

The Trades Are Free, the Data Will Cost You

By [Matt Levine](#)

May 21, 2019, 11:05 AM CDT

Stock exchange economics

I have written before about my model of stock-exchange fees. The basic idea is that big professional traders—bank trading desks, high-frequency trading firms—mostly get to choose which exchanges they *trade at*, but they don't get to choose which exchanges they *connect to*. If you see 100 shares for sale at the same price at the same time at different exchanges, you can buy them at whichever exchange you like most, probably the one that charges you the lowest trading fee. But for various reasons—to maximize your arbitrage opportunities and minimize risks, to know what the market looks like, to comply with best-execution obligations—you pretty much need fast connections and direct data feeds at every exchange.

ADVERTISING

One implication of this is that traders will be price-sensitive to per-trade variable trading costs, but will have no choice but to pay whatever fixed data and connectivity charges the exchanges demand. Another implication is that a smart exchange will charge basically zero trading fees 1

(in order to maximize its volume and importance [2]) and extremely high fees for fast connections and data (in order to maximize its revenue from captive customers). That is not *quite* what happens—trading costs are not zero [3] —but it’s a decent intuition; “revenues from ICE’s and Nasdaq’s various data-related businesses rival what the companies make from trading.”

Here is a paper from Eric Budish, Robin Lee and John J. Shim about “A Theory of Stock Exchange Competition and Innovation.” Their model is a lot like my model:

Our model, shaped by institutional and regulatory details of the U.S. equities market, shows that under the status quo market design: (i) trading behavior across the many distinct exchanges is as if there is just a single “synthesized” exchange, as opposed to traditional platform competition; (ii) as a result, trading fees are perfectly competitive; but (iii) exchanges capture and maintain significant economic rents from the sale of “speed technology” (i.e., proprietary data feeds and co-location)—arms for the arms race.

If you think of all the U.S. stock exchanges as “a single ‘synthesized’ exchange,” then obviously trading firms (ii) will prefer to trade on whichever part of the synthesized exchange is cheapest but (iii) have to see all of them.

But that’s not Budish, Lee and Shim’s main point. Their main argument is that this situation discourages innovations that could make the stock market better:

We then use the model to examine the private and social incentives for market design innovation. We show that the market design adoption game among incumbent exchanges is not a coordination game, but rather a repeated prisoner’s dilemma. If an exchange adopts a new market design that eliminates latency arbitrage, it would win share and earn economic rents. However, imitation by other exchanges would result in an equilibrium that resembles the status quo with competitive trading fees, but now without the rents from the speed race. This means that although the social returns to market design innovation are large, the private returns are much smaller and may be negative, especially for incumbents that derive rents from the status quo.

Intuitively, if the market worked perfectly and there was no way to make money by moving information from one exchange to another faster than everyone else, then there’d be no pressure to connect to every exchange, and so exchanges wouldn’t be able to make much money selling data and connectivity. Budish is well known as an advocate of frequent batch auctions in stock markets to eliminate “the high-frequency trading arms race,” so he is perhaps more fond of “new market design that eliminates latency arbitrage” than actual market participants are, but if

you are worried that U.S. stock-market design favors fragmentation, complexity, and advantages for high-frequency traders, then this might help explain why.

Elsewhere in HFT

Here is a delightful story about a firm called Potamus Trading, which “shuttered its proprietary HFT business in January 2018, shifting its strategy and expertise to focus solely on serving asset managers as a broker-dealer”:

Potamus uses two-sided trading strategies that play to its strengths as a former market-maker. While most brokers typically use a one-sided algorithm to move the necessary amount of stocks – either buying or selling in various increments – Potamus will play both sides of the market.

As an example, if a client wanted 10,000 shares of Apple, Potamus might buy 50,000 shares and sell 40,000 shares in the process of filling the order. Doing so allows Potamus to fly under the radar.

"We find that the size of the client orders that we fill are much harder for the other HFTs out in the market to detect, because our system does not systematically send 1,000 shares every second that HFTs are built to seek out," Linnell said.

Potamus doesn't charge clients a commission fee, another uncommon practice amongst brokers. Instead, it takes a share of the profits made from the market-making that took place while handling the order. The client also gets a cut.

Intuitively, the way many high-frequency trading firms work is that they run essentially market-making-type strategies: They post bids and offers to buy and sell a stock, and try to buy at the bid and sell at the offer, capturing the small spread between them while carefully managing the risk of the stock moving against them. The goal is to make a tiny bit of money on each trade, but never to own a lot of any one stock for any long period of time. You're not getting paid to pick stocks that will go up. You're getting paid for providing liquidity, being willing to buy at the exact microsecond when someone else wants to sell.

More from

Go Ahead and Spend Your Children's Inheritance

Robinhood Is Not Gamifying Markets. It's Democratizing Them.

Give Citigroup's Bankers a Three-Month Break. Why Not?

Exxon's Suddenly Trouncing Chevron. Coincidence?

But you could tweak that, right? Instead of buying 100 shares and selling 100 shares within a second or two, you could buy 100 shares and sell, like, 90 shares, capturing some bid-ask spread, carefully managing your risk, and also, slowly, accumulating some stock. Instead of just going to buy stock, you do your market-making with a slight long bias. That way, you can (perhaps) end up both buying the stock *and* getting paid for providing liquidity—for buying when others want to sell and vice versa—rather than paying someone else to give you liquidity (by selling you the stock you want).

A similar basic idea has been kicking around in the bond markets for a while. People complain all the time about bond market liquidity, by which they mean basically that big banks no longer carry large inventories of bonds and so aren't willing to provide liquidity by buying and selling bonds, at their own risk, when investors want to trade. But big bond investors *do* carry large inventories of bonds, and might not care all *that* much about *exactly* which bonds they own: Of course they have preferences, but they own lots and lots of bonds and no one bond is all that important, so if you really want to buy one of their bonds or sell them a similar bond, then for the right price they might be willing to swap. So there is a lot of talk about the buy-side—big institutional bond investors—becoming liquidity providers, augmenting their basic business (of picking, buying and holding bonds) with a little extra income from buying bonds other people want to sell, selling bonds other people want to buy, and clipping a bid-ask spread. Potamus seems like the same basic idea, but for the weird electronic world of equity trading.

Still I feel like there will be objections? People do not like high-frequency traders. One thing that they dislike about HFTs is that they create a lot of unnecessary trading, that the stock market becomes just a game between computers rather than actually reflecting supply and demand; a strategy that lets people buy 10,000 shares by trading 50,000 shares seems a bit wasteful. Another thing that people dislike about HFT is that they think it is synonymous with “spoofing”: HFTs are constantly entering and canceling orders, and there is a popular sense that they are doing this to manipulate prices rather than for normal risk-management purposes. If you want to buy 100 shares of stock, it can make sense to put in an order to sell 10,000 shares to try to drive the price down (as long as that sell order doesn't execute), but that is illegal spoofing. If you want to buy 10,000 shares, and the way you go about it involves selling 40,000

shares in order not to overpay for the 10,000, well, that isn't spoofing, but to a casual observer it might seem to have a certain family resemblance.

Libor v. SOFR

Libor, the London interbank offered rate, is the world's main reference interest rate for floating-rate debt and interest-rate derivatives. People do not like Libor, both because there was a giant scandal in which banks were caught manipulating Libor for their own profit and vanity, but also, and more fundamentally, because Libor is basically a made-up number and so it is easy to manipulate. Libor is the interest rate on short-term unsecured loans between banks, and there are not all that many of those loans, and so when banks are asked—as they are, in setting Libor—what rate they would have to pay to borrow money unsecured from another bank, they kind of have to guess. Why wouldn't their guesses be biased?

And so there is a big worldwide movement, strongly encouraged by regulators, to replace Libor with something more real. The favored replacement in the U.S. is SOFR, the secured overnight financing rate, which is based on the market for overnight loans collateralized by Treasury securities. A big advantage of SOFR is that it is based on real transactions, so it's hard to manipulate. ⁴

Another advantage of SOFR is that it is useful as a risk-free rate: It's for very short-term loans backed by very high-quality collateral, so it doesn't really incorporate credit risk. A major use of a reference interest rate is as a risk-free rate, and in fact Libor was *also* widely used, back in the day, as a sort of stand-in for a risk-free interest rate. Then everyone remembered that banks can be really risky, Libor went up a lot, and it stopped being useful as a risk-free rate.

On the other hand, Libor was kind of useful as a reference rate for bank loans exactly *because* it reflected some of the banks' own costs: It was a way for banks to make loans on a sort of cost-plus basis, reducing their risks and thus encouraging the growth of the loan market. So replacing it with SOFR, and putting the interest-rate risk of loans on the banks, has some negatives. The “mismatch” between SOFR and the banks' own costs arguably “would make banks reluctant to lend during a period of stress.”

There are other problems with SOFR. It's an overnight rate, which is annoying; a floating-rate loan whose rate changes every three months with three-month Libor is more administratively manageable than one whose rate changes every day with one-day SOFR. There are ways—averaging, futures curves—to bootstrap SOFR into a longer-term rate, but they're not as simple and intuitive as just having a ready-made three-month Libor rate.

Anyway here is an article about how SOFR adoption is going slowly. It mentions those problems, but a key problem is something different and dumber:

The Fed's working group is also reaching out to designers of accounting and trading systems about speeding development of products that can function with the new rate. Many investors and companies can't buy or sell debt linked to SOFR because their accounting and trading systems aren't properly configured for it, said Venetia Woo, principal director of North American regulatory strategy at Accenture.

"The reality is, people want to trade and people want to hedge," Ms. Woo said. "Physically their infrastructure is not ready for it."

Imagine that everyone agreed that SOFR was a perfect replacement for Libor, superior in every respect. Theory would tell you that they'd all adopt SOFR, and they probably would. But it wouldn't be instantaneous. There'd be some period—years?—of people going around saying things like "well obviously SOFR is better but we haven't upgraded our computers yet" or "well sure but it takes time to renegotiate 200 contracts" or "yeah half of my clients haven't heard of SOFR yet." In the long run finance tends to proceed through logic and self-interest, but in the short run the computer systems cause a lot of frictions.

Uber collusion

Here's a story about some businesses that conspire together to charge higher prices to consumers:

Every night, several times a night, Uber and Lyft drivers at Reagan National Airport simultaneously turn off their ride share apps for a minute or two to trick the app into thinking there are no drivers available---creating a price surge. When the fare goes high enough, the drivers turn their apps back on and lock into the higher fare. ...

"All the airplanes we know when they land. So five minutes before, we turn all our apps off all of us at the same time. All of us we turn our apps off. They surge, \$10, \$12, sometimes \$19. Then we turn our app on. Everyone will get the surge," one driver says.

This driver explains how they organize it.

"Someone is standing by that corner. I stand by this corner and the other one stands at this corner and we say turn the app off and then go online."

We have talked before about law professor [Sanjukta Paul's discussion](#) of “the firm exemption” in antitrust law, the fact that a business firm is a legal way to coordinate actions in a way that would be illegal for private individuals.

Imagine if instead of a bunch of independent-contractor Uber drivers waiting at the airport, this was just a story about a bunch of drivers employed by Uber, charging prices set by Uber and getting paid a wage by Uber. ⁵ If that was how Uber Technologies Inc. operated—if its drivers were employees—then Uber could set whatever prices it wanted. If its strategy was like “we will raise fares from the airport when a plane lands, and lower them again when there are no planes,” that would be a perfectly legal and sensible strategy; people would complain—people really don't like surge pricing—but it would basically be fine. It's a company, it can charge whatever prices it wants, as long as people are willing to pay.

Similarly, if Uber's drivers felt that they were being underpaid—and they do; that's why they're turning the app off—then they could go to Uber and negotiate for higher wages. They could coordinate to do this, starting a union and bargaining collectively and maybe striking to put pressure on Uber. The result might be higher wages, paid for by higher prices to consumers, kind of like what is happening here.

But that isn't quite what is happening here. Instead of Uber setting prices, telling its employees what to charge, and negotiating wages collectively with employees, there are just a bunch of independent businesses, linked only by the fact that they all use the same software app to market their services to riders, conspiring to push up consumer prices. They are openly discussing, with their independent-businessperson competitors, coordinating to raise prices for consumers—and using the app to implement their agreement. If they were all coworkers this would be legal, but they aren't, so it's, at least, troubling.

It is somewhat strange how much turns on that distinction. People often talk about how important it is for Uber and Lyft that their drivers be classified as independent contractors rather than employees, because the companies don't want to have the legal obligations (minimum wages, overtime, health care) that businesses have to employees. (“Our business would be adversely affected if Drivers were classified as employees instead of independent contractors,” says a risk factor in Uber's initial public offering prospectus.) On the other hand if the drivers *are* just independent businesses, and if the app gives them a way to coordinate with their competitors to raise prices on consumers, isn't that ... also a problem? “The Legal Argument That Could Destroy Uber,” is how [Jalopnik described it](#), citing Paul's work, and I might not go that far, but Uber's setup does seem to raise some weird antitrust problems that would be avoided by just employing the drivers.

SEC fines

There are different skill levels of fraud. There is the kind of fraud where you lie to people to get them to invest with you, and then you steal their money and either run away (skillful) or get caught (bad). There is the more impressive kind of fraud where you lie to people, steal their money, and then explain everything to them so persuasively that they are satisfied with their losses and don't complain to the authorities. Even better is the kind where you then get them to invest more money with you to continue the scam. I particularly enjoy the kind where you lie to people to get them to invest with you, steal their money, and then explain that their money was actually stolen by shadowy nefarious characters and promise to track down the thieves and recover the money, if they just give you some more money first to cover your expenses.

I suppose an additional level would be stealing people's money, getting caught, getting fined by the Securities and Exchange Commission, and then managing to avoid paying the fine:

The Securities and Exchange Commission over five fiscal years that ended in September 2018 took in 55% of the \$20 billion in enforcement fines set through settlements or court judgments, according to agency statistics. During the prior five years, from 2009 through 2013, the SEC collected on 60% of \$14.6 billion.

Some of that is for sort of double-counting settlements in which the SEC imposes a fine but lets the defendant pay it to some other regulator, allowing both the SEC and the other regulator to count the fine in their wins column. But some of it is for "those behind Ponzi schemes who spent the funds they took from defrauded investors." I would think that if you're good enough at stealing money, you should also have an aptitude for hiding it from the SEC.

Still there is a further level of difficulty; really the move is to do a Ponzi, get caught and fined by the SEC, not pay the fine, and *get the SEC to pay you*. I am not sure exactly how it would work but, you know, there's a lot of money in whistle-blowing. There's that guy who did tax fraud, went to prison for it, and got a \$104 million whistle-blower award from the Internal Revenue Service; I bet you could figure out a similar trade with the SEC.

Things happen

T-Mobile, Sprint Get Merger Backing From FCC Chairman. "Contrary to the common assertion that share buybacks damage long-term growth and investment, we found ... that companies conducting larger share buybacks ... showed higher TSR, higher CapEx growth, and higher employee count growth over the subsequent four-year period." Fed chair tempers fears over

corporate debt meltdown. Pimco's Ties to Architect of College-Admissions Scam Ran Deep. Former Credit Suisse banker pleads guilty on Mozambique bribes. "YouTube summer camps for kids." Kim Kardashian Has a 'Serious Complaint' for Jack in the Box. Dad Gets Himself Irreversibly Stuck In A Playground Toy, Has To Get Cut Out By Firefighters.

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 Particularly because the exchange has basically zero variable *costs*; it's not like executing one additional trade on its computers costs it very much.
- 2 Why? I am exaggerating the model a little; not *every* professional needs a fast connection to *every* exchange. Some will just connect to the most important ones, so the more important (high-volume) your exchange is the more data/connection fees you can probably get.
- 3 Actually many exchanges pay traders for offering liquidity and charge them for taking liquidity, or occasionally vice versa, but on *net* they charge a little bit for trading.
- 4 Not impossible! Any index based on transactions can be manipulated by doing uneconomic transactions. But it is risky and expensive, and probably won't work if there are enough real transactions to offset your manipulative ones. (People tend not to manipulate, like, the S&P 500 Index.) An index based on a *survey* is infinitely easier to manipulate; you just say a fake number in response to the survey.
- 5 Ignore Lyft for these purposes. Obviously it would be illegal for Lyft and Uber to coordinate on prices.

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

To contact the author of this story:

Matt Levine at mlevine51@bloomberg.net

To contact the editor responsible for this story:

James Greiff at jgreiff@bloomberg.net

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

[Read more opinion](#)

[Follow @matt_levine on Twitter](#)

In this article

UBER

UBER TECHNOLOGIE

50.63 USD ▲ +0.14 +0.28%

AONE

ONE - CLASS A

10.59 USD ▲ +0.12 +1.15%

AAPL

APPLE INC

128.51 USD ▲ +0.70 +0.55%

ACN

ACCENTURE PLC-A

267.40 USD ▲ +19.95 +8.06%

[Terms of Service](#) [Do Not Sell My Info \(California\)](#) [Trademarks](#) [Privacy Policy](#)

©2020 Bloomberg L.P. All Rights Reserved

[Careers](#) [Made in NYC](#) [Advertise](#) [Ad Choices](#) [Contact Us](#) [Help](#)