

Opinion | Matt Levine, Columnist

Crypto Brothers Front-Ran the Front-Runners

MEV spoofing, nonfungible Bitcoins and more BLS leaks.

May 16, 2024 at 12:22 PM CDT

By **Matt Levine**

Matt Levine is a Bloomberg Opinion columnist. A former investment banker at Goldman Sachs, he was a mergers and acquisitions lawyer at Wachtell, Lipton, Rosen & Katz; a clerk for the U.S. Court of Appeals for the 3rd Circuit; and an editor of Dealbreaker.

MEV spoof

Sometimes there are profitable trades, but not that often. Sometimes someone is looking to sell 100 widgets at \$20 on the American Widget Exchange, and someone else is looking to buy 100 widgets for \$21 on the National Widget Exchange. And so you can buy on the AWE and sell on the NWE and make an instant \$100 profit. But only *one* person can do that. Once you buy on the AWE and sell on the NWE, the window is closed.

Who gets to do the trade? The most intuitive answer, in most cases, is “whoever does it first.” If you see the same thing trading for two different prices on two different exchanges, and I don’t, and you send in the orders to the exchanges before me, then you do the trade and I don’t. For most human-scale trades, this makes sense and is easy to administer and feels basically fair.

Bloomberg Opinion

Musk Finally Launched His Robotaxi — Kind Of

Medicaid Cuts Will Hurt Nearly Half of America’s Kids

Foreign Terror Has a Price in US Courts

Siemens Is Battling Big Tech for AI Supremacy at Factories

The US stock market mostly does not operate at human time scales, and there these intuitions break down. In the stock market, often, one algorithm

spots trade opportunities a microsecond before another algorithm, and people get mad about that. They worry that this microsecond-scale competition between algorithms is socially wasteful, and they worry that it is unfair to people with worse algorithms. They think “yes, fine, the first person to see a trade should get to do it, *within reason*, but there’s no social benefit to slicing it that fine.” They propose things like “frequent batch auctions,” where everyone who sees the same trade within (say) a second of each other gets to compete *on price* to do it, rather than competing on *time* down to the microsecond.

Crypto, however, has a different problem. Time in electronic stock markets is maybe *too* continuous; the timeline of a stock exchange can be sliced much more finely than the human brain can comprehend. But time in crypto is oddly discrete. Time in crypto is measured in *blocks*. Intuitively, people submit orders to do transactions on a crypto blockchain, and then periodically a batch of those transactions is enshrined in the official ledger of that blockchain. In Bitcoin, a block contains roughly 10 minutes’ worth of transactions; in Ethereum, it’s seconds.

So intuitively, Ethereum transactions happen in big simultaneous bunches every 12 seconds. But the bunches can’t really be simultaneous: If there is one rare nonfungible token for sale, and two people want to buy it, only one of them can. One transaction has to be first. So the transactions within a block are *ordered*; they happen in sequence. But they are not necessarily ordered by *time*.

How does the blockchain decide which transactions to record, and in what order? In Ethereum, the answer is: with money. People who want to do transactions on the Ethereum network pay fees to execute the transactions; there is a flat base fee, but people can also bid more – a “priority fee” or “tip” – to get their transactions executed quickly. Every 12 seconds, some computer on the Ethereum network is selected to record the transactions in a block. This computer used to be called a “miner,” but in current proof-of-stake Ethereum blocks are recorded by computers called “validators.” Each block is compiled by one validator, selected more or less at random, called a “proposer”; the other validators vote to accept the block. The validators

share the transaction fees, with the block proposer getting more than the other validators.

The block proposer will naturally prioritize the transactions that pay more fees, because then it will get more money. And, again, the validators are all computers; they will be *programmed* to select the transactions that pay them the most money. And in fact there is a division of labor in modern Ethereum, where a computer called a “block builder” puts together a list of transactions that will pay the most money to the validators, and then the block proposer proposes a block with that list so it can get paid.

And so if you see 100 widget tokens trading for \$20 on one Ethereum decentralized exchange, and 100 widget tokens trading for \$21 on another Ethereum decentralized exchange, you can buy them at \$20 each and sell them for \$21 each and make \$100. But if I *also* see that arbitrage, I will *also* put in those orders. Which of our trades will execute first? Whoever pays more in execution fees. How much should I offer to pay? Oh, you know, it’s a competitive auction. So roughly \$99. The validators should get most of the money from the arbitrage. This is called “MEV,” for “miner extractable value,” though now Ethereum doesn’t have miners and the acronym stands much less informatively for “maximal extractable value.”

What are the chances that both of us discovered the same arbitrage within 12 seconds of each other? Well, pretty good, in a competitive market with a lot of arbitrageurs. But also, traditionally, when we want to do these transactions, we submit our orders *publicly* to the Ethereum network. (To the “mempool,” the name given to the place where orders reside before they are included in a block.) Everyone can see the mempool. So if you are a clever speedy arbitrageur, watching a bunch of decentralized exchanges for mispricings, you might see a mispricing and submit your arbitrage orders. And if *I* am a clever speedy *front-runner*, watching the mempool for arbitrage trades, I would submit the same trades a second or two after you, and pay more to execute them. And then I get to do the arbitrage and you don’t.

But if I am a clever speedy front-runner, I don’t even have to wait for arbitrage trades. Let’s say you are not a clever arbitrageur, but just a person

who really likes the Shiba Inu token on the Ethereum blockchain. You submit a huge order to buy SHIB on a decentralized exchange. That will, quite predictably, push up the price of SHIB. ¹ If I see that order, I can just jump ahead of it: I can pay a bit more in transaction fees to get my trades to execute first, and then I can buy the SHIB tokens ahead of you and then *sell* them to you for a profit, all within the same block.

That seems kind of rough on you? This sort of trade is sometimes called a “sandwich attack”: I see your order to buy SHIB, and I sandwich it between my order to buy SHIB (before you do, at a lower price) and my order to sell SHIB (after you do, at a higher price). And people talk about “generalized front running” as a strategy:

A generalized front-runner bot will search the mempool for profitable transactions, then copy the transaction and replace the sender address with their own, then increase their bid (gas price) to price + x to be included in a block first front running the original searcher.

I am giving a simplistic and somewhat old-fashioned description of MEV, and modern Ethereum has a whole, like, institutional structure around it. There are private mempools, where you can hide transactions from bots. There is Flashbots, “a research and development organization formed to mitigate the negative externalities posed by Maximal Extractable Value (MEV) to stateful blockchains, starting with Ethereum,” which has things like MEV-Boost, which creates “a competitive block-building market” where validators can “maximize their staking reward by selling their blockspace to an open market,” and MEV-Share, “an open-source protocol for users, wallets, and applications to internalize the MEV that their transactions create,” letting them “selectively share data about their transactions with searchers who bid to include the transactions in bundles” and get paid.

Basically, transaction data on Ethereum is informative, and the value of that information (in making profitable trades) is quantifiable, and there are cutthroat auctions among people who want to use that data to make their own trades, and the people doing the transactions that create the data can themselves get paid for their data.

There is a sort of cool purity to this. In stock markets, some people are faster than others, and can make money by trading ahead of a big order, ² and people get mad about this and think it is unfair and propose solutions. And when money changes hands for speed advantages – “payment for order flow,” “colocation” – people complain about corruption. In crypto it’s like “let’s create an efficient market in trading ahead of big orders.” I once wrote: “Rather than *solve* this concern about traditional markets, crypto *made it explicit*.” That feels almost like a general philosophy of crypto: Take the problems of traditional finance and make them, *worse*, sure, but more transparent and visible and explicit and subject to unbridled free markets.

In traditional financial markets, people sometimes get in trouble for “spoofing,” for submitting trade orders designed to fool people (algorithms) about their true trading intent. Every so often, the spoofer’s defense will be “no, all these algorithms are trying to front-run me, they are trying to use my order information to make money by trading ahead of me. So I created fake order information to fool them. I’m the hero in this story; the bots trying to front-run me are the villains.” This tends not to work.

Here, maybe, is the crypto version of that. Bloomberg’s Ava Benny-Morrison reports:

Two brothers who studied at MIT were charged with exploiting a weakness in the Ethereum blockchain and stealing \$25 million in 12 seconds, in what prosecutors called a first-of-its-kind caper.

Anton Peraire-Bueno, 24, and James Peraire-Bueno, 28, were charged by federal prosecutors in Manhattan with fraud and money laundering offenses. They are accused of carrying out the lightning-fast heist, plotted over the course of months, from their keyboards last year.

“The brothers, who studied computer science and math at one of the most prestigious universities in the world, allegedly used their specialized skills and education to tamper with and manipulate the protocols relied upon by millions of Ethereum users across the globe,” Damian Williams, the US attorney for the Southern District of New York, said in a statement.

Here are the Justice Department [press release](#) and [the indictment](#). The gist of it is that, allegedly:

1. The brothers set up “a series of Ethereum validators” so they could build blocks on the Ethereum blockchain.
2. The “targeted three Victim Traders, ... who are searchers who operate MEV Bots that specialize in cryptocurrency arbitrage trading.” They “tested a series of bait transactions” to see how the arbitrageurs operated.
3. Then, “after receiving notification that one of their 16 Validators had been selected to validate a new block,” they proposed “at least eight specific transactions (the ‘Lure Transactions’)” that they knew the arbitrageurs would like.
4. Specifically, transactions that the arbitrageurs would like to front-run: “The Victim Traders effectively bought substantial amounts of particular illiquid currencies (the frontrun trades), whose price the Victim Traders expected to increase as a result of the Lure Transactions, for approximately \$25 million of various stablecoins ... or other more liquid cryptocurrencies. The Victim Traders also included a sell transaction in each bundle, whereby the Victim Traders would sell their newly acquired cryptocurrency – immediately after the Lure Transaction – at a higher price than what they bought it for.” It was a sandwich trade: The arbitrageurs saw the Peraire-Bueno’s “lure” trades, which involved buying a bunch of illiquid crypto, and they front-ran them by buying the illiquid crypto before them and selling it after them.
5. But then they exploited a bug in Ethereum to get the full contents of the proposed block and to tamper with it by (1) “allow[ing] the Victim Traders to complete their buy transactions (i.e., their frontrun trades)” and (2) “replac[ing] the Lure Transactions with tampered transactions” in which they “sold the same illiquid cryptocurrencies that the Victim Traders had recently purchased,” and which they apparently already held. ³
6. That is, they front-ran and spoofed the front-runners: They put in a big order saying “buy a ton of [some nonsense tokens],” front-running bots

said “ooh, big bid for nonsense tokens, let’s get ahead of that,” and the front-running bots bought the nonsense tokens ahead of them. But then they tampered with the transactions to replace their own buy order with a *sell* order, *selling* the nonsense tokens – effectively, *to* the front-runners – at inflated prices for \$25 million. “In effect, the Tampered Transactions drained the particular liquidity pools of all the cryptocurrency that the Victim Traders had deposited based on their frontrun trades.”

Now. Is this ... bad? Sure, I guess. In Step 5, the Peraire-Buenos do seem to have exploited a real vulnerability in how Ethereum’s consensus mechanism worked. I do not pretend to fully understand the technical details, but here is a Flashbots postmortem of the exploit. (They sent a signed but invalid block header to the relay, which caused the relay to publish the contents of the proposed block, which they used to construct their own block exploiting those contents as the proposer of the block.) The Justice Department calls it “an alleged novel scheme by the defendants to exploit the very integrity of the Ethereum blockchain,” and that seems right: If people can reach in to alter the order of transactions to swipe \$25 million, that does sort of undermine the integrity of the blockchain.

Also they did some bad googling. Before they did the trade, Anton Peraire-Bueno allegedly “searched online for cryptocurrency exchanges with limited ‘know your customer’ protocols and ways to launder cryptocurrency, including searches for ‘how to wash crypto’ and ‘cefi exchanges with no kyc.”^[4] And after the trade, as they were allegedly “laundering the fraud proceeds from the Exploit,” James Peraire-Bueno allegedly “searched online for, among other things, ‘money laundering,’ ‘exploit,’ ‘computer fraud abuse act,’^[5] and ‘does the united states extradite to [foreign country].’” The Justice Department, reasonably, cites these searches as evidence of a guilty conscience. Also though I wonder how well they worked? Like, has anyone ever (1) acquired \$25 million of ill-gotten money, (2) googled “money laundering” and (3) successfully used to results to launder the money? On the other hand I bet googling “cefi exchanges with no kyc” works great; really exchanges should pay a lot for those keywords.

But to me what is wild about this case is that the Justice Department is bringing down the full weight of US federal criminal law to protect *Ethereum front-running bots*. The word “front-run” is right there in the indictment! The Justice Department has a long history of prosecuting people for front-running! Because it is traditionally a crime! In traditional finance, “front-running” is understood to mean violating a fiduciary duty to your customers (or employer, etc.) by trading ahead of their orders, and it is a crime. 6

But here the Justice Department is protecting the front-runners. And the trades that the Victim Traders were trying to do look, to the naked eye, bad? They were “we see you are trying to buy this illiquid cryptocurrency, so we’re going to buy it ahead of you so we can rip you off.” In Ethereum, that is pretty normal, and there is a large institutional structure to make it work in a predictable and fair-ish way. But in traditional markets it just seems bad.

We have talked a few times about Avi Eisenberg, the crypto trader arrested (and convicted) for fraud because he did some market manipulation on a decentralized crypto exchange. He absolutely did do that manipulation, but his defense – which did not work – was that the crypto exchange did not have any norms *against* it. What he did was allowed by the exchange’s code, he argued, so it was allowed. External norms like “don’t do market manipulation” did not apply.

This is controversial, and a lot of people in and out of crypto have complaints of the form “look, if decentralized finance exchanges do not want to be subject to US regulation, why does the federal government need to protect them from market manipulators?” I am sympathetic to those complaints.

But here! I mean! Ethereum and its decentralized exchanges have a market structure that is like “bots can look at your transactions and front-run them if that’s profitable.” And these guys, allegedly, front-ran the front-runners; they turned the market structure around so that *they* could get an early look at the front-running bots’ front-running transactions and front-run them instead. By hacking, sure, sure, it’s bad. But it leaves the Justice Department in the odd position of saying that the integrity of crypto front-running is important and must be defended.

Rare sats

Crudely speaking there are two sorts of crypto tokens. There are *fungible* tokens, like Bitcoin and Ether and Tether and Solana, which have approximately cash-like properties: If a thing costs two Bitcoins, it doesn't matter *which* two Bitcoins you use to pay for it; any two Bitcoins are just as good as any other two Bitcoins. And there are *nonfungible* tokens, NFTs, like Bored Ape Yacht Club, which have approximately art-object-like properties: Each Bored Ape is different, and some – based on their rarity or aesthetics or provenance – are worth much more than others.

Intuitively, fungible tokens are indistinguishable coins, while NFTs are unique images. But *really* neither of those things is true. Really they are both entries in detailed permanent transaction ledgers. An NFT, for example, *isn't* the little drawing of the monkey; the NFT is simply a numbered entry in a ledger that points to the drawing of the monkey. I once wrote that “an NFT consists of a series of numbered tokens, and the thing that makes it an NFT is that it has a different number in its tokenId field from the other tokens in its series.”

Meanwhile a fungible token like Bitcoin is *not* a *numbered* entry in a ledger; there is such a thing as “Bored Ape Number 3” but there is no such thing as “Bitcoin Number 3.” But there is a permanent immutable computer ledger containing every Bitcoin transaction, which means that you *could*, sort of, trace the ledger back to find what the third Bitcoin was and who holds it now. 7

More generally, you could, if you wanted to, treat Bitcoins as completely nonfungible. You could trace back the history of any Bitcoin, or fraction of a Bitcoin. And then you could say things like:

- “I am willing to pay more for old Bitcoins, ones that were mined in like 2009, than I will for new Bitcoins, ones that were only mined in 2024.”
- “I am willing to pay more for Bitcoins with cool provenance, like ones that were once owned by Satoshi Nakamoto or used to pay for the 2010 Bitcoin Pizza, than I will for regular old Bitcoins with no fun history.”

- “I am willing to pay *less* for Bitcoins with a troubled regulatory history, like the ones that were stolen in the Bitfinex hack, than I will for Bitcoins with a clean regulatory history.”

Bitcoins are fungible simply because it is standard convention among Bitcoin users to *treat* them as fungible, but they *are* distinguishable from one another, and if you wanted to discriminate among Bitcoins you *could*. People mostly don't, but the last thing on that list – “try not to accept Bitcoins that have been stolen or are otherwise of interest to US law enforcement” – *does* have some traction. Some Bitcoins really *are* less valuable than others.

And some are more. Here is a delightful story from Joel Khalili at Wired:

In the same way a dollar is made up of 100 cents, one bitcoin is composed of 100 million satoshis—or sats, for short. But not all sats are made equal. Those produced in the year bitcoin was created are considered vintage, like a fine wine. Other coveted sats were part of transactions made by bitcoin's inventor. Some correspond with a particular transaction milestone. These and various other properties make some sats more scarce than others—and therefore more valuable. The very rarest can sell for tens of millions of times their face value; in April, a single sat, normally worth \$0.0006, sold for \$2.1 million.

There are two interesting points here. One is simply the existence of “rare sats”: To at least some ... Bitcoin collectors? ... some satoshis are worth more than others, due to their history or provenance.

The other is that, while some Bitcoin enthusiasts value rare sats highly and will pay a lot for them, others just cheerily assume that Bitcoin is a fungible currency. And so there is a trade:

[Bill] Restey is part of a small, tight-knit band of hunters trying to root out these rare sats, which are scattered across the bitcoin network. They do this by depositing batches of bitcoin with a crypto exchange, then withdrawing the same amount—a little like depositing cash with a bank teller and immediately taking it out again from the ATM outside. The coins

they receive in return are not the same they deposited, giving them a fresh stash through which to sift. They rinse and repeat.

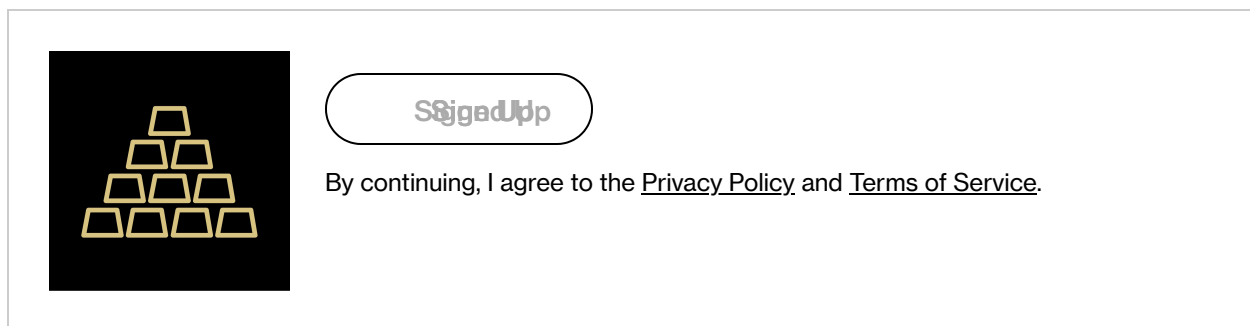
If you deposit one Bitcoin on a crypto exchange and then withdraw it, you will get one Bitcoin. But you will not get the same Bitcoin you put in. The exchange will say “ehh Bitcoins are Bitcoins, it doesn’t matter.” But you might disagree.

Super ... users ... ?

Some companies post their earnings releases every quarter on their website with a URL format like

“xyzcorp.com/investorrelations/earnings_2023q4.html.” And they put out a press release earlier saying, like, “we will release our first-quarter earnings after the close of trading at 4:05 p.m. on May 16.” And then at 4:05 p.m. they send the earnings release to the business wires. Ideally, also at 4:05 p.m., they upload the earnings to their website. But it is not *absolutely unheard-of* for a company to upload the earnings to the website a bit early, but not tell anyone until 4:05. And so it is not absolutely unheard-of for hedge fund analysts to figure out the URL format of a company’s earnings releases and point their browser to

“xyzcorp.com/investorrelations/earnings_2024q1.html” at, like, 3:45 p.m. Most of the time, this turns up nothing. Occasionally, early earnings!



I mean, I should say, this is a thing I have heard of. In 2024, it’s reasonably well-known, and many companies try not to post their earnings early on easily guessed web pages, so it’s not clear that it’s worth looking. 8
Probably some companies still do this, but you cannot really rely on systematic widespread URL guessing as a source of edge. You’re not going to get rich by typing URLs into a browser 15 minutes before earnings are due.

And so here's this:

The US Bureau of Labor Statistics inadvertently published Consumer Price Index data 30 minutes early on Wednesday, raising fresh questions about how the agency releases some of the world's most sensitive economic information.

While there were no obvious signs that the early publication moved markets, the episode is likely to prompt a close look at the dissemination of data that has implications for global asset prices and Federal Reserve policy.

“In advance of today's CPI and Real Earnings releases, BLS inadvertently loaded a subset of files to the website approximately 30 minutes prior to the release,” the agency said in a statement posted on its website Wednesday evening. ...

US stock-index futures jumped and Treasury yields fell immediately after media outlets including Bloomberg News reported the official CPI data at 8:30 a.m., with the S&P 500 Index ending the day at an all-time high. There were no sharp market movements in the half-hour stretch between the early data release and the scheduled one, suggesting the initial figures went unnoticed by investors.

Here is the statement, which is vague about what subset of files it was. Perhaps they were not the market-moving files. Or perhaps they were, and they sat on the website, but nobody was looking, because what were the chances they would have found something?

Things happen

Dow 40,000. With a BlackRock CEO, \$9 Trillion Vanguard Braces for Turbulence. LME Warehouse Queue Returns After Jump in Aluminum Orders. Starwood's \$10bn property fund taps credit line as investors pull money. Warren Buffett's Berkshire Reveals Its Mystery Stock: Chubb. Paramount: the takeover battle that could reshape Hollywood. Top-Ranking Lenders See Diminishing Recoveries in Bankruptcy. How Coinbase Is Seeking to Ease Crypto's Credit Crunch. Europe's Banks Find Breaking Up

With Russia Is Hard to Do. Billionaire Frank McCourt Plans Bid for TikTok in U.S. ‘Finfluencers’ charged for promoting unauthorised trading scheme. Milei Targets Labor Law That’s Set to Hand Banker \$10 Million Severance. Raccoon pitch invader caught with a trash can during Philadelphia Union-NYCFC match.

If you'd like to get Money Stuff in handy email form, right in your inbox, please [subscribe at this link](#). Or you can subscribe to Money Stuff and other great Bloomberg newsletters [here](#). Thanks!

1. Especially because crypto DEXes generally run mechanically so the price impact of a transaction is quite knowable. [View in article](#)

2. Do not email me about this, which is not the point. “Probabilistically make money by trading sort of in the middle of the execution of what they identify as a big order,” maybe. [View in article](#)

3. “Which the defendants already held as a result of information gathered through the bait transactions,” says the indictment. Presumably the intuition is like: You buy some illiquid crypto slowly over time, not moving the price. Then you send a big order to buy a lot of it, which gets front-runners interested. Then you pull it away. [View in article](#)

4. This trade relies on two sorts of exchanges, DeFi (decentralized finance) DEXes (decentralized exchanges) and CeFi (centralized finance) exchanges. The actual \$25 million trade here has to happen *on the blockchain*, because that is where they are spoofing and tampering with transaction ordering; DEXes are where you trade crypto on the blockchain. But then getting the \$25 million out requires a company that can wire money to your bank, and that is a centralized exchange. So he searched for CeFi exchanges to be the on -and off-ramp for the trade, which they did on decentralized exchanges. [View in article](#)

5. In fact they are charged with wire fraud, not with hacking under the CFAA, but it was probably reasonable to worry about getting charged with hacking too. But the lesson is probably that wire fraud is more general than CFAA hacking — everything is wire fraud! — so it’s the easier thing to charge. [View in article](#)

6. Of course, it is often used loosely to mean “using order book data to try to trade faster than someone else,” which is not traditional front-running and not a crime. And that’s mostly the meaning in the indictment here. [View in article](#)

7. In actual fact the first 50 Bitcoins were minted all at once, and that set of Bitcoins, [according to Blockchain.com](#), “is unredeemable, as it was omitted from the transaction database. This means any attempt to spend it would be rejected by the network. Whether this was intentional or not still remains unknown.” Also I am being a little loose with terminology here; there is not such a thing as “a Bitcoin,” but rather any amount of Bitcoin that you get is the residue of various prior transactions that can be traced back in time to when batches of Bitcoin were mined. It’s not like individual coins have passed from hand to hand, but the Bitcoins in your account reflect the history of the transactions that brought them there. [View in article](#)

8. Also there is, in the US, some risk that prosecutors might treat guessing URLs as illegal hacking? And thus, perhaps, also insider trading? [View in article](#)

This column does not necessarily reflect the opinion of the editorial board or Bloomberg LP and its owners.

How easy or hard was it to use Bloomberg.com today?

[Share feedback](#) 

©2025 Bloomberg L.P. All Rights Reserved.
