider these issues in its future work.

The Chair then closed the session.