

[SQUEAKING]

[RUSTLING]

[CLICKING]

SPEAKER: So it's my pleasure to introduce Andrew Gunstensen, who I had the honor and pleasure to work with for a long while at Morgan Stanley. And Andrew is actually MIT alum. He got his PhD in Geophysics from MIT. And after that, he worked in a variety of roles in financial-- senior quantitative roles in financial industry. And now Andrew is head of quantitative strategies at Mizuho in New York. So Andrew, welcome, and please take it away.

ANDREW GUNSTENSEN: OK. Thanks, everyone. So I'm going to have a quick talk today about linear models, and rates, in particular, which is area I work in. So we'll just have a quick overview.

We're going to talk a little bit about some basic markets, interest rates-- what are they? What are some of the more liquid rates products, yield curves, a little bit about hedging-- how you hedge your positions? P&L attribution, and then end with a little bit of e-trading and bonds.

I'm going to caveat, this is going to be a very nonmathy presentation probably for this class. So that's probably my personal bias is I don't find math necessarily that illuminating in some of these problems versus what problems are you actually solving?

As I told Peter, we were talking about how to build a house, we wouldn't spend 80% of the time talking about hammers. We're talking about how big you want your rooms and how the connectivity is. I would highly encourage you to ask questions.

I'm going to ask Peter, if nobody else does, if he can pipe in a couple just because I like to keep the discussion going. And I'm generally happy to yack on and on about any particular topic that anyone is interested in. So please talk away.

So start with the basics. What are linear products? They are things that are not nonlinear. Nothing too surprises there. Some of the major categories are things like bonds, high interest rate swaps, interest rate futures, CDs. These are the more traditional relatively simple products.

These are not CDO-squareds or callable spread range accruals, or any of the super exotic stuff that has been around for a while and still continues to be done. No, the modeling is in some sense relatively simpler here. These are not huge, multi-factor, stochastic calculus models you have to do and term structure models and things. It's relatively straightforward.

So the question, then, is why do we care? As quants, why is this interesting? Anybody want to hazard a guess as to why this stuff is still relevant today? Somebody take a guess. You look like you're thinking.

AUDIENCE: Because it's still related to important stuff like interest rates.

ANDREW GUNSTENSEN: It is. It's an important thing. It's the core underlying. What are some other reasons? Maybe you might care about this?

AUDIENCE: It's traded a lot.

ANDREW It's a very big market and a very liquid market. True. A couple other ones are it's also generally used for hedging
GUNSTENSEN: of the more complicated stuff, and it's also an input to all the more complicated models, which generally, you start from the simple stuff and you build your models up.

These are very large. These are the largest markets on the planet. Just a couple of statistics-- in the last few years, equities is getting close now, particularly with the run-up in global equities last couple of years. Bonds have something north of \$120 trillion outstanding. Interest rate swaps are more like \$500-plus trillion dollars outstanding. The main futures contract, more than \$10 trillion.

These are very large numbers. You don't need to have a huge profit margin in a market this size to do very well. In a market of this size, it's very actively-- it trades a lot. The daily volumes, Treasury's approaching \$1 trillion a day. Interest rate swaps-- comparable.

Again, futures, something like four million contracts a day in the SOFR futures, which are the main short-term interest rate contracts. And then I think a lot of the more complicated models start with the simple stuff and then build up on top of that.

And again, I have a tendency to talk faster and faster and faster if nobody comes back. So I'm going to apologize in advance. And if you find me doing that, somebody just tell me to stop.

And again, I said-- they mentioned they're-- basically, they're complex and inputs to the more complicated model. Previously, I think they were thought quite simple. I think the global financial crisis in 2008 brought a number of factors to the fore that people hadn't considered before. In particular, around things like funding.

And what was previously, say, a very simple product-- so I'll pay you \$1.00 in 10 years. How to value that became quite difficult and subject to a lot of thinking and discussion. And there's not a unique answer to that, which is also interesting.

A couple of the other things that came in. There's a lot of regulation that came in after 2008. In some sense, somebody said the regulators wanted to make banking boring. And it's not boring, but they've changed where the interest is in it. And the super complicated exotic products that were really very much in vogue in the early naughts and mid-naughts, really, those things ended poorly for a lot of people.

You think about the causes of the financial crisis, at least on the market side, a lot of it ultimately boiled down to US residential real estate, which had a huge run-up and then an enormous crash. But off of that, there had been issued a lot of very complicated securitization products, very highly levered and multi-layered. And things were very opaque and hard to value. This was not financial modeling's finest hour.

So I think there was a lot of regulation that came in on top of that in terms of model risk management, and that really of ended that section of the business. So people went back a little bit to basics and started thinking a little more about some of the fundamentals, which I think was a good thing because I think things were kind of silly there.

And I remember talking to one of my friends at a soccer game or something. Our kids were playing. And he was mentioning they'd just done a mortgage securitization issue, which I believe defaulted without making a single payment. So this stuff was just so poorly supported and thought out. But because it was so opaque, a smiling, fast-talking sales guy could push a lot of this stuff to areas that really probably should not have been involved in this area.

And you find little municipalities in Norway having to declare bankruptcy because they bought a whole bunch of CDOs from Iowa strip malls or something crazy like that. And you just think, how did that happen? So I think there was quite a lot of regulation came in, really, to put a little bit of fences around some of this stuff, which is quite powerful, but it can be used for good or bad in some sense.

So let's talk a little bit about, what is an interest rate? An interest rate, at its most basic, is a rate at which you can either borrow or lend money on a principal amount. So in this case, I'm going to have my interest, I , payment will be N , which is a notional typically, times the rate, r , accrued over whatever time I'm borrowing it for.

So I have-- I'm borrowing \$1,000 at 5% for a year. In a year, I'll have to pay back \$1,000 plus 5% times a year times \$1,000 with \$50 of interest. Relatively straightforward. How do I value that in terms of what I'm willing to pay for that today? I'm not willing to pay \$1,000 for it because-- well, let me reverse that.

So if I get \$1,000 in a year, how much am I willing to pay for that today? Now, whatever I'm willing to pay for it today, I can invest. And if you work out arithmetic, it works out to be 1 over 1 plus the rate times the time. So at a 5% rate, I'll pay something like \$950-odd dollars for that. And then in a year, get \$1,000.

Now, this depends a lot on the discounting rate, for example. So what rate am I going to discount at? That's hard to observe. Rates I can get for investing are easy to observe. You walk into a bank, and you look at the signs in the front door-- buy our CDs, 4.5% rates for a year, whatever. How much to discount is harder.

Typically, in our world, we're going to assume one of the institutional rates, which is something like a SOFR rate or an OAS rate, which we'll talk more about in a little bit in the future. But these are a standard assumption.

Now, one of the interesting things practically is you find every person has a slightly different rate that they should be discounting because it depends on where they can borrow money, and that depends on your credit worthiness. So you may have a good credit rating, and you can borrow money at 4%. You-- no offense. I'm not picking at you. You may have a terrible credit rating, so you have to pay 8% to borrow money. You may have a medium one, so you're borrowing at 6%. And it all depends on the particular issues of the issuer.

So this makes it interesting because, at this point, a contract for you guys to pay me \$1,000 bucks in a year has different prices depending on who you are. So there's no single price here makes it interesting. A lot of the banks will deal with this because it's kind of a pain in the ass for our trading desks to have to manage 1,000 different counterparty relationships, each with their own funding.

So usually, you'll centralize that in a central funding desk of some sort. They will then assume all of the individual traders, can fund at a relatively generic rate. And the funding desk will deal with all these horrible little funding bases and small collateral exposures and things like that.

Going back a little bit in history, you may or may not have heard of LIBOR. Probably not, because it now is gone. But this was one of the first really common user rate. Stands for London Interbank Offered Rate. It was developed, I believe, by a Greek guy in the '60s, as a method of standardizing syndicated lending in London.

Most bank lending to corporates is floating rate lending. So you won't do a fixed rate loan for 10 years. You'll do a rate that resets every month, every three months at the then current rate. And the question then was, what rate should we be using to set the interest payments on these loans at?

So what he came up with is a notion we're going to poll the banks as to where they lend money short term. And we'll take the average of that, kick out the highs and lows, and then that'll be a standard rate everyone can use as a reference index. It was a very, very good idea. Incredibly successful.

First interest rate swaps, which are really synthetic versions of those, came out in the early '80s. First currency swap, I believe, was 1980 or '81, which is a cross-currency swap with IBM. And then the first pure single currency interest rate swap came out the next year.

These products had phenomenal growth in the '90s and naughts. Currently, there's something about 350 or 400 trillion of them outstanding as of LIBOR cessation and still continue to grow a little more than that. So very, very popular product, very liquid, very actively traded.

LIBOR is a rate, though-- again, another one of the fallouts from the financial crisis showed that it had some issues. LIBOR was set by polling banks. So you'd go ask the funding department at each of the big banks-- I think there was 20-odd people in the LIBOR pool-- where are you going to lend and borrow-- where are you going to borrow money from? They've come up with a number. They'd quote it.

Note that it's not based on actual transactions. So they're just making a number up. They're trying to, in theory, if they're all good players, hitting it where they would actually do the trades. But this is not actually observable trades. And I think in the late naughts, early teens, a group of traders of less than honorable intent realized that this is, and this was a manipulatable rate.

Now, you may have-- for example, every day you're going to look at how much money-- these rates set every day. And based on your portfolio, you'll have some impact on that. Some of them, you can have millions of dollars a basis point, a basis point being 100th of a percent. So if you can get that rate up or down just a little bit in your favor, it can end up making you a lot of money.

So what they did was they essentially all chatted with each other, and they were chatting with whoever the people at their banks were responsible for setting these rates and convincing them, look, I need you to make this a high rate today, or I need you to make this a lower rate today. And can move this thing a little bit because it's not observable because nothing trades. And they were doing that for a while to their advantage.

This was discovered. A number of them were in jail for some time. There was very substantial multibillion dollar fines to the banks doing this. But I think the realization came that this rate didn't have any support.

One of the things is LIBOR is-- it's an unsecured rate, if you will. It doesn't have anything backing it like a repo or collateral rate. And there's just not a lot of transactions. Post-2008 nobody wanted to lend unsecured because you had a number of bankruptcies. People like Lehman's going out, and don't want to lend to money who's nothing backing it. So people weren't happy with that.

And some of the things and some of the tenors and markets, I think, had three transactions in a six-month period. That's not enough. But you have to quote this rate every day. So there's just nothing fixing it. So that was a real issue.

So the Fed set up-- or the Fed plus the other regulators-- I guess it came out of London, originally, was the big driver for this, the FSA there, LB, London and LIBOR, really decided you needed a new rate. It was going to be defensible backing up based on actual transactions and a number of other criteria to make it hard to manipulate. So would be a genuine real observed rate and not something that was just being made up by banks to favor their trading desks.

AUDIENCE: See-- just a comment.

ANDREW Yep.

GUNSTENSEN:

AUDIENCE: This LIBOR rate may sound a bit esoteric at this point. But at the time, mortgages, mortgage interest rates were often quoted as LIBOR plus a spread. And my experience was during the time of the naughts and the teens, looking at possibly refinancing my mortgage to a lower rate because rates were coming down, and the LIBOR did not move much at all during that period, which was rather frustrating to watch. But then when understood why it was not moving.

ANDREW Well, it was partly that and partly the LIBOR to Treasury spread-- which is the difference between LIBOR rates, **GUNSTENSEN:** which are the bank rates, and the Treasury rates, which are government rates-- blew out in that period, around 2011, '12.

The fine institution I was at at that point wrote off \$8.4 billion in structured notes because of that, for example, because they were hedging their structured notes with LIBOR products and assuming the two would stay together, and they went, whoop. Most of it came back in the next couple of years, but it was quite unpleasant for a while there and made them want to change how they were doing that.

When we went through the LIBOR retirement, when you look at how much of what kind of products are out there, there's all sorts of weird stuff. Swaps are the biggest piece, but then you have floating rate mortgages. You have a fair amount of corporate bonds out there.

You have some really bizarre, slightly esoteric stuff that is like-- I don't know how you redo that, going to go knock on somebody's door and tell my mother that her LIBOR plus 100 mortgage is now a SOFR plus 150. I don't know what she's going to do with that other than be confused and call me and ask for help.

So really, the LIBOR had to go because it was not a good rate anymore. And this is also true of the equivalent LIBORs in-- at least, Japan, Canada, US, Switzerland, Europe, and the UK, I think all have replaced their rates with something a little more defensible.

In the US it's called the SOFR rate, which is the Securitized Overnight Funding Rate, which is essentially the average of overnight repo trades. So a repo rate is if I have a bond, I can just buy it outright if I want to. But it's not usually an efficient way to do.

I can borrow money against the bond. So if I buy 100 million US treasuries, I can then lend that to somebody, post that to somebody and borrow, say, \$98 million of cash. He'll hold my bond. And then if we want, we can swap it back in again.

But the rate at which you're going to pay in that borrowing is known as the repo rate, or reverse repo rate, depending on which direction you're going. And that's a real rate. Real money changes hands. US has about \$800 billion of that a day of the main repo facilities. So that's something that there's a lot of transaction data behind. You can look at it, measure it, add it all up, average it and come up with a rate. So that's what they did.

SOFR does have some features that are not so fabulous. First of all, it's a daily rate. Nobody wants to have a loan where you have to pay interest every day. It's just a pain in the ass operationally and not good. You want to pay every three months or six months or annually, for example. That just makes it just easier mechanically to deal with.

You can compound it into term rates. So we'll talk about that in a slide or two. Another thing, if you do compound it into term rates, so you do like the average SOFR over a three-month period and then pay that rate. You don't know the rate until the end of the period because you haven't finished doing all your daily features.

That's a bit annoying for people, because a lot of people want to know the rate upfront so they can just arrange this operationally, set all their payment systems, make sure the payments get their, schedule that, have enough time to sort out, make sure they're doing the right rate.

And the other interesting thing is SOFR is a secured rate. The banks actually prefer to lend unsecured. So I think when a bank makes a loan-- so it'll do a floating rate loan to a corporate. They'll give them \$100 million. Corporate will pay them the floating rate, say LIBOR plus spread or whatever.

Now, if you think about what happens when bank credit goes bad. So if there's a crisis, typically the bank spreads will blow out, which means LIBOR will blow out. And if the majority of the income is coming from LIBOR products, because that's what you've loaned it, you're not making more money.

So you loan money at LIBOR plus 50 LIBOR is now 100. You got an extra 50 basis points of income because your credit's deteriorating. So it's a right-way risk kind of feature that is-- I don't know if it's worth a whole lot of the bank. But it's a nice feature to have, that if bank credit goes bad, you end up making more income. So people like that.

Said LIBOR has mostly been phased out in the last few years. US ended LIBOR last June, so June 2023. They're one of the biggies. Sterling went earlier. Canada went two months ago. Mexico is going into this year. And then there's a couple of features in yen that are also happening I think early next year. This was an enormous migration. And you should hope in your careers you never deal with this kind of stuff because it is just painful and tedious. Yes?

AUDIENCE: Do you mind explaining the relationship between bank credit and the funding rate?

ANDREW GUNSTENSEN: So I'm a bank. Majority of my corporate lending portfolio is floating rate loans, which are going to pay me an interest of LIBOR plus a spread, fixed spread. So it's going to float. If LIBOR is higher, I make more money. If LIBOR is lower, I make less money. I'm going to fund myself as something that's typically not-- less floating the [? rat. ?]

So credit goes bad. So banks get into trouble, or somebody perceives that there's a systematic issue out there with banking, like 2008, or a couple other episodes. The LIBOR spreads are going to increase because larger spreads, higher rates, lower quality, generally.

So what that means is the spreads are going to increase. My interest I'm going to earn on my loans has now gone up, which helps me because I get more income. Make sense, or--

AUDIENCE: Sorry. Yeah, thanks.

ANDREW Yeah. No, please. Again, please-- I'm talking fast on a big range of topics, so please jump right in. You're looking
GUNSTENSEN: confused or anything? No questions? No? Yeah, yeah.

OK. You hope in your career you don't have to deal with this because it's a pain in the ass. It's just an enormous mechanical process. There's been a few of them around. Euro conversion was another one in the late '90s we did, where all of the European currencies converted to euros.

Yeah, I remember sitting in Canary Wharf with that one at 3:00 in the morning because we'd screwed something up. And it was like, oh, dear. But it all happened. And then Y2K is another classic one. There's been a bunch of other ones like that. But it all happened, and fine. So no more LIBOR.

So as I mentioned, SOFR is a daily rate. You don't want to pay daily. So you can compound it up. Two standard ways-- I'm going to do daily compounding. So each day I'll work out my discount factor for that day, which is $1 + \frac{\text{rate}}{360}$ plus the rate over 360. And you just compound those up, and then subtract 1 off the end. And that'll convert it to an equivalent term rate, and you can pay the term rate.

Or I can just do a simple average. So just average the SOFR, or average the daily SOFRs. In this case, the little d's are the number of days it applies for. So a rate setting in Friday will apply for three days over the weekend. And a rate setting on Monday only applies for one day. Just do that. Add it all up. Pay that rate.

The majority of products in the dollar market go on the compounding SOFR because it's a simpler rate to calibrate. It also doesn't have the averaging you need to make a slight correction due to volatility, which is just annoying. It's a few basis points running, so most people won't care. But if you're trading this a lot as a market maker, you want to get all this stuff right because the bid offers are a couple tenths of a basis point. So if you have something that pushes you outside of that, you're going to lose money.

There's lots of other rates out there. Term SOFR is one the CME created a couple of years ago. They actually fit something. They have a little model. They do a best fit to all the futures prices, which are again, are very liquid, well-traded stuff. And they work out one month, three months, six month rates and then publish those.

It's a set in advance rate. So I set it now, and it applies for the next three months. So I'll fix a term SOFR now. Pay in three months at that rate plus whatever. This one is quite common in corporate loans.

And again, corporate lending-- the banks, in general, have much better systems than the users. Not always the case. But they've been at this for a while, where if I'm a corporate, I don't necessarily want to deal with a lot of the complexities and nuances of this stuff. I just want to have something fairly simple. Every three months, I'll look I rate up on a screen. I'll put that on my loan, and I'll pay that three months later. Yep.

AUDIENCE: I don't know if this is like-- if I'm not understanding the English this correctly. But let's assume we go back to when rates were highest from now, like when rates were-- when there was more of a risk of rates going above 6%, that would directly influence how the duration of the futures contracts. Is that-- am I--

ANDREW I'm not sure I'm understanding the--

GUNSTENSEN:

AUDIENCE: The cheapest to deliver for the futures contract changes because the price, like the 6%.

ANDREW Sure. So at 6%, your futures contract will [INAUDIBLE] at 94 because the 100 minus the price is generally the

GUNSTENSEN: rate. Fine. Yep.

AUDIENCE: So then in terms of thinking about how-- do you ever think about how that affects the term SOFR. Theoretically, that would change the duration of the cheapest deliver, which means that your term SOFR represents a different rate.

ANDREW Term SOFR-- it is a slightly different rate. So let me doodle here for a little bit, if you don't mind. So that's time,

GUNSTENSEN: and that's rate.

So what it is, the futures will each set a block of, say, three-month futures. We'll set an average over a three-month rate. So I might have a rate that does that, and another future that does that and that. Might have a little stub period in here, which can go back to that future. So each future will be set. And it's essentially the average over the term of the future.

So what CME then did was they did a fit of this model to the futures contracts. And they then interpolate the daily rates, say, for six months, which will run from here to about here. So they look at that, and they'll work out what the equivalent six-month term rate is based on these rates to compute that. Does that make sense?

AUDIENCE: Yeah, I think I'm not asking this question. I did mean 1%.

ANDREW OK. Anyway. So CME came up with that. The loan guys like it. Again, end user corporates don't want to deal with all the mechanics of daily sets and operational stuff. It's just a pain in the ass for them. So they're happy to let somebody else do that. It's the usual set of events-- pay six months or three months from now. Fine. That makes everyone happier.

The Fed, however, doesn't want this to become a de facto new LIBOR rate. So they actually have quite severe restrictions on how much the dealers can trade in it. Essentially, dealers are only allowed to trade it to hedge their loan portfolio, which also is an interesting feature, which means it's an entirely one-way market.

Because banks issue loans. They have to hedge it. It's all in the same direction. So there's a very substantial basis between regular vanilla SOFR swaps and the term SOFR swaps, something around four or five basis points, which, again, in a market context, where that's 20 or 30 times a bid offer, it's way outside the bid offer range.

There's another older version of SOFR called FedFunds, which is really an average of slightly different FedFunds drawing, similar to SOFR, that's still in wide use. And there's a whole lot of other crazy rates-- prime, DISCO, municipal or military housing rates and all sorts of other stuff. DISCO's my favorite rate-- not for any reason other than I like the name-- which is a Fannie Mae discount rate.

So now discounting rates. One of the questions is, how do you work out what the discounting rate is going to be? For example, what rate should you use to discount a future payment?

So let's assume we have a simple case of a fixed payment-- F for fixed, at a time, t . We assume some discounting rate. And we'll value the payment today, give a present value of P . Now, we have to fund this rate. Generally, in this world, you assume that you are not just spending money out of your back pocket. You're going to go borrow money from somewhere and use it to purchase this thing.

So we'll buy P , and we're going to buy at whatever our funding rate is. So whatever we can borrow money at, that's what we we'll do to fund it. And now just think a little bit about from one day to the next day, what are the cash flows involved here?

You have two of them, essentially. You're going to have to pay money to borrow in your-- you will have to pay interest on your borrowing. So whatever interest rate you've done for one day, you're going to pay that. And then your underlying asset, which is the fixed cash flow in the future, will increase in value a little bit because you're getting one day closer to its payment.

So the interest paid-- the interest is just going to be whatever the present value of that is, be it t_0 times the borrowing rate over 360. So that's how much you're paying to fund this. And the increase in value is going to go up by yesterday's value times $1 + r$ over 360. That's the compounded rate accrual factor. Or the change in that is, again, the value times r over 360.

Notice that the first and the third equation have very similar forms here, and the only thing different are the two rates. So this is actually going to imply those rates have to be the same, or there's an arbitrage here. So that means that your discounting rate that you should be using to discount on future cash flows is whatever your funding rate is. So whatever you can borrow money at, that's the appropriate rate to discount future flows up.

Now, if those rates are different, there's a trade. You can either borrow cheaper than you can invest, or you can invest cheaper than you can borrow. And you do one or the other. And that'll come back. This is a very-- if you only pick one thing away from the hour and a half we're spending here, remember this because this is-- funding and discounting are the same.

Let's talk a little bit about some of the products. Just going to talk about the vanilla one. So the question is, how do I compute what I think forward rates should be and what projections of those should be? So what I'm going to do essentially is build a curve to do that.

I'm going to start with by thinking, what are the easily observable instruments in the market? What can I see out there that I can think-- whatever my pricing model is going to be, these are the things it needs to hit because I can trade these things.

Some of the common ones we're talking about are CDs, which are Certificate of Deposits. These are very short-term bank lending, borrowing. You can walk into a bank and buy a one-month CD at 4.5% if you want. Maybe a three-month one, a six-month one. It doesn't go much beyond that. Maybe a year. Again, fairly liquid, quite common. Interest rate futures-- very liquid products. You can do that. And finally, interest rate swaps.

Now, so for each of these, we're going to work out what is the price of this as a function of my discounting curve. And then given that, the cash rates, I can invert that to find what my discounting implied curve is that prices all of these instruments back to market.

So CDs-- nothing too hard here. I buy a CD. I get the amount I invested plus the rate accrued over the time. The value of that is that forward interest amount plus the forward return of my principal discounted back or inverted. I just get-- z is 1 over 1 plus r delta t . That's again, that's the thing you're going to see all over the place.

These rates are well-observable. They're a little bit tricky because the swaps desk typically can't necessarily trade them because these CDs are funded products. I have to actually buy them. If I want \$100 million CD, I need to somebody \$100 million. I have to get that.

Swap desks don't like that. They prefer derivative instruments which are not unfunded instruments because you just have to put less cash up front. And everybody likes less cash up front because it means you can do more, leverage yourself up more and so on.

Second one is interest rate futures. Again, very liquid contracts. US ones trade on the CME exchange in Chicago. Every country generally has a relatively active short-term rate. Canada trades in Montreal. The UK contracts trade on, I believe, the LME. The Euro contracts trade on the Eurex Japan contracts trade on Tokyo, I think, Stock Exchange.

So two main contractors-- there's one-month contracts and the three-month contracts. They're slightly different, which is a bit annoying. Three-month contract trades on a compounded rate. One-month contract trades on an average rate. Three-month contract goes from IMM date to IMM date. IMM is International Monetary Market, which was an old market in Chicago that set the standards here, which is the third Wednesday of every March, June, September, December. It's just when the contracts settle, and that's fine.

The one-month contracts trade on a calendar month. So today's contract started in September 1. We'll settle on October 1. I think I may be off by a day there.

The futures settle at 100 minus the rate. So if I have an interest rate of 5% when I'm settling a future, that'll be a 95 price. And again, people just did it that way because they like to know markets rallying, contracts go up. Markets selling off, contracts go down.

For bonds, price and yield move in opposite directions. So yields go up, prices go down and so on. So that's why they set the 100 minus that.

There's a small convexity adjustment here as well dealing with futures margining, which I'm going to just wave my hands and gloss over, but it's something you have to be conscious of. Futures do daily margining. It's a standard part of a futures contract.

So if I have an interest rate futures contract, and I enter it into it, and generally, you'll enter into a futures contract at no upfront cost, aside from some initial margin, perhaps. Market moves up. Futures contract goes from 90 to 91. They'll give me one back in margin. If it goes from 90 to 89, I have to pay one to remargin it. So at the end of every day, you settle up the net move on that day, and they do this to minimize credit exposure between them.

But I think if I have an interest rate contract, when the price of the contract goes up, I'm getting money back. But the fact that the price went up implies rates have gone down. So that means I can reinvest my money now at a lower rate.

Or if the contract goes down in value, I have to fund that more because I need to ante up more margin to the exchange. But I'm having to do that at a higher interest rate. So this correlation between the flows of money and what it's going to cost me to borrow or lend, that introduces a small valuation impact.

So let's talk about interest rate swaps, which are probably the biggest single product on the planet. This is basically an agreement, a custom agreement, between two counterparties to exchange two sets of cash flows. In the most simple case, one of these is a fixed rate.

So I'm going to pay you 5% a year for 10 years. And the other one is a floating rate. I'm going to pay you in compounded SOFR. I'll pay you annually every 10 years. We'll do that for 10 years, and then we'll call it a day.

There are many, many variations of these structures. There's all sorts of crazy stuff. And imagination is the limit here, as long as you convince somebody else to do the other side of it. They're a tremendously useful instrument.

They're great if I need to hedge something or change my risk profile, it's a very efficient way to do it. These are all derivative products, means they generally don't use balance sheet. They don't take a lot of upfront capital. They're cheap to fund. So they're quite popular for those reasons. And I said, they are the most liquid interest rate product out there.

They're one of these things that there's a lot of mechanics to get it right. In particular, if you enter-- you work in this field, you will get very good at data arithmetic, which is a stupid thing.

I remember once I was interviewing a guy, and he had just done a PhD in numerical methods, I think. And one of the questions he asked me, which was a great question, so keep this in mind when you're interviewing. He said, if I come and work for you, a year from now, what will I know then that I don't know now? And I said will know how to do data arithmetic. And he looked at me kind of funny, went away. And they actually gave him a job, and he took it. And he's a fabulous guy.

And he came back a year from now, and said, you know, you were right. I know how to tell between my birthday and today how many days are there. It's not a unique answer to that one, by the way, just to make it [INAUDIBLE]. It depends on your conventions.

And I'm also going to guess in interest rate swaps, at least 90% of all problems are date-related. That's the first thing you want to look at. Do I have the dates correct? Am I using the right holiday calendars?

And holiday calendars-- again, it's another one of these things that's occasionally annoying. They change. We had Juneteenth added a couple years ago, for example. We had the Queen's funeral added. Every time a president or ex-president dies, there's a holiday. And just these-- Canada had Truth and Reconciliation they added recently. So you have to make sure you build your models to account for the fact that the holidays may be dynamic.

So I just put on there a simple standard dollar swap rollout. And there's a half dozen different steps to actually get what are the accrual periods and the cash flows and the payment dates and the reset dates. You start with today. You go to the trade date, if it's a forward starting date.

From that, you typically add two days. You go to the effective date. You then add the tenor of the swap to get the unadjusted maturity date. You then roll that back in increments of the payment frequency, get the unadjusted roll dates, and then you adjust everything based on a holiday calendar, which is typically going to be New York banking or something. And again, you've got to be really precise on the calendars.

Another one of the questions I like to ask people is, how many distinct dollar calendars do you think exist in terms of swaps monitor, which is a standard source for all calendar data? How many dollar calendars do you think they have? Anyone want to take a guess?

So I have New York banking, for example. I might have New York Fed. I might have-- somebody guess. There's no penalties for getting it wrong. [LAUGHS] OK, Peter, what's your guess? You've been around for longer.

AUDIENCE: Let's see. I would guess maybe 150.

ANDREW GUNSTENSEN: You're in the right ballpark. Last time I checked, there was 294. And every exchange has got its own calendar, and they're all different. Every major business center has its own calendar. And it's just if you get the wrong calendar, you're going to price stuff on the wrong date, and you're going to misprice things. And then winner's curse rule says if you price something wrong, you only get the trade when it hurts you. And again, dates are just an enormous pain.

To value the swap, we're going to do it in two pieces. We'll start with the fixed side. So fixed side is relatively straightforward. We roll out all the cash flows. We work out how much we're paying and what dates we're paying. We then, from our curve, get the discount factors there. Sum it all up.

So in this case, I've just got C, which is the coupon rate. Delta is the accrual fraction. So if it's an annual swap, that'll be around one notional. And then times the discount factor. Sum it up. That's the present value of the fixed side.

Floating side's a little harder but same idea. But instead of the coupon rate, I've got the forward rates, which are little f. So I have to work out what the forward rates are, but we'll get to that in the next slide.

Some mechanics about what are the particular specifications. Is it set at the beginning? At the end? Is it forward? Maybe a spread in there. What you end up with, there's two different curves in here. I need one curve to generate the forward rates from, and another curve to generate the discounting factors from. In the vanilla swap, those are both the SOFR curve. So they're actually only one curve

So what do you want to do here? There's another bit of magic here that I can actually generate an implied floating rate from fixed flows only. So if I assume I can borrow fixed rate bonds, I can borrow one at a given-- at a first date, invest it, and pay it back at a later date, invest it at the unknown floating rate, and then pay it back at the end.

And I'm going to assume I'm going to invest at a rate that's a fair trade. So PV of that is zero. So I end up here with-- and the initial term there is me borrowing the money at time t_0 , me getting the interest rate at time 1, and me getting the principal back at time 1. So given all that sums to zero, I can work out what the implied forward rates from this discounting curve are. So I use that, plug it through the swaps, and now I've got a swap market.

So for an implied swap markets, generally when you do a swap it's usually done at a 0 PV, or something pretty close to a 0 PV. So when I look at the two legs and sum them up and make them zero, I've got the fixed side plus the float side, both just the sum of the fixed rates times the discounting factors. And then subtracting the float side, I'm just giving you factors. Do some arithmetic, and I get what the implied swap rate coupon is.

Now, this is useful because swap rates are observed. I go to my favorite Bloomberg screen, and I can see all sorts of swap rates. So that's what I have to work out my zero curve to fit those. Any questions? I feel like I'm just motoring through this here right now. Yeah?

AUDIENCE: Are these slides [INAUDIBLE]?

ANDREW Sorry. What was the question?

GUNSTENSEN:

AUDIENCE: I was going to ask if we could get the slides. I don't know if they posted.

ANDREW I can certainly send them. I don't think I have sent them.

GUNSTENSEN:

SPEAKER: No, we don't have them yet, but we'll get them.

ANDREW There is nothing too top secret in any of this. So that's how you [INAUDIBLE].

GUNSTENSEN:

AUDIENCE: So Andrew, one comment. I'm not a fixed income expert, but the zero present value at trade initiation is quite an important feature, I think, where it encourages people interested in either side of the swap to be comfortable with making the trade.

ANDREW Yeah, that's the so-called par rate, which is the rate that prices are at par, a.k.a. for a swap that's zero. You could
GUNSTENSEN: price it at something other than that. And again, bid offers in this market are usually a tenth of a basis point, 2/10 of a basis point. So it's around zero.

Some of the structured transactions will have swaps that are way off market. And then that may get accounted for elsewhere. Or that just may be where you're taking your P&L out of.

AUDIENCE: Just another comment that the-- a long time ago, this was like the late '80s, I was at the Sloan School, and swaps were becoming very popular. And one of the reasons they were becoming very popular was that certain corporations could access certain fixed rates at very low rates. And less creditworthy companies could not have access to such good rates for borrowing, but did have access to floating rates. So these corporations would pair up with a swap. I don't know to what extent the swap market now is.

ANDREW

I think it's well past that. But at the time, if I'm a lower rated guy, and I can maybe borrow short term, but

GUNSTENSEN:

nobody's going to lend me money long term because they don't trust I'll be around in 20 years, but maybe they're comfortable that I'll be here in three months, and another guy who can then fund for 20 years because he's a high credit rating, we can both borrow in our respective countries and then do a swap.

So I now have fixed rate swaps, and he has floating rate swaps. And depending how they want to manage their treasuries, that may be a great trade for all sides. So yeah. Again, swaps are very flexible and let you move the risk around in quite a flexible way and quite a customizable way.

So let's talk now about the core valuation model really in interest rate swaps, which is the yield curve. So we're going to parameterize the yield curve. And the yield curve is used to-- it's just to get a function of the discounting as a function of time, which then feeds all of the other valuation models.

So we're going to parameterize it by a set of times and zero coupon bond prices. You can parameterize it in a number of different ways, but this is one of the easy ones. I also need to choose a method to figure out what is my zero-- the Zs, if I'm between knot points.

So maybe I have annual knot points. So I'll do a 1, 2, 3, 4, 5, up to 10 years, for my times, and I'll have the Zs at each of that. But suppose I need a 4 and 1/2 year discount factor? How do I do that? Lots of possibilities here.

First of all, I'll talk about the knot points. How do you choose which times to use to interpolate that? Simplest one, which is the most common one, is you just use the maturity dates of all your instruments. So I'm going to build it with, say, 1, 2, 3, 5, 7, 10-year swaps. I have knot points-- 1, 2, 3, 5, 7, 10 years. Use that. That's the easy one.

There's other possibilities, like FOMC dates, which we'll talk about in a minute. You also have to be a little careful on picking your knot points to work well with your interpolation method. Because some of them, there's some-- you can make bad choices here, which will give you unstable hedging.

We then are going to fit a set of market observables here to get the Zs. So given that we've chosen an interpolation, our knots. We have, hopefully, n knots and n points we can solve.

Let's talk a little bit about CDF spline. This is a constant daily forward. This is one of the easy ones, which you assume that between your knot points, your daily forward rates are constant. So you end up with a curve that looks like a skyline or a stair step or something.

Just an example there, a simple one, where in this case, you can see what my knot points are. I think I have 6 months, a year, 2 years, 3 years, 5 years, 7, 10, and 20, I believe, and 30. And the forward rates are just flat in between them.

This is perfectly good. This is really fast. This is a local interpolation method, which is one of the key features. And what I mean by that is if I'm looking at-- I want to find a discount factor say, or there, it only depends on the value there and the value there. It doesn't depend on anything here or anything there. So I'm just going to interpolate between those two.

AUDIENCE:

Why do you not choose to not [INAUDIBLE] based on minimizing error?

ANDREW On what?

GUNSTENSEN:

AUDIENCE: On minimizing error.

ANDREW How are you going to do that?

GUNSTENSEN:

AUDIENCE: Like Chebyshev, I guess.

ANDREW You could. What you want to do-- the thing-- on the simple case is you want to maintain a connection with your
GUNSTENSEN: knot point and your calibration instruments. So in this case, I'm choosing the maturities.

You can do lots of other stuff. But you could-- for example, let's suppose you had this set of instruments according to the knot points here. If I choose, say, knot points at 5, 10, 15, 20, 25, 30, I'm overspecified here, and I'm underspecified there. So you need to choose it with regards to how you're going to calibrate, particularly.

And the other thing that comes out is when I work out in the simplest cases what my hedges are, it's in terms of my yield curve input instruments. So I want to choose a good set of liquid, observable stuff that I can actually transact in to do the hedging.

There's a zillion ways to do this, and that's still-- 30 years in, that's still a very active area of research and discussion. So lots of stuff. And some of the areas, like some of the computer graphics, for example, has a lot of advanced interpolation stuff, B-splines and beyond. And all of those are likely entirely relevant here. Just haven't looked at it.

So anyway, so this one gets you your very characteristic shape. So this is probably the least smooth, simplest, fast method. So it's got popularities for that reason. On the other side, the most smooth, little slower is cubic splines.

So in this case, I'm just going to fit a piecewise cubic polynomial here. So I'll work out what my little parameters are. And then I'm just going to sum it up. I have four terms, cubic in each interval. I impose constraints at the edges of the intervals on continuity and smoothness and do that. Again, this is still a reasonably fast method. Gives a very smooth curve.

Now, this curve is not local. This is a global curve. And if any of you go into interest rate quant space, you need to make up little laminated cards saying, I can have fast, local hedges, or I can have smooth, global hedges. And it can be anywhere in between.

But I've had this argument for 30 years with traders. Traders always want-- if I'm going to do a hedge of an eight-year instrument, they want it to be only seven and 10-year hedges. They don't like it when the hedges spill over to five years or 15 years. But smoothness is a global property.

So if this is my curve, and I have a point, some instrument in this area and I'm pricing it, it's affected, obviously, by these two. But if this point here was down here, I would have a different shape through that area because I'm maintaining smoothness, again global. So what the global thing implies is that I'm sensitive to points much farther out from my maturity point. And again, this is an argument we've had for decades.

Some of the smooth methods are better than others. Cubic spline is probably the worst because it is the most smooth. I guess you could do a quintic spline or something if you really wanted to be pathologic, but that would be a weird thing to do.

OK. So then, how do I calibrate? I can write my error function. I'm going to price each instrument as a function of my proposed curve. I'm going to from that take the difference between that and what the actual observed price is.

So for-- I can then just-- if I have n knots and n calibration instrument, I just solve that exactly. Pick your favorite numerical solver, and off you go. Or, in some cases, you can actually bootstrap it and be fast.

Or if I want to relax that constraint that I have to have the same number of calibration instruments as points, maybe I'm going to do a least squares fit. Again, just pick my favorite minimizer and let it go and see what we come up with.

Now, you can also include lots of other stuff in your objective function other than just the errors. So one of the [INAUDIBLE] we built with, for example, are [INAUDIBLE] smoothness constraints. So I'm going to underspecify a curve. But then I'm going to impose additional smoothness constraints because I want my curve to be smooth. That's fine.

I can also drop the exact fit requirement a little bit. There is no single price in it. There's always a bid offer. So maybe I'm happy with my curve prices things to within a bid offer. That's close enough. I can't distinguish that from an exact fit. And any number of other extra constraints you can add here.

Now just one of the other ones. And actually, when we take a look at rates-- so this is SOFR rates over the last, I think, three or four years-- this doesn't look very smooth. This, in fact, looks very stairsteppy. And what you'll notice here is the rates are-- this is a little-- this is an interesting spike-- cuts for tax, receipt issues, which caused a lot of consternation in the industry at the time. But we'll ignore that for now.

It's generally flat, and then it goes up. It stays flat. It goes up. These are the central bank meeting dates. So these are when the FOMC gets together and decides we're increasing rates. The daily rates will generally stay constant until the next FOMC rate change. You can tell this graph actually needs a 50 basis point drop-off of the right side of that because they just cut rates last week. So this looks very like the CDF function, for example.

So there's another curve model that I think is probably pretty much-- pretty widely used at this point in the industry, which is I'm going to build my knot points based on the central bank meeting dates, because those are the only dates that rates are going to change, in general, or going to change materially. And I'll assume it's flat in between.

So that's one. Again, I think that started six or seven years ago. And it's pretty-- far as I know, almost everybody uses this methodology now, where you have generally the central bank will publish the next two years meeting calendar. If you really want to, you can apply beyond that. But two years is about the stuff you know. And then we'll assume, say, a spline beyond that.

You end up with a little bit of issues on, how do you choose your instruments? There's eight meeting dates, eight central bank meeting dates a year, four quarterly futures a year. So you have a slight under constraint issue here. So you need to fill that in somehow. You can include the one-year-- or sorry-- one-month futures, smoothness constraint. You have to do something else. But that's fine.

Also recently, in the last couple of years, actual swaps between meeting dates have started to trade. So basically, this is a bet on the implied rate after this meeting date. So if you think the Fed is going to cut it the next meeting, which I think is early November, you can buy a swap from November through, I think, mid or late December at this rate. And that's essentially the implied rate after the meeting.

So I don't know what the current-- I think-- probably thinking is they're going to drop 25 basis points. Fine. They surprised me a little bit by dropping to 50. Fine. But there's instruments that will directly probe that and reflect that rate. So that's also something you can do. So I said, we can make a hybrid curve, and this works pretty well.

Lots of other stuff in curves, and this is still a very active area of research. There's other features that you sometimes want to account to, for example, turns is an interesting feature where there's dates that are significant accounting dates, year-ends, for example. Generally, corporates will post their earnings and their balance sheets, return on assets, on December 31 or quarterly.

So what that does is it puts a premium to borrow money over that rate. If I'm a bank, I want to minimize my balance sheet when I'm reporting what it is because I'm going to be judged by the investor community and things like return on assets, on total assets, and things like that. So I don't want my assets to spike over year-end.

On the other hand, money is money. And if somebody wants to borrow money over year-end for some other reason, I'll do it. But I'm going to charge them a premium for that. So you find the curves actually-- if I actually look at a curve, it might be like doo, doo, doo, doo, doo, but then I'll have a big spike every year. So maybe that's one year, where this essentially is the premium for borrowing over year-end. So those you can include in curves as well.

And you'll find some of the hedge funds, in particular, are slightly nasty. So they'll do a trade where they'll want to do one swap to this date and the second swap opposite direction to that date. So they're really just picking out whatever curve-building you're doing over year end. And it's very easy to get that wrong. And you do it in enough size, and it's real money.

And the other way they-- also, other approaches you can do, for example, do a PCA method. One of the things on the curve is I may build a curve. Typically, it's about 50 inputs for \$1 curve. There's widely varying liquidity among all those 50. Typically, there's going to be about a half dozen that are the super liquid points. And then the rest of the stuff will slowly become less liquid.

So you can, if you want, say, look at historical data, work out what my factors are, pick the first half dozen factors which represent shapes of curve moves. Have those driven by my most liquid stuff, and use that for a curve. That works quite well in bond world, for example. So lots of other possibilities.

Talk briefly about hedging. So one of the things you want to do if you're running a trading book of some sort or asset management or something, you want to know, what is my risk? And risk, generally, in the trading desk view of the world is, how much do I make or lose if my, say, rates go up or down? So I'm going to use just simple PV01 is my main risk measure, which is just a sensitivity to one basis-point increase.

Nothing too-- I've got a two-sided difference here. Just PV at rates plus a bit minus PV at rates down a bit, normalized by two times the size of that. Do this for each parameter, and I get a delta strip. And that's typically what I'm going to use for hedging. That's what I'm going to have limits put on me from risk management. That's what I'm going to use, and an assortment of other things. Very standard approach.

And you can do that for any parameter you want. You work out just what is my sensitivity? By generally, a brute-force bump in reval. Change the parameter reprice. Take the difference. You can get a bit cleverer than that if you want to go to analytical hedge computations and things like that, but the core method is still the same. So that's first order.

I can also-- depending on my product, everything has got a little bit of convexity. And depending on what my trading style is may or may not matter. So I might want to do a gamma, which is my second order risk. Generally, the change in PV squared by D rate squared.

Many definitions of gamma. And at least in equity world, it's a little simpler. In the interest rate world, it's a little more annoying because interest rates, I said, you have a yield curve that might have 50 inputs on it. So my delta will have 50 computations. Gamma is going to be that squared because I need to bump each of them in pairs. 2,500 computations is very expensive, so you probably don't want to do that.

So again, you can do a parallel gamma or factor gamma, a number of other approaches, or just use scenario analysis. Lots of ways to get your higher order terms in there.

As soon as you start actually hedging in practice, you hit a whole lot of practical issues that you had to be aware of. How much is the higher order terms contribute? You can have third order terms, fourth order terms, cross terms, and so on, all sorts of weird complexities.

How do I choose my bump size? The standard is a basis point, but that's really just for reporting. And this depends a little bit on your numerical method and what you're doing. So if I'm in a Monte Carlo world, for example, I have a Monte Carlo pricing engine, I probably want to bump a little bigger. Because there's Monte Carlo noise because it's random numbers. So it's a little noisy. So my bumps are too small. My difference is going to be quite small.

Some of the other methods, for example, some of the PDE-based methods, if I'm doing a large bump, I can see grid effects starting to hit it. So there's a lot of numerical issues here, and everyone's a little bit different. So there's some sequential bumps. There's all sorts of little convexities come in.

And I said, you're exposed to your interpolation method. If I'm doing a local interpolation, my hedges are concentrated around my maturity. If I'm doing a global one, they're going to be more distributed. So this was something that, again, you get into lots of arguments with people about pretty quickly.

One of the other things to do, depending on your jurisdiction may be a regulatory requirement, is to do your P&L attribution. So I come in the morning. I've got a set of positions, set of whatever market is whatever.

At the end of the day I go home. I've made or lost some amount of money. So one of the key questions is, why? Where did that come from? Am I a genius trader? Did I get the market right? Did I just have a lot of client flow that I took a bid offer on? What happened. It's also a very good check on how well your model is working, which is interesting.

So you know the P&L, I said, from yesterday close to today's close, because you can value everything. So how did you get that? Generally, there's three main components. And this is also something that you will have arguments about until the [INAUDIBLE] death of the universe on, particularly on the theta component.

You have a change due to the portfolio. So did I buy and sell trades? Did I do unwinds? Did things get assigned to me? Do I have resets? Did I exercise options? All of that will have an impact.

There's the passage of time. That's the complicated one. So I went home yesterday, and I came in today. Market hasn't moved. But what does market not move actually mean? It's a very subtle effect because interest rate curves are a function of time versus this.

So when I move this date to that date, what do I hold constant, and what do I let change? And there's more than one intellectually defensible answer to do this. So anytime I tell people you do-- be very precise. Make sure you know exactly what you're doing. And again, there's lots of defensible answers, but you all have to agree on your notation.

And then finally, there's the more common one that people think about, which is, how does the market move? So I'm long a billion bonds. I got a PV01 of-- what is that-- \$10 million and one, or whatever. Rates move five basis points. I made \$50 million. Well done, me. Rates went down three basis points. I lost \$30 million. Less well done, me. So market impacts easily.

And I said, the portfolio change is also one of these ones. It's mechanically just a pain in the ass. You get introduced to booking systems and all your operational flows and things like that, which can confuse things a lot.

Finally, just still bouncing between subjects, so we still got a few more minutes. Clearing is another thing that came in post-2008. One of the features in 2008 is people became acutely aware of their counterparty exposure.

So I do a trade with you, and we do it at OPV initially. Time passes. Now it's worth \$10 million to you. Great. This is not realized money. This is money I owe you in the future. If I go bankrupt, you lose \$10 million. So people become very aware of-- and if you go bankrupt, I get a \$10 million freebie, assuming I can argue with your bankruptcy judge. Fine.

So there's people are very acutely aware, and Lehman really brought that home because they were a huge dealer. They had many, many billions, hundreds of billions of dollars of derivatives with people. It, by the way, is why JP Morgan is doing so well now because they were the only bank still standing in the fall of '08. So people had to replace an enormous amount of Lehman hedges, and JP was the only one still actively trading.

Goldman's was in trouble. Stanley was in trouble. Bank of America, where I was, had just bought Merrill. Merrill went bust. All the-- Citigroup was having all sorts of issues with guaranteed credit exposures and things. So JP Morgan just made out. The bid offers went up by about a factor of 20 for a month or two there, And They just raked it all in. Good for them, I guess, but not good for anybody else.

So one of the things the regulators decided after that was people would have to use clearing for vanilla swaps. The clearing is a concept that's been around in futures forever, but in swaps it's relatively new.

So what that means is, say, I'll do a trade with you, and we'll do a trade. Fine. You guys, we agree on the terms. We will then both assign our side of the trade to a central clearing house, say, a CME and LCH. London Clearinghouse and the CME are the two big ones.

So then both of us now face, say, the CME. And you and I are now separated, and we don't care what happens to each other. You can go bust, and I'll be sad, but I'm not losing money. So that came in.

And clearinghouses are generally big. One of the worries is we've traded off small and large potential disasters for a very small number of enormous disasters. If LCH goes bust, it's time to get guns and ammo and head for the hills because it's going to be just a complete disaster in the financial world. They are so big and so integrated with everything. And all the members post various margins to them and membership fees. So they're pretty robust. But if something did weird happen, it would be quite unpleasant there.

Now, one of the things about the clearinghouses is that dealers are required, but you end up with this now situation where \$1 I'm paying to LCH is not the same as \$1 I'm paying to the CME. These are not fungible amounts, even though they're both dollars. So it's quite interesting because I can't-- if I have a swap with LCH, and I swap with CME, I can't offset these two. I can in my books, but not in terms of operationally and collateral and everything.

So what happened probably eight or nine years ago, it was quite a significant spread opened up between these two. So this is from 2017. I guess Tradition is one of the big brokers. This is the difference between if I want to pay fixed on CME versus pay fixed on LCH.

So again, if I'm doing a 30-year swap, it's now 3 and 1/2 basis points between them. And again, this is a market where the bid offer is 0.2 of a basis point. So the question is, why did this-- and this persisted for a long time. It's actually flipped sign recently. But this was around for years.

Anybody want to guess why this kind of-- no. Why can't I just arbitrage this away? I'll go receive fix in CME. I'll pay fix in LCH, and I'll collect 3 and 1/2 basis points and go buy my island in the Caribbean. Why can I not do that? Or why did nobody do that? Anybody want to hazard a guess even?

It's quite interesting. And it's one of the pretty core things in a post '08 world, which makes things a lot harder. No guesses? Yeah?

AUDIENCE: Like a zero spread.

ANDREW Pardon?

GUNSTENSEN:

AUDIENCE: Like a zero spread. Like, why are [INAUDIBLE] something. Like, you receive versus on your short, it's different.

ANDREW But why does that persist? Why can't I just-- there's enough people out there that are happy to make money.

GUNSTENSEN: Why can't they just take that out? If you told me I could make 3 and 1/2 basis points running for 30 years, I'll do that all day, and that'll make me a ton of money.

There's a couple of reasons. You have to post margin on both sides. The clearinghouse, as part of doing clearing, require you post first swap. It's typically a percent or 2% of the notional. So you have to post money to them on both sides. You're now double posting margin. That costs you, and that's the main driver of it.

Some of the other features are other products trade on these clearinghouses, and you can net them. For example, CME is the big futures counterparty. So if I'm long a ton of interest rate futures and short a ton of swaps, I get at least partial netting in my margin between the two. Then there's a couple of non, less economic things.

This thing was largely driven by some very large US asset managers, people like PIMCO, who is ginormous, that wanted to pay fixed all day to hedge their loan portfolio. They are a little less trusting of the LCH because the LCH is owned by a consortium of dealers. CME is owned a private company. So their thought was the LCH is going to be more favorable to dealers because the dealers own them. So if it comes down to a legal issue or something, they're going to be biased.

The other one is they want it to be in US jurisdiction. LCH is in London, L being London. And they just want-- they want to deal with the legal issues in the US because that's where they have their expertise and where they're based. So a couple of things. Again, it's quite interesting.

But one of the things this introduces, the fundamental thing, is that because of the margining requirements, there's a spread here that can exist between things that would normally be arbitrageable. And it can persist. Fx has all sorts of these, too.

And this is coming to the point that it actually to correctly and completely calculate the prices, is if you need to include a lot of things outside of the core pricing of the instrument, how are you going to margin it? What capital do you have to post it? Or is there any other kind of things like this that are going to cost you money that have not traditionally been accounted for?

And I would say one of the huge quant efforts over the last 15 years is to work out what are all these extra costs, and how do I push them back into the upfront pricing? If I do a swap for a very small P&L with a counterparty, and then it turns out I have to post an enormous margin that's going to cost me way more than that, I've just done a money-losing trade.

So I don't want to do that. You can't do too many of those, or you're not trading anymore. So you have to correctly account for all of these extra costs, and that's quite non-trivial, just both mechanically in their system side, as well as conceptually. And a lot of these things are portfolio effects. So how do I then allocate that to individual trades is not necessarily super obvious.

Good. Coming to the end at about the right time, quick blurb about electronic trading, which is-- if any of you guys get into this area in your careers, this is one of the very quite interesting, very hot areas. Electronic trading has been around in equities for quite a while. My standard metric is equities is probably 10 years ahead of foreign exchange, which is 10 years ahead of fixed income.

And one of the graphs there is from one of the data providers-- the percentage of markets in some of the fixed income products that are traded electronically-- two, basically. There's electronic trading, which I set on my keyboard, and I go. Computer does everything. Or I can pick up the phone, and I can call somebody-- voice trading.

So again, there's a spectrum here. On the right, I've got the very liquid stuff-- US equity products, cash equities. Those are all 90%-plus Fx spot-- all traded electronically, basically. And then I start to get less and less. European governments, US governments-- both predominantly electronic. But we're talking 60% to 80% of the market here.

What you tend to see is a lot of the smaller trades we'll go through electronically for a couple of reasons. These are not market moving trades. So people are happy to let the machine do its thing.

And if I'm buying \$1 million of a Treasury, that's not going to affect a market that trades \$1 trillion a day. If I'm buying \$100 million, maybe if I'm buying \$500 million, I need to be careful because that is going to move the market, and it's going to move the market against me, and I'm going to end up paying more than I would otherwise do that.

And the other thing is I get down there. So I have some of the investment-grade corporates, high-yield bonds. And I get down to the left, which is really not electronic. It's things like emerging market bonds and municipals. Again, munis is one of these ones-- quite an interesting market.

Before I was at Mizuho, I was at TD, Toronto-Dominion. We acquired a muni bond electronic trading firm from Chicago, which was tremendously successful. And they have about half of the electronic market in munis. So munis is something that there are huge number of bonds. There's something like 3.5 million issuers in the country, which is the states, the towns, bridge projects, school boards, all sorts of things that issue municipal bonds. They're all a little different.

So there's a lot of nonstandardization in the actual bond structures. And the other thing, there's a tax advantage. I don't know if you guys-- any of you guys have a PA at this point, or hopefully in the future, not too much longer. Muni bonds, you don't pay federal income tax on them, which means typically, the yields are about 70% of a taxable issue because you get the same net after tax yield.

This means it's a retail market. Institutional investors don't get the tax break. So they are less interested in munis. Retail investors typically are not buying things in blocks of 100 million. They're buying things in blocks of \$5,000 or \$10,000, something like that. So the average trade size is very small.

So at TD, they did something like 5,000 to 10,000 trades a day. But the average trade size, I think, was \$28,000, which in a government market is absurd. That's not even an odd lot. That's something the janitor picks up at the end of the day and just does something with. So it's quite interesting. But on the other hand, you do enough of it, you can still make good money.

Emerging markets is another one where the bonds are very idiosyncratic, and there's a lot of very special stuff. There's not really a generic Chilean bond. You have to know exactly what it is, how is it paid, what's the credit support, all of that. So these things are generally done over the phone with somebody because not for the faint of heart, though, I would say.

But as time goes on, things are slowly moving up here, and things are becoming more electronic. Electronic is just-- it's faster. It's cheaper. You're less sensitive to people taking advantage of you. You have markets. You can see where the market is. There are quite a variety of participants here.

You have the trading shops. There's people like XTX or Jane Street or Hudson River Trading or folks like that. They don't have customers. They just trade on their own account. They're very good. They're very fast. They generally make markets in a lot of this.

Telling Peter a story I had at some point. I was at a Thanksgiving dinner or something with a big group of people. And one of the other people there had just graduated from college with a PhD in physics in microwave guides, and he'd been employed, I think, by Jump Trading to work out the optimal microwave transmission to reduce Chicago-New York latency. Because Chicago is where the exchanges are, and here's where the customers are.

And I just remember talking to that guy and thinking, Jesus-- I think I was at Bank of America at the time-- we are never getting in front of this. These guys are just-- they are so advanced in this stuff that you have to be very careful to pick what you're doing. So these guys are super fast, super good. This is the high-frequency traders, people like [INAUDIBLE].

Then you have the dealers, which are people like the Bank of Americas, or TDs, Morgan Stanleys. They generally are trading. They're pretty good. They're fast. They're not building their own microwave towers. So they're doing that. They have customers. Customers are good. Customers want to trade with you. The customers will pay you a bid offer to do trade with you. So you like that.

And there are of low latency, but microseconds and not nanoseconds. And then there's the customers who just want efficient execution and to get sense that they're not being taken advantage of, and that are seeing a full market, and they're able to decide what they want to do. And there's a bunch of other guys, but those are some of the main market participants here.

So price making here, at least on the dealer side, which is what I'm familiar with. I've never worked at a prop shop, so I don't know if those guys are much different. We start by getting a good mid-market price, either from liquidity providers, if it's something that trades actively, or from a model, if it's something that's a little off the run.

You're then going to skew the mid-market price to what you're actually going to offer the guy based on a whole lot of stuff. Maybe risk-- what is your current position on this product? If this is a risk-reducing trade, you're willing to be more aggressive just to get it off your books. If this is a risk-increasing trade, you're generally going to want to be a little more conservative because it's going to cost you to manage it.

Who's the client? Is this a good client that it's important for you to keep good relations with, you have other transactions as the firm, in which case, you're willing to be a little more aggressive with him.

Is this a toxic client? That's quite an interesting study of how do you evaluate client toxicity? And the classic toxic client is, again, is some of fast money guy who will put out 10 requests to 10 different dealers for the same trade at the same time.

So if you execute that, you'll then find nine other guys have executed the same. So your hedge has moved away from you. So it's going to cost you more money to hedge, and you're probably losing money on the trade. But how do you determine if that's likely to happen or not? So a lot of market microstructure things.

Also, volatility. If the market is in a particularly volatile period, you're going to want to be a little more conservative in your pricing, widen out your bid offer because the market may move away from you before you can hedge.

And then there's a whole world of prediction models, and how do you look at that? And what's your estimate of where the next tick is up or down and things like that. And there's all sorts of market microstructure there. That's a quite interesting area you can do.

And generally, the trading requires very good integrated systems. You need to have everything talk and everything super fast. You can't have-- one of the banes of our existence at a previous firm was the credit check.

So I'd get a request from somebody to do a trade. I had to determine if we had open credit lines with them. By the time that answer came back from our credit systems, the deal had gone away. So that was a bit annoying.

I think I'll end there and just open it up for any more questions. I think we've got a couple of minutes still, so I'm happy to yack about anything, as I said. So otherwise, I think we'll call it a day there. Skip through some bonds.

AUDIENCE: I guess like one of the main things that people have been talking about [INAUDIBLE] is the interest rates in the recent quarter. What is your opinion on the implications that changes in interest rates can have? And just kind of understand what type of considerations are there going to be?

ANDREW GUNSTENSEN: The biggest one from the guys actually managing interest rate policy is really inflation. We've had a burst of inflation here driven by things like the COVID funds. We pumped an enormous amount of money into the economy to stop things collapsing in COVID. One of the ways that showed up was inflation. Russia invading Ukraine also popped energy prices. So that also popped inflation. So that's one of the big ones.

Now, the downside is the Fed's trying to balance that with employment. Generally, if rates are high, people are not willing to borrow money because it costs you more. If I'm going to start a business or something, I'm going to want to, quite often, borrow money.

The other huge connected thing in that is home prices. When you look at in my lifetime, what have I spent money on? My house is overwhelmingly the most expensive thing I've bought. And then also, every time you move, you buy furniture and paint and a new roof and this. So that's also a huge driver of economic thing.

So I think in terms of we've just, I think, just kicked into an interest rate cutting cycle. I said a few weeks ago they dropped it. There's some arguments that they kept the rates too high for too long. But I think they've started to cut.

So I think the kind of things that you would look to see is things like mortgages, home sales activity, business startups, stuff like that. It's quite pervasive in terms of how much-- how expensive is money is a huge driver of a lot of activity. I don't know if I answered your question or not, but OK.

So anything else? Nope?

AUDIENCE: Just one question with the fixed income modeling. I would guess that forecasting or predicting rate changes would take a real back seat to hedging.

ANDREW Well, again, that's a subtlety--

GUNSTENSEN:

It's a-- there's a subtlety--

--a subtlety difference here. There's two kinds of fixed income research. There is the economics client-facing research, which is a lot about that stuff. Here's my prediction for the economy. Here's if the rates cut, this is what I think you should do. Here's the positions.

And then there's the internal area, which is I am, which is more about pricing and hedging of transactions, which we don't make any predictions. We don't care. We just want to know how to hedge it. And at that point, we're good.

So they're quite different in outlook, even though they're both fixed income research. And I have no opinion on where rates are going. If I did, I wouldn't be standing in front of you here. I'd be sitting in my Caribbean island, with the piña coladas. [LAUGHS]

AUDIENCE: Could you also comment on the municipal bonds, what the analytical challenges are there?

ANDREW So a couple of the challenges there are, again, there's a set of bonds that are very idiosyncratic, maybe the
GUNSTENSEN: project bonds. So this is a bond where the revenues come off of, say, tolls on the Pennsylvania Turnpike. They'll issue bonds supported by that.

I don't know about you. I have no idea on the traffic in the Pennsylvania Turnpike and what that's like. Could it be-- so a lot of these things got a big whack in COVID, for example, because nobody went anywhere for a year. So there was no revenues to support them. So they all got [INAUDIBLE].

A lot of other munis, it's about what do you think future tax policy is? So let's suppose the average federal rate is, I don't know, 30% right now. So that means I'm willing to take 70% of the Treasury rate for a muni bond.

When I look forward, do I think tax rates are going up or down? I have no idea. I would suspect up, but you never know just given the deficit situation. But who knows? So there's that kind of issue.

A lot of the other issues more on the trading and the quant side is, how do you figure out who's got them? So I want-- I have somebody who wants to buy a bunch of AA New Jersey bonds. How do I source those? So that's a lot of data.

That's one of the reasons munis are still very voice-driven is the sales guys know where the bonds are, so they can get them out for you. Whereas, the machines just have no idea who's got the bonds right now. So it's hard to source the product at that point. So that's some of the complexities.

SPEAKER: Well, thank you very much.

ANDREW Thank you very much.

GUNSTENSEN:

[APPLAUSE]