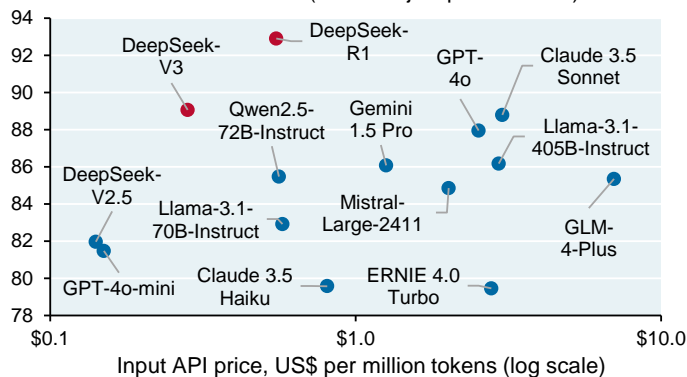


**The sincerest form of flattery: on DeepSeek, NVIDIA, OpenAI and the futility of US chip bans**

The DeepSeek episode can be two things at once: (i) a reflection of impressive Chinese AI innovation in the face of US chip bans and other restrictions, and (ii) the by-product of probable terms of service and copyright violations by DeepSeek against OpenAI. A Shakespearean irony: OpenAI may have had its terms of service violated after spending years training their own models on other people's data.

**Cost vs performance of select AI models**

MMLU Redux ZeroEval Score (multi-subject performance)



Source: DeepSeek, JPMAM, January 2025

How does a small Chinese company like DeepSeek with less than 200 employees offer strong language model performance at low cost?

The short list:

- Mixture of experts
- Multi-head latent attention
- Lower precision parameter storage during model training
- Multi-token prediction
- Reinforcement “chain of thought” learning without human feedback

**What did DeepSeek announce? Let's start with its V2 model introduced last December:**

- DeepSeek appears to have trained its models 45x more efficiently than other leading-edge models. To be clear, most of DeepSeek's approaches already existed. It's greatest accomplishment: figuring out how to deploy them all at once in the face of a chip ban, and introduce its own self-reinforcement learning
- **Mixture of Experts:** GPT-3.5 uses its entire model for both training and inference to solve problems despite the fact that only a small part of the model might be needed. In contrast, GPT-4 and DeepSeek are mixture of experts (MoE) models which only activate the parts of the model that are needed to solve each problem. DeepSeek V3 is quite massive with 671 billion parameters, but only 37 billion are active at any given time
- **MLA** refers to “multi-head latent attention”, which is jargon for how DeepSeek maintains a smaller memory cache while running
- Other DeepSeek efficiency approaches: while parameters are stored with BF16 or FP32 precision, they are reduced to **FP8 precision** for training purposes<sup>1</sup>. The models also use multi-token prediction (**MTP**) rather than just predicting the next token, which reduces accuracy by ~10% but doubles inference speed
- DeepSeek claims that V3 was very cheap to train, requiring 2.7 mm H800 GPU hours which at a cost of \$2/GPU hour is just \$5.6 million<sup>2</sup>. The comparable number of GPU hours for the Llama 3.1 405B final training run was ~10x higher<sup>3</sup>. DeepSeek made clear that this was the cost of the final training run, excluding “costs associated with prior research and ablation experiments on architectures, algorithms or data”
- DeepSeek V3 performance is competitive with OpenAI's 4o and Anthropic's Sonnet-3.5 and appears to be better than Llama's biggest model with lower training costs. DeepSeek provides API access at \$0.14 per million tokens while OpenAI charges \$7.50 per million tokens<sup>4</sup>; perhaps some degree of loss leader pricing
- DeepSeek may have “over-specified” its model: it tortured it to do well on the MMLU benchmark but when the questions changed slightly, its performance declined at a faster rate than other models did<sup>5</sup>. More analysis is needed to determine whether this overspecialization is a broader issue
- DeepSeek just announced another release this morning: a multi modal model (text, image generation and interpretation). Unsurprisingly, DeepSeek makes no pretense of data privacy and stores everything

<sup>1</sup> “DeepSeek-R1 and FP8 Mixed-Precision Training”, Colfax International, January 27, 2025

<sup>2</sup> “DeepSeek V3 and the actual cost of training frontier AI models”, Nathan Lambert, January 9, 2025

<sup>3</sup> “Meta claims world's largest open AI model with Llama 3.1 405B debut”, The Register, July 2024


<sup>4</sup> Venture Beat, January 27, 2025

<sup>5</sup> Bridgewater Daily Observations, January 27, 2025

### Did DeepSeek V2/V3 models benefit from “distillation”, which entails training a model by accessing other AI models? It sure looks that way

- DeepSeek trained its models on 14.8 trillion tokens, which is a massive sample similar to Llama
- Some AI analysts believe that DeepSeek sent prompts to a GPT-4 or Chat GPT teacher model, and then used the responses to train its own student model, at least for part of the training process<sup>6,7</sup>. Companies like OpenAI do this when deriving GPT-4 Turbo from GPT-4, but they are training their **own** models. Companies like OpenAI and Anthropic typically make clear that it would be a terms of service violation to use their models to train another model (although start-ups and researchers probably do this all the time such as the Stanford Alpaca project, which disclosed what it did)
- Going forward, will OpenAI and other LLM companies more aggressively monitor how/who/when/why their models are being used and control access to them via IP address banning or rate limiting? And will start-ups figure out ways to mask themselves?
- As for the open-source approach DeepSeek is taking, we wrote about such risks to closed source models a year ago. The chart below shows how adapted open-source models performed just as well as closed source models across multiple domains. We also cited the leaked and infamous Google memo entitled “We have no moat...and neither does OpenAI”
- The open-source issue may be a catalyst for the gradual divorce between Microsoft and OpenAI<sup>8</sup>. Microsoft presumably wants to provide inference to customers, but may be reluctant to fund billions for data centers to train models that may end up commoditized

**LLM performance on domain-specific multiple choice exams**



Source: “Adapting large language models via reading comprehension”, Huang et al, Microsoft, September 2023

**DeepSeek** also used its reinforcement learning models to distill Meta’s Llama and Alibaba’s Qwen into smaller versions to demonstrate how they can outperform GPT-4o and Claude 3.5 Sonnet in select math benchmarks. DeepSeek’s distilled models are open source and available on Hugging Face under an MIT license. Such distilled models may be a violation of Llama licenses as well, depending on how many active monthly users the end-product has

<sup>6</sup> One example: “**DeepSeek: is it a stolen ChatGPT?**” from Jan Kammerath, January 27, 2025

- When answering questions about VW car sales in China, ChatGPT, Grok and Gemini all gave very different answers, while DeepSeek’s answer was almost identically worded to ChatGPT
- Formatting is another highly identifiable LLM footprint. When asked to program an impossible graphics function, DeepSeek’s answer was 95% similar to ChatGPT but very different from the garbage that Co-Pilot, Grok and Gemini produced
- Why would a Chinese chatbot be trained on what happened at Tiananmen Square in 1989, and be so easily cajoled into talking about it? Why does it talk about Presidents and “best cities to live in” by talking about American ones, even when asked in German?
- Why would a Chinese chatbot refer to a single-party state as being “a dictatorship” and reject the one-party system unless it was trained on Western data with strong ideological beliefs?

<sup>7</sup> In December 2024, DeepSeek’s V3 was asked “**What model are you?**”. In five times out of eight, V3 actually answered by saying that it was ChatGPT. When asked about DeepSeek’s API, V3 responded by giving instructions on how to use OpenAI’s API. DeepSeek V3 even tells some of the same lame jokes at ChatGPT, including the punchlines. Source: TechCrunch, “**Why DeepSeek’s new AI model thinks it’s ChatGPT**”, December 2024

<sup>8</sup> “**Money, Power and AI: The Crumbling Alliance Between OpenAI and Microsoft**”, Alinvest.com, October 2024

**Why did DeepSeek's R1 announcement clobber NVIDIA, and what are implications for OpenAI and Anthropic?**

- DeepSeek's R1 is a chain-of-thought reasoning model like OpenAI's o1. It can think through a problem and produce higher quality results in areas like coding, math and logic<sup>9</sup>. As shown in the chart on the first page, R1 offers similar performance at lower costs
- The most important aspects of DeepSeek's R1 model were already known a month ago when DeepSeek V2/V3 was released. Equity markets started to pay more attention when DeepSeek's app became more popular than ChatGPT in the App store
- **One market shock:** even after acknowledging DeepSeek's probable piggy-backing off of OpenAI, China is further along on AI-LLMs than many market participants thought. AI-LLM breakthroughs are no longer just in the US domain
- **Another market shock:** more efficient training/inference processes and possible alternatives to NVIDIA software could eventually affect long run projections of NVIDIA's order book. One example: a company could conceivably run inference models on AMD GPUs which are half the price of NVIDIA on a \$/FLOP basis, if DeepSeek coding disclosures help users mitigate AMD's inferior chip-to-chip communications capabilities
- I've read in a few places that the US chip ban on China indirectly led to DeepSeek's success: by forcing China to innovate with less cutting edge hardware and software, Chinese engineers figured it out and developed innovations along the way<sup>10</sup>. One thing's for sure: DeepSeek's intention to make everything public stands in stark contrast to OpenAI's pronouncements at the time of GPT-2's release that they would not release datasets, training codes or model weights due to concerns of such data being misused by the great unwashed proletariat
- Where to from here for OpenAI, Anthropic, Cohere, Mistral, etc? The questions on how closed source AI models will monetize IP become more challenging to answer. Even Sam Altman acknowledged last night that "DeepSeek's R1 is an impressive model, particularly around what they are able to deliver for the price"

**What does the long run look like for big tech and consumer companies?<sup>11</sup>**

- Model commoditization and cheaper inference is probably good for Big Tech and large consumer-facing companies in the long run. The cost of providing inference models to customers would go down, which could increase AI adoption. That said, I cannot stop thinking about the massive amounts of money spent already on AI compute infrastructure, which we discussed on the first page of last week's piece
- Amazon could benefit; it hasn't created its own high-quality model, but can now benefit from low-cost, high-quality open source models like DeepSeek
- Apple's hardware could benefit from cheaper and more efficient inference models
- Meta could benefit as well since almost every aspect of its business is AI related at this point, although it will be important to follow the impact on Llama<sup>12</sup>
- Google may be less well positioned: in a world of possibly decreased hardware requirements, Google's TPUs are less of an advantage. Also, lower inference costs may increase the viability and likelihood of products that displace Google search
- All of these implications depend on whether DeepSeek and other low-cost, open-source models can thrive in a world where training data might not be as readily available

---

<sup>9</sup> During training, DeepSeek observed an "**aha moment**": a phase when the model spontaneously learned to revise its thinking process mid-stream when encountering uncertainty. This wasn't explicitly programmed and arose naturally from the interaction between the model and the reinforcement learning environment

<sup>10</sup> "*China's DeepSeek Shows Why Trump's Trade War Will Be Hard to Win*", Tyler Cowen, Bloomberg, Jan 9, 2025

<sup>11</sup> Stratechery Research, Ben Thompson, January 27, 2025

<sup>12</sup> "13 individuals working on Llama each earn more per year in total compensation than the combined training cost for the DeepSeek-V3 models which outperform it": Jeffrey Emanuel, GLG Insights, January 25, 2025

### How powerful are NVIDIA moats?<sup>13</sup>


- Most AI projects rely on NVIDIA’s CUDA software, which only works on NVIDIA chips. NVIDIA drivers are battle-tested and perform well on Linux (unlike AMD which is notorious for low quality and instability of their Linux drivers), and benefit from highly optimized open-source code in libraries like PyTorch. Nvidia also has a huge lead in terms of its ability to combine multiple chips together into one large virtual GPU. NVIDIA’s industry-leading interconnect technology dates back to its purchase of Mellanox in 2019
- But there have been competitors circling around NVIDIA for a while: Cerebras (create one massive chip rather than a lot of little ones, thus eliminating interconnection challenges); Groq (deterministic computing chips that can offer better economics if GPU utilization rates are high enough); and several companies that are attempting to design code that works on a variety of different GPUs and TPUs (MLX, sponsored by Apple; Triton, sponsored by OpenAI; and JAX, developed by Google)
- Yesterday was a “shoot first, ask questions later” market response; NVIDIA P/E based on forward earnings expectations declined towards the very low end of the range since 2020, assuming no material changes to NVIDIA’s order book...and that’s the big question

**Nvidia forward price to earnings ratio**



Source: Bloomberg, JPMAM, January 27, 2025

**US electricity demand forecast by application**



Source: EI, LBNL, Rystad, Evolved Energy, NREL, JPMAM, 2025

### What about implications for energy consumption due to more energy efficient training and inference models?

- We should all dial down the frenzy about increased electricity demand from data centers. Even before DeepSeek, there were already strong incentives to reduce training and computation costs by developing more energy efficient chips and to develop and apply software innovations that require less training, fewer model solutions and much less movement of model solutions between nodes/chips on the network
- Politics may slow US electricity demand growth as well. We will cover Trump 2.0 energy policies in more detail in the energy paper in March. The short version: solar, wind, battery, EV, carbon capture and other tax credits might be reduced through a Congressional reconciliation bill in which these reductions pay for tax cuts. Remember: tariffs don’t count towards reported fiscal outcomes unless they’re legislated (if tariffs are simply imposed by the President, they would not count as revenue offsets in a reconciliation process)
- The low end of the US electricity demand forecast above is growth of just 7%, even after including EVs, electrification of home heating and new data centers

Michael Cembalest  
JP Morgan Asset Management

Note: I was a French and Russian literature major in college from 1981-1984 so I apologize for any technical errors resulting from my liberal arts background

<sup>13</sup> “The short case for NVIDIA stock”, Jeffrey Emanuel, GLG Insights, January 25, 2025. Opinions on the relative merits of NVIDIA and AMD hardware/software differ; I am citing Jeffrey’s views in this section.

**IMPORTANT INFORMATION**

This material is for information purposes only. The views, opinions, estimates and strategies expressed herein constitutes Michael Cembalest's judgment based on current market conditions and are subject to change without notice, and may differ from those expressed by other areas of JPMorgan Chase & Co. ("JPM"). **This information in no way constitutes J.P. Morgan Research and should not be treated as such.** Any companies referenced are shown for illustrative purposes only, and are not intended as a recommendation or endorsement by J.P. Morgan in this context.

**GENERAL RISKS & CONSIDERATIONS** Any views, strategies or products discussed in this material may not be appropriate for all individuals and are subject to risks. Investors may get back less than they invested, and **past performance is not a reliable indicator of future results.** Asset allocation/diversification does not guarantee a profit or protect against loss. Nothing in this material should be relied upon in isolation for the purpose of making an investment decision.

**NON-RELIANCE** Certain information contained in this material is believed to be reliable; however, JPM does not represent or warrant its accuracy, reliability or completeness, or accept any liability for any loss or damage (whether direct or indirect) arising out of the use of all or any part of this material. No representation or warranty should be made with regard to any computations, graphs, tables, diagrams or commentary in this material, which are provided for illustration/ reference purposes only. Any projected results and risks are based solely on hypothetical examples cited, and actual results and risks will vary depending on specific circumstances. Forward-looking statements should not be considered as guarantees or predictions of future events. Nothing in this document shall be construed as giving rise to any duty of care owed to, or advisory relationship with, you or any third party. Nothing in this document shall be regarded as an offer, solicitation, recommendation or advice (whether financial, accounting, legal, tax or other) given by J.P. Morgan and/or its officers or employees; J.P. Morgan and its affiliates and employees do not provide tax, legal or accounting advice. You should consult your own tax, legal and accounting advisors before engaging in any financial transactions.

**For J.P. Morgan Asset Management Clients:**

J.P. Morgan Asset Management is the brand for the asset management business of JPMorgan Chase & Co. and its affiliates worldwide.

To the extent permitted by applicable law, we may record telephone calls and monitor electronic communications to comply with our legal and regulatory obligations and internal policies. Personal data will be collected, stored and processed by J.P. Morgan Asset Management in accordance with our privacy policies at <https://am.jpmorgan.com/global/privacy>.

**ACCESSIBILITY**

For U.S. only: If you are a person with a disability and need additional support in viewing the material, please call us at 1-800-343-1113 for assistance.

This communication is issued by the following entities: In the United States, by J.P. Morgan Investment Management Inc. or J.P. Morgan Alternative Asset Management, Inc., both regulated by the Securities and Exchange Commission; in Latin America, for intended recipients' use only, by local J.P. Morgan entities, as the case may be.; in Canada, for institutional clients' use only, by JPMorgan Asset Management (Canada) Inc., which is a registered Portfolio Manager and Exempt Market Dealer in all Canadian provinces and territories except the Yukon and is also registered as an Investment Fund Manager in British Columbia, Ontario, Quebec and Newfoundland and Labrador. In the United Kingdom, by JPMorgan Asset Management (UK) Limited, which is authorized and regulated by the Financial Conduct Authority; in other European jurisdictions, by JPMorgan Asset Management (Europe) S.à r.l. In Asia Pacific ("APAC"), by the following issuing entities and in the respective jurisdictions in which they are primarily regulated: JPMorgan Asset Management (Asia Pacific) Limited, or JPMorgan Funds (Asia) Limited, or JPMorgan Asset Management Real Assets (Asia) Limited, each of which is regulated by the Securities and Futures Commission of Hong Kong; JPMorgan Asset Management (Singapore) Limited (Co. Reg. No. 197601586K), which this advertisement or publication has not been reviewed by the Monetary Authority of Singapore; JPMorgan Asset Management (Taiwan) Limited; JPMorgan Asset Management (Japan) Limited, which is a member of the Investment Trusts Association, Japan, the Japan Investment Advisers Association, Type II Financial Instruments Firms Association and the Japan Securities Dealers Association and is regulated by the Financial Services Agency (registration number "Kanto Local Finance Bureau (Financial Instruments Firm) No. 330"); in Australia, to wholesale clients only as defined in section 761A and 761G of the Corporations Act 2001 (Commonwealth), by JPMorgan Asset Management (Australia) Limited (ABN 55143832080) (AFSL 376919). For all other markets in APAC, to intended recipients only.

**For J.P. Morgan Private Bank Clients:****ACCESSIBILITY**

J.P. Morgan is committed to making our products and services accessible to meet the financial services needs of all our clients. Please direct any accessibility issues to the Private Bank Client Service Center at 1-866-265-1727

**LEGAL ENTITY, BRAND & REGULATORY INFORMATION**

In the **United States**, **JPMorgan Chase Bank, N.A.** and its affiliates (collectively "**JPMCB**") offer investment products, which may include bank managed investment accounts and custody, as part of its trust and fiduciary services. Other investment products and services, such as brokerage and advisory accounts, are offered through **J.P. Morgan Securities LLC ("JPMS")**, a member of **FINRA** and **SIPC**. JPMCB and JPMS are affiliated companies under the common control of JPM.

In **Germany**, this material is issued by **J.P. Morgan SE**, with its registered office at Taunustor 1 (TaunusTurm), 60310 Frankfurt am Main, Germany, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB). In **Luxembourg**, this material is issued by **J.P. Morgan SE – Luxembourg Branch**, with registered office at European Bank and Business Centre, 6 route de Treves, L-2633, Senningerberg, Luxembourg, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Luxembourg Branch is also supervised by the Commission de Surveillance du Secteur Financier (CSSF); registered under R.C.S Luxembourg B255938. In the **United Kingdom**, this material is issued by **J.P. Morgan SE – London Branch**, registered office at 25 Bank Street, Canary Wharf, London E14 5JP, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – London Branch is also supervised by the Financial Conduct Authority and Prudential Regulation Authority. In **Spain**, this material is distributed by **J.P. Morgan SE, Sucursal en España**, with registered office at Paseo de la Castellana, 31, 28046 Madrid, Spain, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE, Sucursal en España is also supervised by the Spanish Securities Market Commission (CNMV); registered with Bank of Spain as a branch of J.P. Morgan SE under code 1567. In **Italy**, this material is distributed by **J.P. Morgan SE – Milan Branch**, with its registered office at Via Cordusio, n.3, Milan 20123, Italy, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Milan Branch is also supervised by Bank of Italy and the Commissione Nazionale per le Società e la Borsa (CONSOB); registered with Bank of Italy as a branch of J.P. Morgan SE under code 8076; Milan Chamber of Commerce Registered Number: REA MI

[2025 Eye on the Market Outlook](#)

2536325. In the **Netherlands**, this material is distributed by **J.P. Morgan SE – Amsterdam Branch**, with registered office at World Trade Centre, Tower B, Strawinskyalaan 1135, 1077 XX, Amsterdam, The Netherlands, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Amsterdam Branch is also supervised by De Nederlandsche Bank (DNB) and the Autoriteit Financiële Markten (AFM) in the Netherlands. Registered with the Kamer van Koophandel as a branch of J.P. Morgan SE under registration number 72610220. In **Denmark**, this material is distributed by **J.P. Morgan SE – Copenhagen Branch, filial af J.P. Morgan SE, Tyskland**, with registered office at Kalvebod Brygge 39-41, 1560 København V, Denmark, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Copenhagen Branch, filial af J.P. Morgan SE, Tyskland is also supervised by Finanstilsynet (Danish FSA) and is registered with Finanstilsynet as a branch of J.P. Morgan SE under code 29010. In **Sweden**, this material is distributed by **J.P. Morgan SE – Stockholm Bankfilial**, with registered office at Hamngatan 15, Stockholm, 11147, Sweden, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Stockholm Bankfilial is also supervised by Finansinspektionen (Swedish FSA); registered with Finansinspektionen as a branch of J.P. Morgan SE. In **Belgium**, this material is distributed by **J.P. Morgan SE – Brussels Branch** with registered office at 35 Boulevard du Régent, 1000, Brussels, Belgium, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE Brussels Branch is also supervised by the National Bank of Belgium (NBB) and the Financial Services and Markets Authority (FSMA) in Belgium; registered with the NBB under registration number 0715.622.844. In **Greece**, this material is distributed by **J.P. Morgan SE – Athens Branch**, with its registered office at 3 Haritos Street, Athens, 10675, Greece, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB); J.P. Morgan SE – Athens Branch is also supervised by Bank of Greece; registered with Bank of Greece as a branch of J.P. Morgan SE under code 124; Athens Chamber of Commerce Registered Number 158683760001; VAT Number 99676577. In **France**, this material is distributed by **J.P. Morgan SE – Paris Branch**, with its registered office at 14, Place Vendôme 75001 Paris, France, authorized by the Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin) and jointly supervised by the BaFin, the German Central Bank (Deutsche Bundesbank) and the European Central Bank (ECB) under code 842 422 972; J.P. Morgan SE – Paris Branch is also supervised by the French banking authorities the Autorité de Contrôle Prudentiel et de Résolution (ACPR) and the Autorité des Marchés Financiers (AMF). In **Switzerland**, this material is distributed by **J.P. Morgan (Suisse) SA**, with registered address at rue du Rhône, 35, 1204, Geneva, Switzerland, which is authorised and supervised by the Swiss Financial Market Supervisory Authority (FINMA) as a bank and a securities dealer in Switzerland.

In **Hong Kong**, this material is distributed by **JPMCB, Hong Kong branch**. JPMCB, Hong Kong branch is regulated by the Hong Kong Monetary Authority and the Securities and Futures Commission of Hong Kong. In Hong Kong, we will cease to use your personal data for our marketing purposes without charge if you so request. In **Singapore**, this material is distributed by **JPMCB, Singapore branch**. JPMCB, Singapore branch is regulated by the Monetary Authority of Singapore. Dealing and advisory services and discretionary investment management services are provided to you by JPMCB, Hong Kong/Singapore branch (as notified to you). Banking and custody services are provided to you by JPMCB Singapore Branch. The contents of this document have not been reviewed by any regulatory authority in Hong Kong, Singapore or any other jurisdictions. You are advised to exercise caution in relation to this document. If you are in any doubt about any of the contents of this document, you should obtain independent professional advice. For materials which constitute product advertisement under the Securities and Futures Act and the Financial Advisers Act, this advertisement has not been reviewed by the Monetary Authority of Singapore. JPMorgan Chase Bank, N.A., a national banking association chartered under the laws of the United States, and as a body corporate, its shareholder's liability is limited.

With respect to countries in **Latin America**, the distribution of this material may be restricted in certain jurisdictions.

Issued in **Australia** by **JPMorgan Chase Bank, N.A. (ABN 43 074 112 011/AFS Licence No: 238367) and J.P. Morgan Securities LLC (ARBN 109293610)**.

References to "J.P. Morgan" are to JPM, its subsidiaries and affiliates worldwide. "J.P. Morgan Private Bank" is the brand name for the private banking business conducted by JPM. This material is intended for your personal use and should not be circulated to or used by any other person, or duplicated for non-personal use, without our permission. If you have any questions or no longer wish to receive these communications, please contact your J.P. Morgan team.