

Internet/e-Commerce

BofA Frontier AI Data Tracker: Model rankings, usage data, and token/GPU costs

Industry Overview

Key signals shaping Internet AI stocks sentiment

Frontier AI capabilities & pace of AI investments remain key drivers of sentiment and valuation across large-cap Internet stocks. We are launching the BofA Frontier AI Tracker that monitors key indicators across the AI value stack, including frontier model intelligence rankings, usage trends, token pricing, GPU rental rates & memory costs.

Intelligence: Claude leading, followed by GPT 5.6 Sol

Model intelligence remains a key driver of competitive positioning & AI sentiment. According to Artificial Analysis, Anthropic currently leads the Intelligence Index with Claude Opus 5 ranking first, followed by Claude Fable 5 and OpenAI's GPT-5.6 Sol. Meta's MuseSpark 1.2 (xhigh) was ranked 7th and Gemini 3.6 Flash was ranked 13th, suggesting Meta (Buy) is making progress with its AI labs initiatives.

Model usage: A mixed bag based on platform end market

Vercel data suggests DeepSeek gained usage share m/m and leads at 30% share, while Anthropic usage share was second (25%) and was first for spend (65%), and OpenAI usage share was at 16% and spend share at 11%. OpenRouter data suggests open Chinese models are leading for usage, with Xiaomi's MiMo-V2.5 #1 at 32.8tr tokens MTD, followed by DeepSeek V4 Flash at 26.4tr. For closed models, OpenAI's GPT-5.6 Luna lead for usage with 8.6tr tokens, followed by Claude Opus 4.8 at 4.8tr. We note this data does not incorporate model usage directly on leading hyperscaler platforms.

Token pricing falls m/m after 3 month run up

As of August 13, AI Token price Index (TPI) was down 9% m/m (still up 87% y/y) to \$2.21 and LLM Token Expenditure Index was down 27% m/m to \$1.15. OpenAI price cuts, including an 80% reduction for GPT-5.6 Luna and a 20% reduction for GPT-5.6 Terra, as well as Gemini 3.6 Flash replacing Gemini 3.5 Flash in the TPI index (reportedly delivering 17% better cost efficiency) likely contributed to m/m declines.

GPU rentals and memory pricing mixed m/m

According to Silicon Data, B200 rental prices were down 2% m/m (+7% y/y) to \$5.63 in Aug. H100 were up 2% m/m (+33% y/y) to \$2.77 and A100 were flat m/m (+17% y/y) to \$1.65. In Aug, DRAM memory prices (DDR5 – 5600/6400 24GB 3Gx8) were up 8% m/m (+483% y/y) and NAND memory prices (QLC Flash 1TB) were up flat m/m (+432% y/y). Data suggests server useful lives remain healthy (A100 shipments started in 2020).

Healthy ecosystem, token costs worth watching

We think August data suggests: 1) Anthropic and OpenAI are setting the pace at the LLM frontier (Meta's Watermelon and Google's Gemini 4 launches are potential catalysts), 2) Lower-cost open models had modest recent usage share gains (Which may have prompted Meta's launch of open-weight models), 3) Token pricing was very strong through July, which likely aided cloud margins, but price competition increased in August, 4) GPU rental pricing remains firm, suggesting AI infra demand remains healthy & server useful lives are likely meeting/exceeding expectations, and 5) no NT relief from memory cost inflation. We think hyperscaler model launches and token pricing trends and announcements could be a key mega cap Internet sentiment drivers going forward.

BofA Securities does and seeks to do business with issuers covered in its research reports. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of this report. Investors should consider this report as only a single factor in making their investment decision. Refer to important disclosures on page 12 to 16. Analyst Certification on page 11. Price Objective Basis/Risk on page 11.

13009068

Timestamp: 17 August 2026 04:30AM EDT

17 August 2026

Equity

United States

Justin Post

Research Analyst
BofAS
+1 415 676 3547
justin.post@bofa.com

Nitin Bansal, CFA

Research Analyst
BofAS
+1 415 676 3551
nbansal7@bofa.com

MTD: Month to Date

GPU: Graphic Processing Unit

API: Application Programming Interface

TPI: Token price Index

DRAM: Dynamic random-access memory

BofA Frontier AI Tracker - Aug Overview

AI Index Dashboard

Exhibit 1: Anthropic models leading in frontier AI capabilities
BofA Frontier AI Tracker (August 2026 Data Overview)

Frontier AI Capabilities	
Top Model by Intelligence Score	Anthropic Opus 5 (max)
Top Model by Agentic Score	Anthropic Opus 5 (max)
Top Model by Coding Agentic Score	Anthropic Opus 5 (x high)
Model Usage	
Highest usage on OpenRouter (Closed model)	OpenAI GPT 5.6 Luna (8.6tr tokens MTD)
Highest usage on OpenRouter (Open model)	Xiaomi's MiMo-V2.5 (32.8tr tokens MTD)
Highest token volume share (Vercel AI Platform)	DeepSeek
Highest spend share (Vercel AI Platform)	Anthropic
Token Pricing	
AI Token Pricing Index	Down 9% m/m (up 87% y/y)
LLM Token Expenditure Index	Down 27% m/m
GPU Rental Prices	
B200 (Blackwell)	Up 2% m/m (up 7% y/y)
H100 (Hopper)	Up 2% m/m (up 33% y/y)
A100 (Ampere)	Flat m/m (up 17% y/y)
Memory Prices	
DRAM	Up 8% m/m (up 483% y/y)
NAND	Flat m/m (up 432% y/y)

Source: Artificial Analysis, Bloomberg, Silicon Data, OpenRouter


BofA GLOBAL RESEARCH

Frontier Model Intelligence Ranking

Intelligence Index

According to Artificial Analysis, Claude Opus 5 (max) was ranked highest in intelligence index, followed by Claude Fable 5 (w/fallback), and GPT-5.6 Sol (max). Meta’s MuseSpark 1.2 (xhigh) was ranked 7th and Gemini 3.6 Flash was ranked 13th.

Exhibit 2: Claude Opus 5 (max) was ranked highest in intelligence index
Intelligence Index - Top 15 Models



Source: Artificial Analysis

BofA GLOBAL RESEARCH




Agentic Index

According to Artificial Analysis, Claude Opus 5 (max) was ranked highest in Agentic index, followed by GPT-5.6 Sol (max) and Qwen 3.8 Max. Meta's MuseSpark 1.2 (xhigh) was ranked 9th, Grok 4.5 (high) was ranked 10th and Gemini 3.6 Flash was ranked 14th.

Exhibit 3: Claude Opus 5 (max) was ranked highest in Agentic index

Agentic Index Top - 15 Models



Source: Artificial Analysis


BofA GLOBAL RESEARCH

Coding Agent Index

According to Artificial Analysis, Claude Opus 5 (xhigh) was ranked highest in Coding Agent index, followed by GPT 5.6 (Sol max), and Claude Fable 5 (max). Grok 4.5 (high) was ranked 5th and Meta's MuseSpark 1.1 (xhigh) was ranked 11th.

Exhibit 4: Claude Opus 5 (xhigh) was ranked highest in Coding Agent index

Coding Agent Benchmark - Top 15 Models



Source: Artificial Analysis

BofA GLOBAL RESEARCH

Top Models by Intelligence Score by BenchLM

According to BenchLM, Anthropic's Claude Mythos 5, Claude Opus 5 and Claude Fable 5 lead the overall intelligence rankings, each with a score of ~83, followed by OpenAI's GPT-5.6 Sol at ~82 and Kimi K3 at ~80. Meta's Muse Spark 1.2 ranked 8th with a score of ~77, while Google's Gemini 3.6 Flash ranked 10th with a score of ~75.

Exhibit 5: Anthropic Claude models lead in the intelligence score

Top 20 Models by Intelligence by BenchLM

Rank	Model			Pricing		Intelligence Score					
	Model	Creator	Type	Input Price	Output Price	Overall Score	Coding	Agentic	Reasoning	Knowledge	Math
1	Claude Mythos 5	Anthropic	Proprietary	\$10	\$50	83	80	76	-	94	-
2	Claude Opus 5	Anthropic	Proprietary	\$5	\$25	83	78	81	91	93	-
3	Claude Fable 5	Anthropic	Proprietary	\$10	\$50	83	80	75	-	67	-
4	GPT-5.6 Sol	OpenAI	Proprietary	\$5	\$30	82	79	68	93	83	97
5	Kimi K3	Moonshot AI	Pending	\$3	\$15	80	78	74	-	85	-
6	Qwen3.8 Max	Alibaba	Proprietary	-	-	80	67	79	96	67	-
7	Claude Opus 4.8	Anthropic	Proprietary	\$5	\$25	77	72	61	76	88	67
8	Muse Spark 1.1	Meta	Proprietary	-	-	77	64	65	-	93	-
9	Grok 4.5	xAI	Proprietary	\$2	\$6	76	58	62	60	67	-
10	Gemini 3.6 Flash	Google	Proprietary	\$2	\$8	75	65	46	-	68	-
11	GPT-5.4	OpenAI	Proprietary	\$3	\$15	73	59	57	78	77	66
12	GPT-5.5	OpenAI	Proprietary	\$5	\$30	73	71	64	87	79	71
13	GPT-5.6 Terra	OpenAI	Proprietary	\$2	\$12	72	66	63	86	82	97
14	Claude Opus 4.7 (Adaptive)	Anthropic	Proprietary	\$5	\$25	72	65	64	79	83	-
15	Claude Opus 4.7	Anthropic	Proprietary	\$5	\$25	72	68	57	-	65	62
16	Qwen3.7 Max	Alibaba	Proprietary	-	-	72	54	39	87	72	82
17	Muse Spark	Meta	Proprietary	-	-	71	63	65	52	75	56
18	MiMo-V2.5-Pro	Xiaomi	Proprietary	-	-	69	55	42	-	71	-
19	MiniMax M3	MiniMax	Open Weight	\$0.3	\$1.2	69	48	40	-	64	-
20	Hy3	Tencent	Open Weight	\$0.0	\$0.0	68	63	61	-	-	-

Source: benchlm.ai

BofA GLOBAL RESEARCH


Model Usage / Adoption

Token Usage (OpenRouter)

Model usage trends from OpenRouter suggest strong adoption of lower-cost open Chinese models, with Xiaomi's MiMo-V2.5 leading at 32.8tr tokens MTD, followed by DeepSeek V4 Flash at 26.4tr and Tencent's Hy3 at 20.1tr.

Exhibit 6: Xiaomi MiMo-V2.5 and Deepseek V4 Flash lead in token usage

Top 15 Models by Token Usage (August MTD, trillion token)



Source: Openrouter.ai

BofA GLOBAL RESEARCH




Closed Models Token Usage (OpenRouter)

Among closed models GPT 5.6 Luna leads in usage with 8.6tr token, followed by Claude Opus 4.8 with 4.8tr token and Claude Sonnet 5 with 4.5tr tokens.

Exhibit 7: GPT 5.6 Luna and Claude Opus 4.8 lead in token usage

Top 15 Closed Models by Token Usage (August MTD, trillion token)



Source: Openrouter.ai

BofA GLOBAL RESEARCH

How to interpret OpenRouter model usage data?

OpenRouter is an AI model aggregation and routing platform that provides developers with access to multiple LLMs through a single API. OpenRouter LLM usage rankings offer a real-time view of model adoption based on token volume routed through the OpenRouter platform, making the data a useful indicator of developer mindshare, cost-performance trade-offs and commercial traction within the frontier AI ecosystem.

However, we note OpenRouter data should not be interpreted as total model usage across the broader AI market. The dataset captures usage from developers and applications that route requests through OpenRouter but does not include workloads that go directly through provider-specific APIs or cloud platforms such as OpenAI API, Anthropic API, Google Vertex AI/Gemini API, AWS Bedrock or Azure OpenAI. It also excludes usage from enterprise private deployments, self-hosted inference stacks, hyperscaler internal workloads and company-owned data centers unless those workloads are routed through OpenRouter.


Token Volume Share on Vercel AI Platform

DeepSeek led the average MTD token usage share on the Vercel AI platform at 29.7% (vs 26.0% in July), followed by Anthropic with average MTD token market share at 24.7% (vs. 30.1% in July), OpenAI's 16.3% (vs. 12.3% in July) and Google 5.1% (vs. 11.5% in July). The most notable share gains were OpenAI (+4.1pt) and DeepSeek (+3.7pt), followed by StepFun (+3.3pt), while the largest declines were Google (down 6.4pt), Anthropic (down 5.4pt), followed by MiniMax (down 1.3pt).



Exhibit 8: DeepSeek lead AI Token Usage on Vercel, capturing approximately 25% share MTD in August (vs 30.1% in July)

Daily Token Volume Share on Vercel AI Platform (%)



Source: Vercel.com

BofA GLOBAL RESEARCH


Vercel's daily token share data captures each AI lab's share of the total tokens routed through the Vercel AI Gateway, regardless of the price paid per token. Similar to OpenRouter, the data reflects only activity routed through the Vercel gateway and excludes AI inference activity handled directly through the respective labs' APIs or through third-party cloud providers.

Spend Share on Vercel AI Platform

Anthropic led the average MTD spend share on the Vercel AI platform at 64.8% (same as July), followed by OpenAI at 11.2% (vs. 13.4% in July), Google 7.8% (vs. 10.5% in July) and Moonshot.ai at 6.4% (vs 2.1% in July). The most notable spend-share gains were Moonshotai (+4.3pt), DeepSeek (+0.5pt), and StepFun (+0.4pt), while the largest declines were Google (-2.8pt), OpenAI (-2.1pt), and Alibaba (-0.2pt). Overall, Opensource models gained some share m/m.

Exhibit 9: Anthropic lead AI model spend on Vercel, capturing approximately 65% share

Dily Spend Share on Vercel AI Platform (%)



Source: Vercel.com

BofA GLOBAL RESEARCH



Vercel's daily spend share represents each AI lab's share of total AI Gateway spend. Vercel calculates spend using market-rate pricing based on providers' published list prices, providing a normalized measure across teams, including those using their own API keys. The metric therefore reflects both the volume of tokens routed to each lab and the market-rate price of those tokens. The data covers only activity routed through Vercel AI Gateway and excludes AI inference handled directly through the respective labs' APIs or through other cloud providers.


Token Pricing

AI Token Price Index- Prices moderate after big run

As of August 13, the AI Token price index indicated token prices were \$2.21, down 6% m/m (up 87% y/y). Potential drivers of the m/m decline in August: 1) OpenAI lowering prices of GPT 5.6 Luna by 80% and GPT 5.6 Terra by 20%, and 2) Gemini 3.6 Flash replacing Gemini 3.5 Flash delivering 17% cost efficiency with same model capability.

Exhibit 10: In August, Token price index declined 9% m/m (+87% y/y).

Token Price Index – Avg cost per 1mn tokens, aggregated across leading providers



Source: Token Price Index


BofA.G. GLOBAL RESEARCH

The Token Price Index (TPI) tracks the average cost of accessing leading frontier AI models across major providers, including Anthropic, OpenAI, Google, xAI, DeepSeek, Mistral, and Alibaba's Qwen. The index is designed to measure how frontier AI pricing evolves over time as new models are introduced and providers adjust pricing. TPI serves as a barometer for the economics of frontier AI. Index levels can rise when higher-performing premium models gain share or enter the market and fall when lower-cost alternatives or pricing reductions improve the overall cost efficiency of accessing advanced AI capabilities.

LLM Token Expenditure Index (Silicon Data)

Exhibit 11: In August, LLM Token Expenditure Index declined 27% m/m to \$1.15

Silicon Data LLM Token Expenditure Index



Source: Bloomberg, Silicon Data LLM Token Expenditure Index

BofA GLOBAL RESEARCH

In August, LLM Token Expenditure Index declined 27% m/m to \$1.15 (vs \$1.57 in July).


The Silicon Data LLM Token Expenditure Index (SDLLMTK) is a daily benchmark that tracks the effective cost of AI inference across the LLM ecosystem, expressed as a normalized blended rate in USD per million tokens. The index is calculated using pricing observations across more than 20 leading models and is weighted towards where inference usage is concentrated, providing a market-based view of what users are actually paying for AI inference. The index incorporates data from frontier API providers, open-weight inference platforms, dedicated-instance markets, and self-hosted deployments, offering a broader measure of AI inference economics.

Cost per Task (Artificial Analysis)

According to Artificial Analysis, Claude Fable 5 (w/fallback) has the highest cost per task among leading frontier AI models at \$3.14 per Intelligence Index task, followed by Claude Opus 5 at \$2.34 and GPT-5.6 Sol at \$1.23. Gemini 3.6 Flash was 5th at \$0.56 and Grok 4.5 (high) was ranked 8th at \$0.36.

Exhibit 12: Claude Fable 5 (w/fallback) has the highest cost per task among leading AI models

Weighted average Dollar Cost per Intelligence Index task, by model



Source: Artificial Analysis

BofA GLOBAL RESEARCH



Artificial Analysis' Cost per Intelligence Index Task measures the effective cost of delivering a unit of AI capability, adjusting for model quality as captured by the Artificial Analysis Intelligence Index. The metric estimates the weighted average cost (USD) of completing a standardized AI task by incorporating all relevant inference costs, including input, cache, reasoning, and output tokens, and weighting each evaluation according to its contribution to the Intelligence Index. Lower values indicate greater cost efficiency, meaning a model can achieve a given level of intelligence at a lower inference cost.


GPU Pricing

Silicon Data GPU Rental Index

According to Silicon Data GPU Rental index, in Aug, B200 (Blackwell) rental prices were down 2% m/m (+7% y/y) to \$5.63, H100 (Hopper) rental prices were up 2% m/m (+33% y/y) to \$2.77 and A100 (Ampere) rental prices were flat m/m (+17% y/y) to \$1.65.

Exhibit 13: In Aug, B200 rental prices were down 2% m/m, H100 +2% m/m and A100 flat m/m

Silicon Data GPU Rental Index (\$ rental price/hour)



Source: Bloomberg, Silicon Data Indices

BofA GLOBAL RESEARCH

The Silicon Data GPU Rental Index tracks the market rental price of AI GPUs. Prices are expressed as a standardized \$USD per GPU-hour and are derived from data across neo-cloud providers, hyperscalers, colocation markets, brokered cluster sales, and private rental platforms.


Memory Pricing

DRAM spot prices

According to Bloomberg, DRAM memory prices (DDR5 – 5600/6400 24GB 3Gx8) were up ~8% m/m (+483% y/y) in August.

Exhibit 14: DRAM memory prices were up ~8% m/m in August

Spot Price of DRAM DDR5 – 5600/6400 24GB 3Gx8



Source: Bloomberg, inSpectrum Tech


BofA GLOBAL RESEARCH

NAND spot prices

According to Bloomberg, NAND memory prices (QLC Flash 1TB) were up flat m/m (+432% y/y) in August.

Exhibit 15: NAND memory prices were up flat m/m in August

Spot Price of NAND DDR5 - QLC Flash 1TB



Source: Bloomberg, inSpectrum Tech

BofA GLOBAL RESEARCH

Exhibit 16: Stocks mentioned

Prices and ratings for stocks mentioned in this report

BofA Ticker	Bloomberg ticker	Company name	Price	Rating
GOOGL	GOOGL US	Alphabet A	US\$ 345.9	B-1-7
GOOG	GOOG US	Alphabet C	US\$ 343.54	B-1-7
AMZN	AMZN US	Amazon.com	US\$ 262.65	B-1-9



Exhibit 16: Stocks mentioned

Prices and ratings for stocks mentioned in this report

BofA Ticker	Bloomberg ticker	Company name	Price	Rating
META	META US	Meta Platforms	US\$ 589.85	B-1-7

Source: BofA Global Research

BofA GLOBAL RESEARCH

Price objective basis & risk**Alphabet (GOOGL / GOOG)**

Our price objective of \$430/\$430 is based on 27x 2027E core Google GAAP EPS plus cash per share. Alphabet has traded at an average multiple of 22x GAAP P/E over the last ten years, and we think our target multiple is reasonable vs history given expectations for double-digit revenue growth, Cloud margin expansion, and opportunity to capitalize on strong AI assets.

Downside risks are 1) loss of search traffic to AI tools from competitors, 2) LLM integration in search may take longer than expected or negatively impact search revenues, 3) revenue pressure from compliance with the EU Digital Markets Act (DMA), and 4) potential for increasing Capex and lower FCF-given AI investments.

Amazon.com (AMZN)

Our \$320 PO is based on our SOTP analysis that values AWS at 9x 2027E sales, 1P Retail at 1.0x, 3P Retail at 2.5x, and Advertising at 5.0x. Our 2027 sales multiples compare to 5.7x for SaaS comps, 1.3x for Retail comps, 2.4x for Marketplace comps, and 3.9x for Online Media comps. Our PO implies 3.6x blended 2027 P/Sales, 13x '27 EBITDA, and 31x '27 EPS.

Downside risks to our price objective are increasing competition from Cloud providers and big box retailers including Walmart, increasing AWS infrastructure costs that will pressure FCF, and Agentic AI usage that could pressure direct traffic and high-margin ad revenues. The stock has been subject to heavy volatility in the past, based on margin trends, and this volatility could increase due to economic uncertainty.

Meta Platforms Inc (META)

Our \$810 price objective is based on 24x 2027E GAAP EPS. On a total company basis, including Metaverse investments, our valuation is at a slight premium to S&P 500 given Meta's higher growth rate and AI opportunity. Historically, Meta has traded at an average premium of 3pts to S&P 500.

Downside risks are: 1) Significant dependence on digital advertising revenue, which remain highly sensitive to macroeconomic conditions, 2) Elevated spend on AI initiatives that will negatively impact margins, 3) Growing base of fixed assets could reduce cost flexibility in a macro downturn, 4) Intensifying competition for user engagement and advertising budgets from AI-native platforms, and 5) Elevated regulatory overhang from new teen safety cases or unfavorable outcomes in ongoing litigation.

Analyst Certification

I, Justin Post, hereby certify that the views expressed in this research report accurately reflect my personal views about the subject securities and issuers. I also certify that no part of my compensation was, is, or will be, directly or indirectly, related to the specific recommendations or view expressed in this research report.



US - Internet Coverage Cluster

Investment rating	Company	BofA Ticker	Bloomberg symbol	Analyst
BUY	Alphabet	GOOGL	GOOGL US	Justin Post
	Alphabet	GOOG	GOOG US	Justin Post
	Amazon.com	AMZN	AMZN US	Justin Post
	Booking Holdings Inc	BKNG	BKNG US	Justin Post
	Chewy Inc	CHWY	CHWY US	Michael McGovern
	DoorDash	DASH	DASH US	Justin Post
	Ethos	LIFE	LIFE US	Michael McGovern
	Expedia	EXPE	EXPE US	Justin Post
	LegalZoom	LZ	LZ US	Michael McGovern
	Life360	LIF	LIF US	Nitin Bansal, CFA
	Magnite, Inc.	MGNI	MGNI US	Omar Dessouky, CFA
	Meta Platforms Inc	META	META US	Justin Post
	Navan, Inc.	NAVN	NAVN US	Nafeesa Gupta
	Peloton	PTON	PTON US	Nafeesa Gupta
	Take-Two Interactive	TTWO	TTWO US	Omar Dessouky, CFA
	Tripadvisor	TRIP	TRIP US	Nafeesa Gupta
	Uber	UBER	UBER US	Justin Post
	Unity	U	U US	Omar Dessouky, CFA
	Wayfair	W	W US	Michael McGovern
	Wix.com	WIX	WIX US	Michael McGovern
NEUTRAL	Airbnb	ABNB	ABNB US	Justin Post
	AppLovin	APP	APP US	Omar Dessouky, CFA
	Carvana	CVNA	CVNA US	Michael McGovern
	Coursera Inc.	COUR	COUR US	Nafeesa Gupta
	eBay	EBAY	EBAY US	Michael McGovern
	Etsy, Inc.	ETSY	ETSY US	Michael McGovern
	Instacart	CART	CART US	Justin Post
	Match Group	MTCH	MTCH US	Omar Dessouky, CFA
	Pinterest	PINS	PINS US	Justin Post
	Reddit	RDDT	RDDT US	Justin Post
	Roblox Corp. Class A	RBLX	RBLX US	Omar Dessouky, CFA
	Snap	SNAP	SNAP US	Justin Post
	Zillow	Z	Z US	Omar Dessouky, CFA
	Zillow Group	ZG	ZG US	Omar Dessouky, CFA
UNDERPERFORM	ACV Auctions	ACVA	ACVA US	Michael McGovern
	Bending Spoons SpA	BSP	BSP US	Omar Dessouky, CFA
	Bumble	BMBL	BMBL US	Omar Dessouky, CFA
	Cardlytics, Inc	CDLX	CDLX US	Omar Dessouky, CFA
	Duolingo	DUOL	DUOL US	Omar Dessouky, CFA
	Ibotta	IBTA	IBTA US	Nitin Bansal, CFA
	Lyft, Inc.	LYFT	LYFT US	Justin Post
	StubHub Holdings	STUB	STUB US	Justin Post
	Yelp	YELP	YELP US	Nitin Bansal, CFA
RSTR	Global Business Travel Group, Inc.	GBTG	GBTG US	Nafeesa Gupta

Disclosures

Important Disclosures

