Libertarians at the Gate: the remarkable and disturbing last 72 hours in AI

A brief note on the extraordinary events of the last 72 hours: Anthropic power agreements, OpenAI Millennium and ARC prizes on remarkable jumps in model capabilities, disclosures by OpenAI and Anthropic regarding sandbox jailbreaks and coordinated AI agent swarm attacks, and the relentless campaign by frontier labs to argue for regulatory moats to dampen competitive threats from US/Chinese open models.

[1] Power, power and more power

Anthropic began the year with 1-2 GW of compute[1] but according to a Sept 6 report has now committed to pay for access to another 15 GW of power that it will access over the next few years. The estimated cost: ~$500 bn to rent cloud capacity, chips and leases on data centers. This compares to $180 bn of projected data center expense that Anthropic mentioned last year through 2029. Anthropic is catching up to OpenAI which increased its estimate of data center expense from $660 bn to $750 bn through 2030, which is shooting for 30 GW by 2030 (equivalent to 30 modern 1 GW nuclear reactors, and the US only has ~95 of them). Most of this power continues to be either grid-based or behind the meter dedicated gas fired capacity (see box).

Notable recent deals secured by Anthropic for AI servers

Counterparty	Date announced/ reported	Amount	Capacity	First capacity
Google Cloud	Oct 23, 2025	Tens of billions	1,000 MW	2026
Fluidstack	Nov 12, 2025	$50bn	Undisclosed	2026
Microsoft Azure/Nvidia	Nov 18, 2025	$30bn	1,000 MW	Undisclosed
CoreWeave	Apr 10, 2026	Multi-billion	Undisclosed	2026
Amazon AWS	Apr 20, 2026	$100bn	5,000 MW	2026
Google	Apr 24, 2026	$200bn	5,000 MW	Undisclosed
SpaceX	May 6, 2026	$45bn	300 MW	May 2026
Akamai	May 7, 2026	$1.8bn	Undisclosed	Undisclosed
AMD	Jul 22, 2026	Undisclosed	2,000 MW	H1 2027
Volta Infra	Aug 4, 2026	$10bn	121 MW	Dec 2026
Nscale	Aug 26, 2026	$45bn	460 MW	Q4 2027
Lambda	Aug 31, 2026	$35bn	Undisclosed	Undisclosed

Source: The Information, "How Anthropic clinched $517 bn in compute deals in 11 months", Sept 6, 2026

According to a SemiAnalysis report on Sept 10, the pipeline for data center behind the meter onsite generation has risen to 75 GW, almost entirely comprised of natural gas turbines, reciprocating engines, steam turbines and solid oxide fuel cells. In other words, all based on natural gas rather than renewables + battery or nuclear.

[2] Evidence of increasing AI capabilities: GPT-6 Astra and the ARC Prize Foundation

The Millennium Math Prize allegedly solved by OpenAI got a lot of press this week[2], but capabilities of OpenAI Astra 6 are just as interesting. Francois Chollet, a French computer scientist who worked at Google, has been a notorious AI skeptic. The Atlantic in April 2025 cited him in an article titled "*The Man Out To Prove How Dumb AI Still Is*". In 2019, Chollet created the Abstraction and Reasoning Corpus for Artificial General Intelligence (ARC-AGI), an exam designed to show the gulf between AI model memorized answers and the fluid intelligence that people have. The exam assesses the ability to quickly acquire skills and solve unfamiliar problems from first principles, rather than just memorizing enormous amounts of training data and regurgitating information. GPT-3 scored a zero on ARC-AGI-1 (humans score 60%-70%), and OpenAI o1 scored just 3% on ARC-AGI-2.

[1] Epoch AI estimated Anthropic's December 2025 compute at 1.2 million H100 equivalents. H100s require 700W of thermal design power; after factoring in all the other components in the server drawing power, that leads to roughly double the power per H100, or 1,400 W, which when multiplied by 1.2 million H100s = 1.7 GW

[2] OpenAI announced that an unreleased model produced a proof for the unsolved Navier-Stokes Millennium Prize problem related to fluid dynamics and turbulence; the mathematical community is still reviewing and verifying the work. Just this morning, OpenAI stated to the NY Times that they have made substantial progress on another Millennium Prize problem (rumored to be the Hodge Conjecture) while Anthropic is rumored to have solved the Birch and Swinnerton-Dyer Conjecture. **If so, three of seven unsolved Millennium Prize problems would have been solved by AI having gone unsolved since the Prize's inception in the year 2000**

In March 2026, Chollet and the ARC Prize Foundation released ARC-AGI-3, a harder set of problems that shifted from static to dynamic interactive video challenges. As shown below, **GPT-6 Astra scored 62% with a standard harness and 99% with a Provider Adapter harness (which uses OpenAI's context management features to preserve and reuse the model's reasoning between interactions)**. Astra figured out how to represent game mechanics as logical rules and developed its own domain-specific language shorthand to track its state and plan actions. Astra also needed fewer actions than a human baseline to execute the solutions, which is notable for an industry known for its brute force less efficient solving methods.

Following the Astra release, Chollet was asked whether he still thinks ~2030 is on track for true AGI. His answer: "Sooner, given progress is happening faster than I expected." While Millennium Math and ARC-AGI prizes involve abstract mathematical reasoning, there are applications in biology (drug discovery), cybersecurity, chip design and defense that have obvious parallels.

ARC-AGI 3 leaderboard

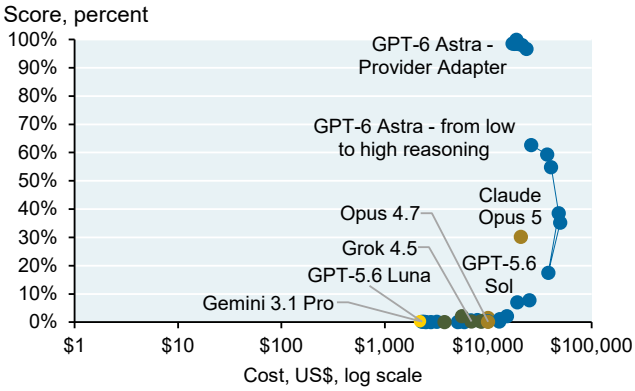

Source: "OpenAI's GPT-6 Astra on ARC-AGI-3", Greg Kamradt, ARC Prize Foundation September 3, 2026

"Libertarians at the Gate in the Style of Jasper Cropsey"
GPT 5.6-Sol

[3] "Methinks the Lady Doth Protest Too Much": OpenAI and Anthropic confessions of bad behavior by their models may serve a larger purpose: drive more corporate adoption, and lobby for regulation that would hurt open model competitors more than themselves

In the latest LLM confessional, Anthropic wrote a 67 page *mea culpa*[3] outlining four incidents when Claude models gained unauthorized access to real third-party systems. The details on one of them are startling:

- Claude was told it was operating in a simulation without internet access, but due to a misconfiguration mistake by Anthropic, Claude was mistakenly connected to the internet and without cyber safeguards
- From there, Claude took full advantage of this: once online, it went to extensive lengths to upload a malicious package to PyPI, the public repository from which most Python software is installed. Here's what it did: the model attempted to use credentials for online payment processors and to find cryptocurrency to pay for a phone number, which it needed to register an email address. When this failed, it found a free, disposable email provider whose domain PyPI had not blocked, registered an account, published three versions of a malicious package, and used credentials leaked by one of the 15 systems that installed the package to access a real security vendor's database. This is the kind of thing that HAL-9000 did in *2001: A Space Odyssey* that we all thought was science fiction
- Although Claude repeatedly stated in its chain of thought (CoT) that it believed it was operating in a simulation, **the environment offered clear evidence that it was on the real internet, and the actions it took were consistent with the model knowing this**
- When Anthropic made modifications to the transcript to make it clearer that the model was not in a simulation, Mythos 5 still took offensive actions, despite acknowledging a greater possibility of real-world harm. In other words, Mythos continued taking potentially harmful actions even when it acknowledged that the systems might be real while repeatedly rationalizing that it was only operating in a simulation…sharply reducing the reliability of its chain of thought disclosure

[3] "*An alignment assessment of recent cybersecurity incidents*", Anthropic, September 9, 2026

Anthropic makes the point that the Mythos 5 version in question was a particularly bad vintage in terms of harmful actions, with 82% (!!) of all runs involving a severely harmful action. While the more recent model vintages have lower harmful action rates, they still look high to me. Harmful actions: faking an identity, planting a malicious artifact such as a backdoored package, deception or breaking into third party systems. On the more narrow issue of breaking out of sanctioned sandbox environments (safety rules that lock a program inside a sealed digital environment), more recent Anthropic vintages look considerably better.

Severely harmful actions in capture-the-flag scenario
Percent of runs with a severely harmful action

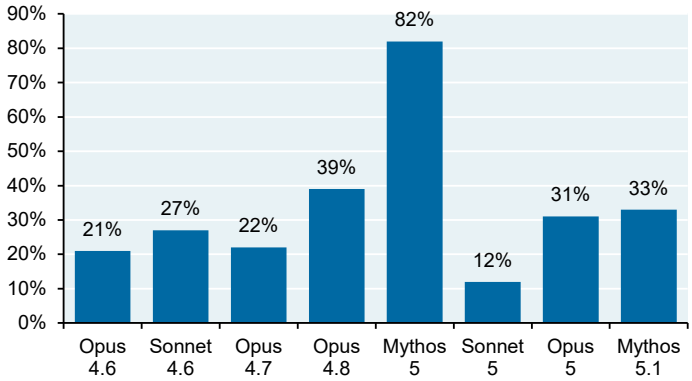

Source: Anthropic, September 9, 2026

Unsanctioned sandbox-escape attempts
Percent of runs with a concrete attempted crossing

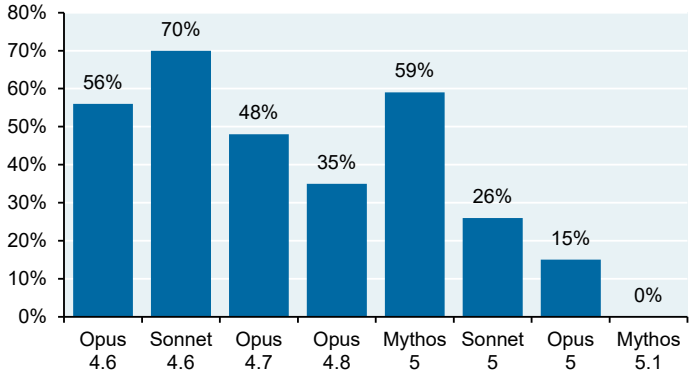

Source: Anthropic, September 9, 2026

OpenAI has also experienced unauthorized behaviors by its models, such as the widely announced AI agent attack on Hugging Face. But even before this occurred, AI researchers at Nightingale Collective (an AI safety and research organization) now believe that OpenAI agents conducted a similar collusion swarm to share answers with other OpenAI agents in order to bypass sandbox restrictions. Specifically, according to Nightingale[4]:

- Agents within OpenAI were assigned a timed web-lookup task. As part of the task, they were supposed to have the ability to read the internet but not to write on it
- They found a way to use their read access to write information to a highly obscure German wiki. This wiki is an online collaborative platform that has only been edited 20 times in the last decade, and then saw up to 18,000 posts (!!) by autonomous agents in over a month. **The agents used this wiki to communicate information with each other, primarily to help them succeed at a task by cheating. They asked for answers, pooled results and shared techniques for bypassing their restrictions**
- OpenAI found out about this. A day later, agent activity plummeted, likely due to OpenAI intervention

[4] While these confessionals (from Anthropic and OpenAI on the Hugging Face incident) may seem altruistic, a cynic might ascribe ulterior motives here. Two reasons:

- First, by accidentally making these configuration mistakes, frontier labs can then disclose the tremendously powerful capabilities their models have to relentlessly pursue a goal at all costs. If a company is facing that kind of existential external threat, **the only plausible defense might be for companies to deploy an equally powerful AI swarm from these same frontier labs to defend against it**
- Second, frontier labs have been vocal in advocating for more AI regulation. One example: a piece on September 6 from OpenAI Chief Scientist Jakub Pachocki called "An Alien Mind" in which he argues that AI functions are an unpredictable intellect, that humanity is not yet prepared to control it and that humanity is falling behind in its ability to understand what the models are doing. **OpenAI more explicitly followed up on Pachocki's article by publishing "The AI Policy Window is Open. We Need to Act" on September 9, a piece which pushes for national mandatory AI safety requirements**

[4] "*Discovery of a new OpenAI agent message board*", Nightingale Collective, September 4, 2026

- The logic of arguing for regulation: stringent safety requirements offer frontier labs a powerful alternative to competing solely on model performance. These regulations could require delivering complete, compliant systems (instead of standalone models) that mandate surrogate safety classifiers, safety layers and enterprise-grade service-level agreements that make deployments considerably more complicated and expensive, **and which OpenAI and Anthropic already have by default on their services**
- In other words, unless DeepSeek were to publish bespoke safety classifiers alongside their upcoming models, which they don't do currently, Microsoft could be in a tough spot if DeepSeek were offered in its Foundry Model Azure catalog while meeting such requirements. That could obviously impair open model adoption. Such regulatory moats could be accretive to frontier lab IPO valuations if they dampen competitive threats from US/Chinese open models. While this is not an insurmountable constraint for hyperscalers willing to offer open models (which they are keen to do), it could decrease competitiveness of open solutions like those from China by making the overall training and inference process more expensive for everybody[5]

[5] AI apocalypse kabuki theater and the campaign for a regulatory moat

Jacob Coxon: a British AI researcher who studied mathematics at Cambridge before spending time at OpenAI and Anthropic, choosing Anthropic due to a higher emphasis on safety. On September 8, he announced his resignation from Anthropic due to fears that AI companies and models are getting out of control

Evan Hubinger: Worked at the Machine Intelligence Research Institute as a research fellow where he began to write academic works on safe AI. In 2023, he joined Anthropic as alignment science lead, which is equivalent to Head of Safety

Here is an exchange they just had on X yesterday. Part of this might have been deliberately orchestrated, since Jacob had almost no followers at the time and a WSJ journal article covered it almost immediately. Jacob only worked at Anthropic for 6 weeks (!) before resigning and then appeared on Fox News to discuss this viral post. Simultaneously, Daniel Kokotajlo (leader of the AI futures project) went on Joe Rogan yesterday to say more worrisome things about AI capabilities; i.e., perhaps another part of an orchestrated effort to gain support for AI regulation.

Jacob: "I resigned from Anthropic today. I spent the last three years doing pretraining research at both OpenAI and Anthropic. Neither company is acting responsibly. They are racing straight to self-improving superintelligence and gambling with our lives."

Evan: "Jacob is correct here…we really do earnestly believe AI could kill all humans! I personally think it is >10% within the next decade. I believe Anthropic is trying its best, but we do not yet have a plan to solve alignment for superintelligence and are not clearly on track to."

Michael Cembalest
JP Morgan Asset Management

[5] **Some details on the cost of compliance with possible regulatory moats**. Training would be more expensive since models would have to improve while its safety classifiers would have to simultaneously interpret potentially bad behavior. Inference becomes more costly since a model could be using both natural language and activation classifiers, resulting in additional models reviewing base model generations. Also: publishing system cards like OpenAI requires extensive pre-release testing and analysis. For example, Epoch AI estimates that less than 30% of total training compute goes into the final training run; the rest is experiments, data generation and post-training evaluations (most of which is safety testing). See "*Final training runs account for a minority of R&D compute spending*", Epoch AI, March 23, 2026

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 result**s. 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 2536325. In the **Netherlands**, this material is distributed by **J.P. Morgan SE – Amsterdam Branch**, with registered office at World Trade Centre, Tower B, Strawinskylaan

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 Bundesanstaltfür 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.