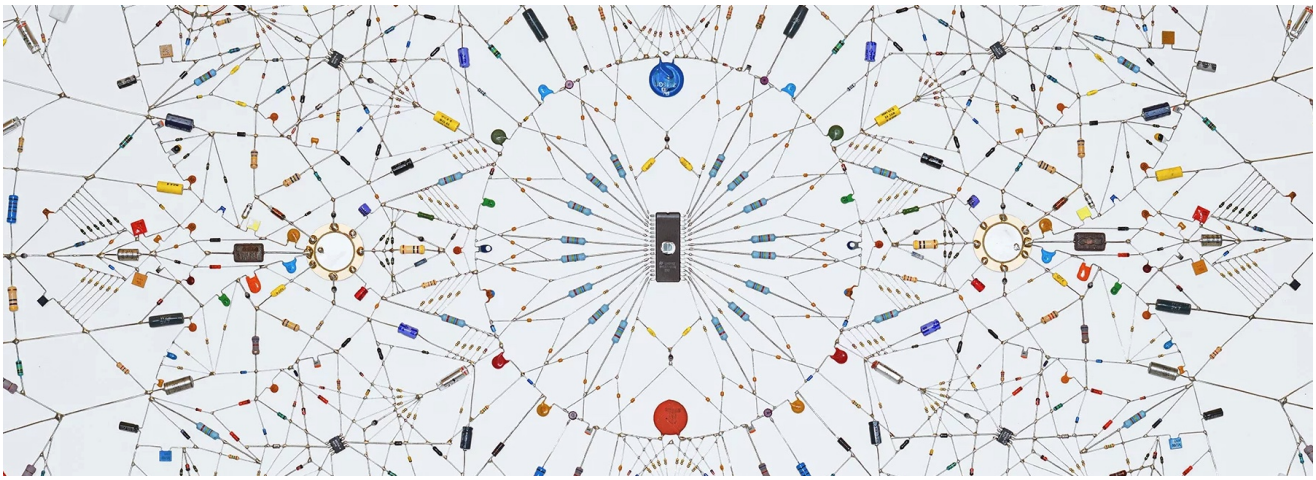

# Market Structure Primer

Get Started!

The purpose of this document is to provide an overview of the US equity trading landscape. It started as an onboarding document to help new hires who aren't as well versed in market structure. It can be read start to finish in order, or you can jump around. The visual Life Cycle of an Order also provides a handy view of the various events and entities involved in the process of a trade - from conception to completion. For those with little to no prior knowledge, the three core sections on market participants, mechanisms of communication between participants, and a birds-eye view of overall market activity should provide a working understanding of trading in the US equities market as it is today. Over time, we will be adding material that delves more deeply into further topics, especially into market stucture history and the conflicts of interest that evolve along with market structure.

# Introduction

The US stock market is the complicated machine that allows people to invest capital in public companies and unwind those investments as they see fit. There are many types of participants that work together to make this possible, including exchanges, brokers, market makers, and regulatory bodies. These participants often have competing interests that drive their actions — some are eager to advance the ecosystem forward, while others prefer to maintain the status quo. Because there is so much money at stake, whenever a policy change is proposed, there are a lot of vocal opinions in the room. As such, change in the stock market is slow.

This primer aims to provide a functional understanding of who the participants in the US stock market are, what goals drive their behavior, how they communicate and interact with each other, and how the overall market behaves and evolves as a result. This is intended to be a living document that will improve and expand over time.

You might ask: "Why is this all so complicated? I don't need a multi-layered ecosystem of stakeholders and middlemen to interact every time I buy something online." Well, actually you probably do - but the hidden realities of retail manufacturing, pricing, logistics, and shipping are a story for another day. Stock market complexity is perhaps more suspicious because there is not a physical product at the heart of it. If the "ownership" of a share of a stock basically corresponds to a database entry somewhere, then what gives? Why are there so many kinds of financial intermediaries?

To address this question, we have to confront some of the implicit assumptions and invisible shortcuts in popular conversations about the stock market. First of all, the phrase "the stock market" is already a bit misleading. It conjures up an image from times gone by of a rowdy trading floor where red-faced men in suits holler and yell. But "the stock market" is now a collection of roughly 40 different for-profit companies that offer the service of matching buyers and sellers of stock. Thirteen of them are "exchanges" and the rest are "alternative trading systems," sometimes referred to as "dark pools." Phrases like "the stock price of Microsoft..."

are also misleading. In a particular place at a particular time, there are actually two prices: the price that buyers are willing to pay, and the price that sellers are willing to accept. And when we write it out like that, we start to see the dirty edges of the complexity being swept under the rug and suspect that there is more: probably not all the buyers agree on the price they are willing to pay! And probably the price for even a single would-be buyer depends heavily on the quantity of stock being transacted. And probably all of these buyer/seller prices change dynamically over time based on what's happening elsewhere, and probably the mechanisms for expressing those prices are imperfect and sometimes incentivize buyers and sellers to hide their real preferences to various extents, and probably the mechanisms of arbitrage that should keep prices in line across different trading venues are also imperfect and sometimes incentivize behavior of questionable value ...

We could go on (and we will!), but you get the idea for now. Suffice it to say, that for every "simple" decision to buy or sell stock, there is a laundry list of sub-decisions: what stock should be bought/sold? (there are over 8000 choices). What quantity should be bought/sold and what is the ideal tradeoff between speed and the cost of implementing the corresponding transactions? Where in the space-time continuum of 50 venues and 23.4 million milliseconds of regular trading time each day should these transactions take place? Each of these decisions requires a different specialized expertise, and perhaps it is not ultimately surprising that they are often made by different people. Naturally the incentives and rules that shape the behavior of these different people often lead their interests to diverge from one another. And then of course we wonder: who should/does have control over these incentives and rules? And what incentives and rules are they subject to? And we haven't even mentioned yet the decisions and rules that a company must navigate to offer stock in the public market in the first place. Or what happens when buyers and sellers don't deliver what they promised on time, and many other crucial details.

There is one popular view of the stock market, that it is an antiquated behemoth of a system with outdated technology and bloated structure that is long overdue for disruption. This is in many ways true. There is another popular view of the stock market, that it is a highly evolved, almost living thing - absorbing the lessons of the past and growing stronger and more efficient over time, perhaps to a point where to tinker with it externally is damaging. This is also in many ways true. The complexity of the US stock market arose historically as a patchwork of responses to real problems and real needs among participants, and it is naive

to ignore the value of imperfect but functioning solutions to real problems. But just because a system evolved in response to real needs does not mean that the long period of its evolution has lead to an optimal or even acceptable status quo.

If we (speaking as a society now and not just as Proof!) are going to help shape the stock market's continued evolution in positive ways, we should first understand how the market functions today, what competing interests have driven and continue to drive its evolution, and what effects (both positive and negative) its current structures have on various participants. Only with this context and a clear articulation of our goals can we meaningfully evaluate whether a certain proposed action is likely to have a positive impact.

We at Proof obviously have specific opinions about how the market could be better, and we will discuss these in the later (planned) sections of the primer. At that point we will have developed sufficient context to explain our goals and why we think certain actions/reforms are helpful in achieving those goals and why we think others are not. In the early sections, we will be sticking as closely as possible to relatively factual descriptions of the relevant stakeholders and their mandates so we can build a shared factual context before venturing into more subjective territory.

If there are topics along the way that you are curious about and want to see a more in depth treatment of, or topics that aren't covered at all that you feel should be, please let us know!

# The Life Cycle of An Order/Trade

This graphic illustrates how the various participants in the stock market interact to form the lifecycle of an order placed by an institutional investor, from its inception through to any trades that result. At the periphery of the process (outside of the stock market itself), end investors give money to an investment manager. The portfolio manager decides what stocks to buy and sell, and in what quantities. When the portfolio manager decides to buy/sell a stock, they communicate this to a trader, typically through a piece of technology known as an order management system (OMS). The trader makes high level decisions about how the trade(s) should be accomplished, like what broker(s) should be used and what algorithms. The trader's instructions are then communicated to a broker-dealer, typically through a piece of technology known as an execution management system (EMS). The broker consumes real-time market data from various sources, and uses this to make low level decisions about how best to implement the instructions in current market conditions. The broker produces "child" orders derived from the "parent" order instructions, and routes them to trading venues such as exchanges and dark pools. On these venues, the child orders may interact (either immediately or after some waiting) with orders submitted by counterparties like themselves or by proprietary traders, so-called because they trade on their own behalf, not on behalf of external end investors. When two orders on a venue are matched in a trade, the venue sends the relevant information back to the brokers/proprietary traders, and also reports it to the "tape," meaning that it becomes part of the market data that others consume.

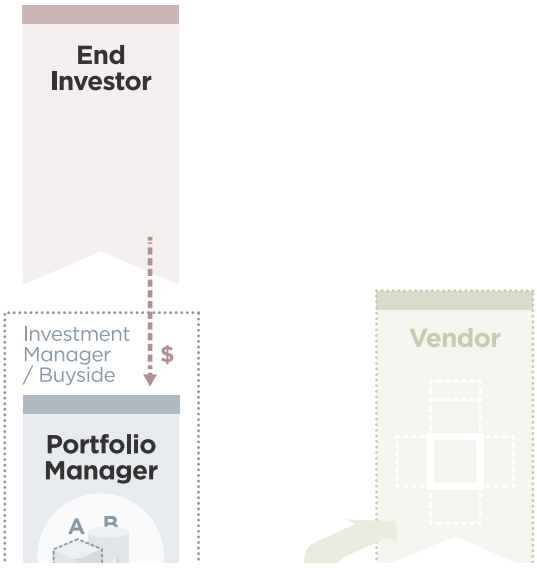

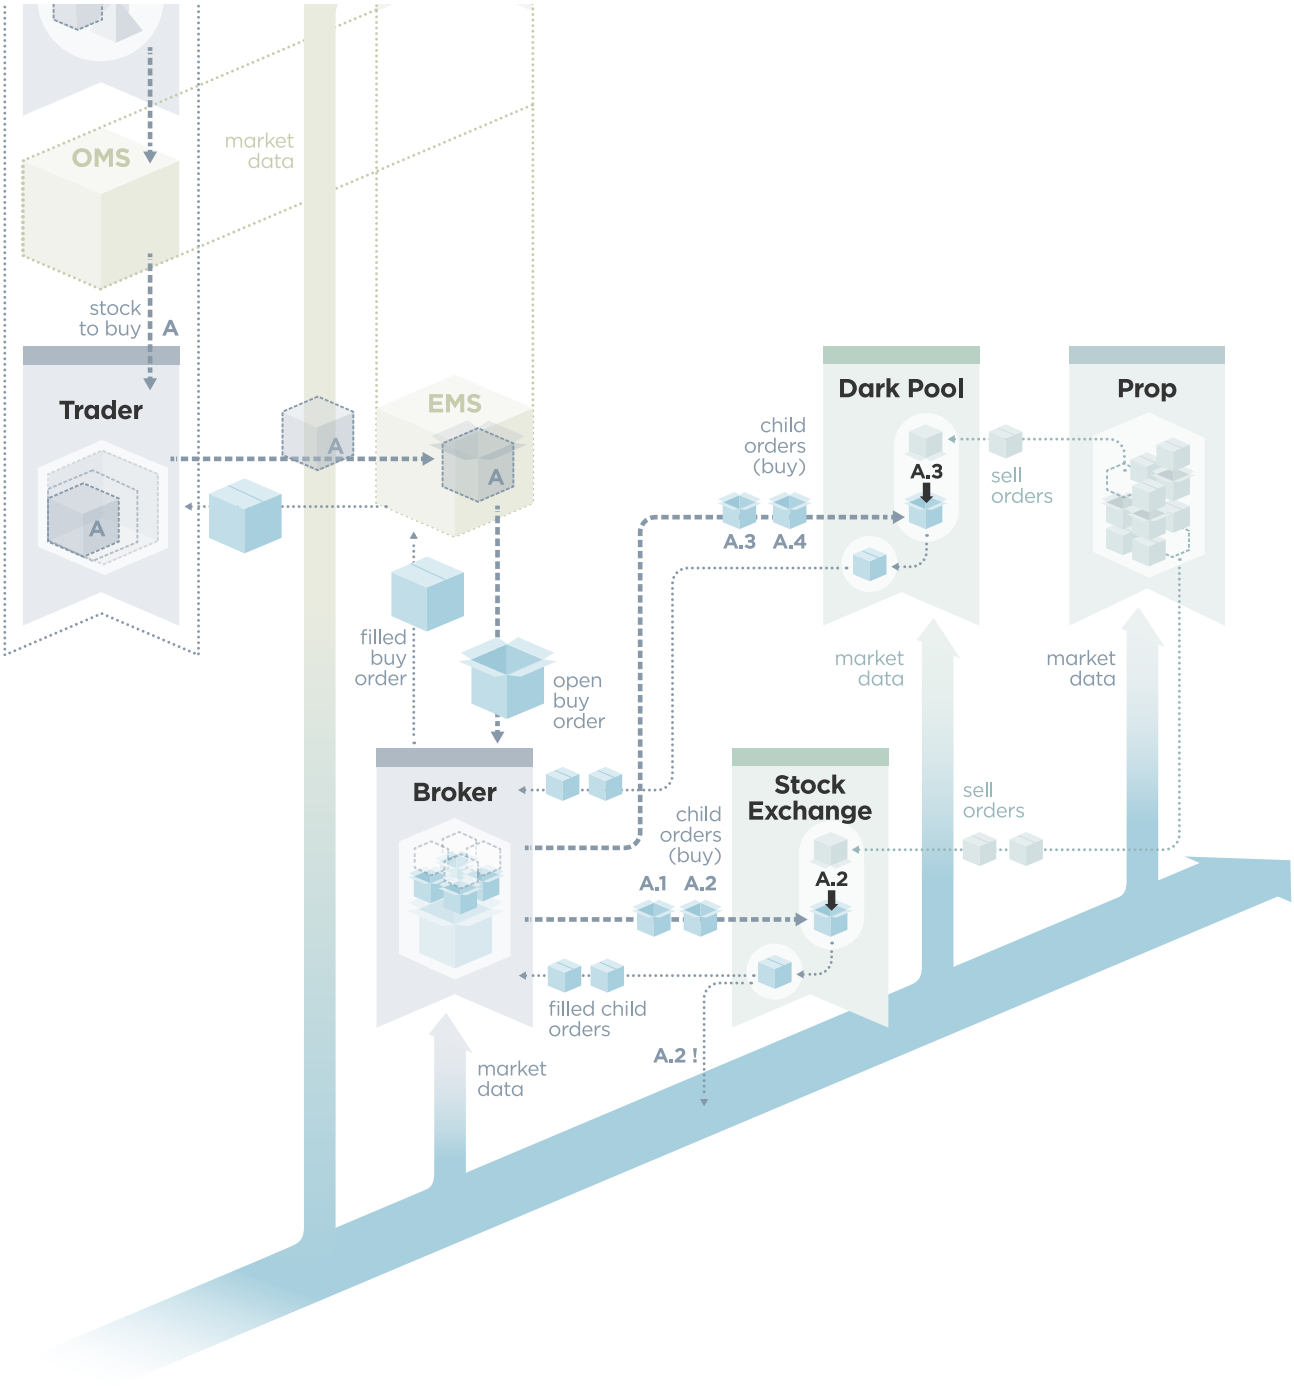
OMS
market data
stock to buy A
Trader
A
EMS
A
A
Dark Pool
Prop
child orders (buy)
A.3
sell orders
A.3 A.4
filled buy order
open buy order
market data
market data
Broker
Stock Exchange
child orders (buy)
sell orders
A.1 A.2
A.2
filled child orders
A.2 !
market data

# Market Participants

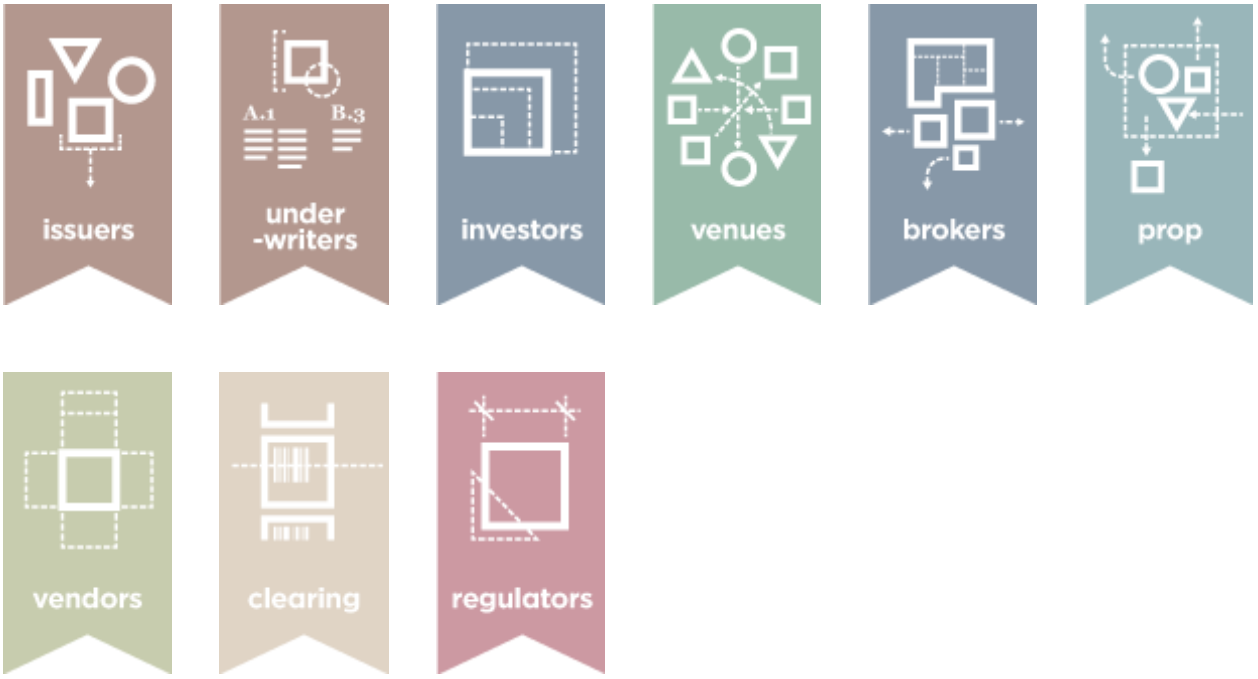


## The Who and the Why

There are many different types of participants in the US equities market, each with distinct needs and objectives. Here we will focus on participants and trading mechanics involved in institutional trading. We will define our unit of a "participant" in terms of a single role, which does not always correspond smoothly to the often-used unit of a company. Within a large financial firm, there may be separate lines of business that fulfill separate roles, and treating the resulting mesh of sometimes conflicting incentives and functions as a "participant" can obscure the underlying mechanics driving behavior in the financial system. Conversely, considering only individual roles neglects the higher level constraints of company ties and relationships that can be needed to make sense of the lower level behaviors. To avoid both sources of potential confusion, we'll define and discuss individual roles as our base units first, and then periodically zoom out a bit to discuss the inter- and intra-company relationships between roles that comprise the overall financial ecosystem.

# Issuers


Obviously, the stock market wouldn't exist if there were no public companies. It's common in the stock trading world to conflate the companies that issue stock (aka issuers) with the stock shares themselves - in fact, the lingo is built on it. It's normal to say "the price of Microsoft went up today" rather than "the price of a share of Microsoft went up today," even though this wordier version is more technically accurate. [Full disclosure: Allison used to work for Microsoft, so all hypothetical stock pricing examples tend to be about Microsoft.] But it is worthwhile to mentally preserve the distinction between issuers of stock and the stocks themselves, as the issuers are entities who have agency and goals of their own.

So what does an issuer of stock want? When a stock is first offered publicly in an IPO (Initial Public Offering), the issuer's goal might be to raise money and/or to allow early private investors to cash out. As a stock continues to trade publicly, an issuer's goal might be to increase or maintain its market captilization, which is the total dollar value of its outstanding shares. Or its goal might be to increase or maintain its share price, which is subtly different. The market capitalization is equal to the number of outstanding shares times the price per share, so the two metrics are certainly related, but they are not quite the same thing. The market capitalization can change because the share price changes, or because the number of outstanding shares changes. The issuer might increase the number of outstanding shares by issuing new stock, or might decrease by the number of outstanding shares by buying back shares for itself that were previously in public circulation. This is called a stock buyback.

There are reasons why an issuer might not want the individual share price to be too high or too low, regardless of the market capitilization. Very low prices per share might violate rules

(like the listing standards of a given stock exchange - more on that later) or lead to an impression of low quality. High prices per share might have a dampening effect of trading, since the share is an indivisible unit and partial shares cannot be traded. An issuer may make it a goal for trading in its stock to be reasonably active, so that its investors don't get worried that trading will dry up and they will be stuck with shares they don't want in the future. Shares trade commonly in "round lots," which are typically multiples of 100 shares. To keep these typical units of a palatable size for active traders, the issuer might decide to do a stock split or reverse split to influence the price per share without necessarily changing the market capitalization. A stock split takes 1 share of stock and converts it to multiple shares (for a specified multiple), thereby lowering the price per share but raising the number of outstanding shares proportionately so that the market capitalization stays the same. A reverse split combines multiple outstanding shares into one, thereby increasing the price per share but lowering the number of outstanding share proportionately.

To sum up, an issuer is an entity whose stock is publicly traded. An issuer may want to control its market capitlization, its price per share, and/or the quality of trading activity in its stock. Some of the main tools it has at its disposal to accomplish these goals include issuing new stock, buying back existing stock, and stock splits and reverse splits.

# Underwriters

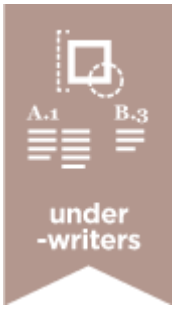


When an issuer plans an IPO, it typically hires underwriters to be responsible for selling the initially offered stock. Underwriters typically make some kind of committment to sell a particular amount of stock at the IPO price. Often the total offering is split among several underwriters, who collectively form a syndicate. The immediate goal of an underwriter is to sell the initial stock and collect the corresponding fees. A more long-term goal is likely to position themselves for future underwriting business. Since underwriters are typically attached to firms that offer other financial services, they may have additional goals like pleasing the firm's other clients by offering them access to valuable IPOs.

# Investors


The term "investors" is a rather vague term, often used in different ways by different people. In the private equity world of companies whose stock is not publicly traded, the term "investors" typically refers to entities who purchase and hold shares in the company. But things are simpler in this context, where it is usually fine to conflate the entities who make the decision to purchase with the entities who then formally own the stock, because it is typical for these to be the same. In the public markets, this is much less true. If you as an individual person purchase 100 shares of Apple (a publicly traded company) in an individual brokerage account, you probably assume you count as an "investor" in Apple. But is your name actually on a stock certificate somewhere saying you own 100 Apple shares? Does Apple know who you are? Not necessarily. Some providers of brokerage accounts put their own names on the stock certificates and hold them "for you."

Things get even murkier if you pool your money with others in a retirement plan, say, and the plan manager decides to purchase shares in a mutual fund, and the mutual fund decides to buy some shares of Apple, a publicly traded company. In this case, who is the "investor"? You? The retirement plan manager? The mutual fund? All of the above? You don't have your name on an actual stock certificate for Apple, but a series of contractual links between financial entities approximately allows you to feel as if you do. And if things work as they should and Apple's share price goes up, you should get a proportional benefit from that.

More generally, in the public markets there are many flavors of separations between those who technically own stock, those that make the decisions of which stocks to buy and sell, and those who have contractual rights to reap the benefits or absorb the losses of stocks that

others buy and sell "on their behalf." This makes it challenging to formulate a clear cut definition of the term "investor" that feels right in all situations. Some people make a technical definition of an "investor" that would isolate the mutual fund in the example as the true investor: they reserve "investor" for someone who is making the final decision of what stocks to buy/sell, and doing the buying/selling with someone else's money. For our purposes, it makes more sense to categorize participants by goals rather than by mechanics, so the distinction of whose money is it is less important than what is the goal. A better definition for our purposes might be: an "investor" is an entity who makes a decision to buy or sell a stock in order to express an opinion on the medium or long-term health of an individual company, a set of companies (e.g. a sector of the economy), or the stock market holistically (e.g. through an index like the S&P 500). The opinion might be positive or negative: a short-seller, for example, expresses a negative opinion about the outlook for a particular stock by selling stock they have borrowed rather than own. If the stock price goes down, they will be able to buy the stock at a low price and return it to the lender, making a profit off the negative trend.

The shrewd reader will note the wishy-washy-ness of the phrase "medium or long-term" - let's say this covers anything from multiple days to months to years to decades. Which is not to say that investors seeking long-term returns won't occassionally buy and sell the same stock within a single day - but we'll make a distinction between investment strategies that are intended (or at least very willing) to hold risk over multiple days and longer time scales and those that are designed to get in and out of positions intra-day (or even in a matter of milliseconds). The short-term strategies we will discuss later below.

## Institutional Investors

Institutional investors are entities like mutual funds, pension funds, and hedge funds that take positions in the stock market with money that is pooled from many individuals. Those individuals trust the insitution to make decisions about which stocks to buy and sell, when to trade, etc., and they also trust the institution to allow them to put money in and out of the pool under agreed upon terms. If the insitution does well, the individuals share in its prosperity, and if it does poorly, they share in its paucity. The goals of the institution are typically stated publicly and known to the individuals contributing, and are typically a mix of growing capital, and controlling and/or hedging risk. To achieve this goal, the institutional investor may do

research and formulate opinions about the overall worth of a certain security, sector of securities, or the stock market as a whole, and choose a mix of financial instruments to express those opinions (in the sense that they will benefit if the opinions turn out to be true), and to hedge the risk of the opinions being wrong (e.g. use counterbalancing financial instruments to lower the probability of significant loss). Institutional investors are often colloquially referred to as the "buy-side" within the US equities market.

### Quant Funds

There is a growing segment of hedge funds whose strategies are driven primarily or entirely by quantitative models. Some of the most well-known quant funds are Renaissance Technologies, D.E. Shaw, and AQR. These types of firms are not to be confused with high frequency traders who also enact computer-driven strategies, but often make trading decisions on a much shorter (e.g. sub-millisecond) timescale, although it is worth noting that some firms employ both high frequency and less speed dependent strategies under the same roof.

## Retail Investors

Retail investors are individuals who make their own decisions about what stocks to buy and sell, and when to place orders. Essentially, retail investors directly reap the rewards or losses of the trading they direct. There is one technical caveat though, since retail investors send their orders to trading venues through brokers, who we will discuss in detail below. As we mentioned above, a broker might put its own company name on the stock certificates as it trades on behalf of a retail investor, and internally keep track of which customers "own" which stocks. This can create issues if a broker declares bankruptcy, and can't make good on everything on its books. There is insurance that mitigates this issue, provided by the SIPC, similar to how the FDIC ensures savings accounts at banks.

The goal of a retail investor may be similar to an insitutional investor, e.g. achieving a particular balance of risk and potential growth of capital. A retail investor may also be seeking to express an opinion about a particular security, sector, or the economy more generally.

# Trading Venues


Trading venues are the entities who match orders to buy and sell stock, and execute the resulting trades. A more common phrasing for that would be: "trading venues are where buyers and sellers find each other, and trades take place." We have deliberately eschewed this common phrasing, as it obscures the active role of the trading venue as an entity with its own goals. Despite the name of "venue," it can be misleading to think of a trading venue as a "place." A trading venue is still a "who."

A trading venue's goal might be to increase its revenue from trading fees or other fees that it charges to those who use its services. A related goal may be to attract as many buyers and sellers as it can in order to execute more trades. Another goal may be for the trading it executes to exhibit certain desirable features. To try to accomplish these goals, a trading venue has many tools at its disposal. It can design/modify its rules for how buyers and sellers express their interest to buy and sell semantically. It can design/modify its infrastructure for how buyers and sellers communicate with it mechanically. It can design/modify its rules for how compatible interests are ultimately matched into trades. It can also design/modify its rules for who can use its services, and what information about active interest and completed trades is provided to whom, up to some regulatory constraints. The constraints differ based on the type of venue, and from a regulartory perspective there are two types of venues:

## Exchanges

Exchanges are heavily regulated by the SEC and are also self-regulatory organizations (SROs) themselves - this means that the exchanges have certain regulatory authority over their customers. Exchanges generally publish firm quotes describing at what prices market participants are currently willing to buy or sell a given stock in real-time. In addition, exchanges publish messages describing the price and quantity of transactions after they're consummated. These trade and quote messages are known as market data.

In some markets around the world, certain securities trade exclusively on a single exchange, but in the US, pretty much all public companies are traded on all venues.

There are currently 3 major exchange families: ICE (NYSE) which operates 5 exchanges, Nasdaq which operates 3, and Cboe (BATS) which operates 4. IEX is the only other current exchange operator, though there are three other companies with exchange plans in the works: MEMX, MIAX, and LTSE.

# ATS's

Alternative Trading Systems (aka ATS's) are much more loosely regulated than exchanges and hence have lower barriers to entry. They typically don't disseminate quotes. ATS's come in two flavors:

## Dark Pools

Dark pools are designed to allow many would-be buyers to interact with many would-be sellers. They function pretty much like exchanges, but with a few major differences. Since they are scrutinized in a more lenient regularatory framework, dark pool operators can pick and choose which participants they allow to use their service. (Exchanges in contrast must provide fair access to any entity that follows the rules.) Dark pools also tend to tightly control information, not revealing who their subscribers are or any details about the trading in their pools that is not required for them to reveal. They do report trades after they happen - something we will discuss in more detail later after we have defined the mechanics of communication among the various players.

The typical explanation for why dark pools exist is that "investors with big orders to buy/sell want to trade quietly in a semi-private pool with other investors, away from the revealing

lights of exchanges which publish more information for all the sharks to see." But it is important to separate the goals of the investors and traders who might use a dark pool from the goals of the dark pool itself. The dark pool operator may have similar goals to an exchange operator: maximizing revenue from trading and other fees, maximizing trading activity, obtaining certain features of trading, etc.

Dark pools must be operated by broker-dealers, and they are typically run by large banks. The largest dark pools are run by UBS, CS, JPM, GS, Barclays, MS, and Level (owned by a consortium of banks). This introduces other possible goals a dark pool operator might have that don't necessarily make sense for the dark pool in insolation, but do make sense in the context of the operator's larger business.

## Internalizers

Internalizers are designed to allow many would-be buyers/sellers to interact with a single possible counterparty. These are crossing engines operated by proprietary trading firms and banks, where the operator trades against the client. Most retail trading occurs in this fashion, and a portion of institutional trading executes through these channels as well. The largest internalizers are Citadel and Virtu.

The operator of an internalizer might have it as a goal to execute more of their own trading internally so that they minimize their fees to exchanges and other venues.

# Brokers


A broker is a firm that facilitates trading on behalf of an investor. The brokers are collectively referred to as the "sell-side." In the US equities market, only broker-dealers are allowed to participate directly on an exchange. One goal of a broker is presumably to provide a service to the investor that meets the investor's needs, at a cost which is profitable to the broker.

The mechanics of the relationship between a broker and an investor can take on many different forms, but a good rule of thumb to keep in mind is that an investor decides what to trade and when to trade at a coarse level, while a broker decides when to trade at a more granular level, as well as where to trade and how to trade. For example, an investor may convey to a broker that they want to buy 10,000 shares of Microsoft sometime over the next few hours. The broker will decide if/how to break this amount into smaller orders and how to put these orders into trading venues, potentially spreading them out over venues and over time. The investor will naturally want the broker's decisions to result in purchasing the target amount of stock within the target time frame at the lowest possible cummulative price. Competition among brokers is supposed to drive brokers to achieve this goal while charging competitive fees for their services.

There can be complicating tensions, however, between the goals of the investor and the goals of the broker that arise from the ways that brokers charge for their services, which may holistically encompass more than just trading. To achieve goals like attracting business and maintaining/increasing profitability, brokers can do more than just tinker with their trading algorithms. They can also design and change their billing practices, as well as their full offering of services.

## Commissions

Unlike many other businesses, brokers often charge different commission rates to their customers, even for the same type of execution. In US equities, broker commissions are usually charged on a flat per-share basis, or on a simple scale based on the securities price (for example, 1 cent per share on stocks below $5; 3 cents per share on stocks above $5). Common equities commissions range from free to 5 cents per share (an enormous range), with a typical rate in the ballpark of 0.5 cents per share.

## Cost Plus

Many brokers offer their clients the option of cost plus pricing, which means charging a relatively small commission and passing through their actual explicit transaction costs. For example, if a broker charged "cost plus 10 mils," this would mean they tally up the exchange and dark pool fees (and rebates) that they incurred while executing this order, and then add a flat 10 mils per share on top of this. (A "mil" is a one-hundredth of a penny, for some cruel reason, despite the usual rule that "milli" as a measurement prefix corresponds to a one-thousandth, e.g. there are 1000 milliseconds in a second. But I digress.) Cost plus pricing has the wonderful benefit of removing the incentive for a broker to route (potentially to the client's detriment) based on its own economic outcome, but it complicates the billing process and many investment firms are unable to process cost plus billing. Without cost plus pricing, brokers are incentivized to route to venues that charge lower fees or give them rebates, which may be in conflict with getting the best trade price for clients.

## Researchers and Corporate Access

Full service brokers also provide services like research and corporate access. Researchers produce detailed research reports on companies and sectors. This role exists at companies called "full service brokers," which provide a wide array of services to institutional investors beyond trade execution. Full service brokers also coordinate direct access to corporate executives for interview.

In the US, brokers generally don't charge fees for these additional services, and instead they expect to receive indirect compensation via their trading commissions. Those payments are

thus considered part of the transaction costs as opposed to operational costs for their investment firm clients, an arrangement which many buy-side firms prefer. This means the typical goal of providing these additional services is to acquire or at least maintain trading business. This leads to some trading behaviors that seem weird in isolation but can be explained in the larger context. For example, an institutional firm may receive comparatively worse trading services for higher commissions, and still be unwilling to move too much of their business to another broker, because they value the research and other services provided by their full service broker.

## Prime brokerage

There are other services that can be bundled together in the offerings a broker makes to clients like hedge funds. For example, a hedge fund who uses multiple brokers may not want to track and manage all of their trading relationships and collateral requirements individually. A prime broker can serve as a centralized manager of their activity to reduce the overhead for the client.

## Stock loan

Brokers can also facilitate stock lending and borrowing for their clients. This can allow clients to sell a stock "short," meaning that they don't own the stock they are selling, but rather have borrowed it, with an obligation to return it at a later point in time. Short sellers are often expecting the price of a stock to go down, and can profit if it does, since they can buy the stock to return at the later time in the market for a lower price than what they previously sold. Brokers may manage the lending and borrowing of stocks for their clients, and may set collateral requirements for securing the loans, etc.

*Further reading*

https://en.wikipedia.org/wiki/Prime_brokerage

https://en.wikipedia.org/wiki/Securities_lending

# Proprietary Traders

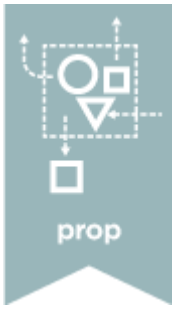


A large segment of the trading ecosystem is proprietary trading firms - companies that use their own capital to buy and sell securities directly in the market. By most estimates, these firms comprise roughly 50% of all trading volume in the US. Some of the largest firms in this segment are Citadel, Virtu, HRT, Tower Research, IMC, TwoSigma, and Quantlab.

In terms of goals, it's not clear why someone trading with their own money should behave any differently than someone trading on behalf of others. Growing capital and managing risk etc., are reasonable goals for either case. However, there is a general differentiation in terms of time horizons. At least most (if not all) of the large proprietary trading firms focus on time horizons within a day or even shorter time intervals, and do not seek to hold positions overnight.

At short time horizons, one common goal is market making, which means connecting buyers and sellers who are separated by time and/or venues, and charging a fee for the service. Market makers may stand ready to buy or sell a stock at any moment, with the price they are willing to buy at being slightly lower than the price they are willing to sell at. The difference between these prices is called the "spread," and it is the compensation that market makers get for providing immediate transactions to buyers and sellers who might otherwise have to wait seconds, minutes, or hours to find each other in the larger market, if at all.

A proprietary trader might also have a goal of expressing a very short-term opinion about the price of a security. The time horizon here might be as short as 1 millisecond or less. Such opinions can be formed on the basis of very fast data feeds and recognition of common

patterns in how price changes tend to travel throughout the fragmented landscape of venues. Some high frequency trading strategies seek to leverage predictive analytics acting on top of expensive high-speed data access at various venues and high-performance architecture to make a profit off of patterns of trading at a millisecond or microsecond time scale. Market makers may also employ such technologies to try to adjust their buying and selling prices ahead of detectable trends at this time scale.

# Technology Vendors


Technology vendors provide services that are useful (sometimes arguably required) for other participants. Technology vendors may provide physical and/or logical connectivity to trading venues, software that is used to communicate and/or track orders and trades, processed data feeds that contain various kinds of quote and trade information, access to microwave towers or fiber optic cables that carry messages quickly between market participants, cloud computing resources, database and analytics tools, etc.

The goal of a technology vendor is typically to maximize its revenue in the long-term, but this may be consistent with low pricing in the short-term to encourage broad adoption and entrench dependency on their services. Especially for technology that connects participants, like execution management systems (EMS) that connect investors to brokers, it can be hard for participants to switch providers, as this has to be coordinated by both sides of the connection. These kind of effects dampen the power of competition among technology providers to drive prices lower and service quality higher.

A technology provider might be in a strongly unique position to offer a particular service if it sits inside a company such as an exchange. For example, NYSE is clearly in the best position to offer optimally fast connection services into a NYSE exchange. Similarly, it is in the best position to offer up-to-date data feeds about what is going on at a NYSE exchange. These natural monopolies of latency-sensitive products and services that orginate at a trading venue are another phenomenon that heavily limits the power of competition to control pricing of services in the securities trading ecosystem.

# Clearing / Back Office




After a buyer and a seller are matched on a trading venue, there are back office processes that need to be done. The buyer needs to pay the seller, and the seller needs to produce the security. The security's title needs to be formally changed so that it now demonstrably belongs to the buyer. The entity who performs these back office processes is called a clearing broker. The clearing broker charges a fee for the services it provides, and manages regulatory compliance with this part of the trade. The goal of a entity performing a clearing service may be to maximize revenue from this service, or it might be to attract customers that it can then pitch on other services offered by the same firm.

# Regulators


Regulators are charged with making sure that the stock market functions in an orderly fashion, and serves its intended purposes. Regulators create and update rules intended to protect the health of the market, and are charged with enforcing these rules. Regulators also must scrutinize and approve or disapprove changes that exchanges propose to their existing policies. The goal of regulators is to ensure that the rules are designed to serve the interests of investors and the market as a whole, and that the day-to-day activities in the market follow the rules.

## SEC

The primary regulator of the US equities market is the SEC, which is a government agency. The SEC's mission is two-fold: investor protection and the facilitation of capital formation. The SEC is tasked with, among other things, policing investment firms, brokers, exchanges, and corporate issuers. The SEC is empowered to make policy that affects the operations of market participants. Sometimes the SEC will test the impacts of new policies before making them widespread or permanent in the form of "pilots." Pilots are policy changes that are applied in a way that is limited in time or in scope, typically designed to yield a meangingful comparison with market behavior happening in parallel that is not subject to the contained changes.

The SEC is also empowered to review and to potentially reject changes to operations that trading venues like exchanges would like to make. The goal here is to ensure that proposed

changes are not counter to the interests of investors or the facilitation of capital formation. This can be a difficult job, as there is often little or no quantitative basis to understand what impact the proposed changes might have until they are implemented.

## FINRA

FINRA (the Financial Industry Regulatory Authority) is a private not-for-profit organization that oversees the broker-dealer industry. Like the exchanges, it is a self regulatory organization and thus has power to create and enforce rules and impose fines on its members.

As per their website, FINRA does five things:

1. Deter misconduct by enforcing the rules
2. Discipline those who break the rules
3. Detect and prevent wrongdoing in the US markets
4. Educate and inform investors
5. Resolve securities disputes

# Mechanisms of Communication between Participants

## The When and the How

Now that we have laid out the different types of participants in the US equities market and what their goals might be, we'll move on to discussing the logistics of communications between them. Any full description of communication and interaction in a complex ecosystem should start with the core, established rules that participants all operate under and assume. These do not need to be actively communicated, but rather form the shared context underlying communications. In US equities trading, there are core rules concerning time and price. After discussing the basics of these, we'll go on to the discuss the logistics and content of active communications between the various participants.

# The Trading Day

Trading in the US equities market is confined to certain hours of the day. The primary concept of the "trading day" refers to the hours from 9:30 am to 4:00 pm, eastern standard time, Monday through Friday, except for certain declared market holidays. At the start of each trading day, there is a single auction in each security that sets the official opening price for the day. The opening auction is facilitated by a single exchange, which is called the "listing exchange" for that security. The listing exchange collects orders from would-be buyers and sellers, and matches many of them at a single opening price. This typcially happens precisely at 9:30 am. From then on throughout the rest of the day until the end, the security can trade at any venue (not just the listing venue). Similarly, leading up to the 4:00 pm the listing exchange will collect orders from would-be buyers and sellers for a closing auction, which will set the official closing price for that security for the day. The listing exchange for a security has a few other duties as well, such as halting trading in that security when certain conditions arise, and monitoring the security for compliance with its listing standards.

Many trading venues offer "pre-market" and "post-market" hours, allowing participants to trade securities ahead of the opening auction or after the closing auction. These are still typically limited in scope (e.g. a pre-market session might open at 8:30 am eastern). One main difference between these supplemental trading sessions and the official "trading day" is that some regulatory constraints do not apply to pre- and post-market trading.

# Pricing Increments

Trading in the US equities market is also confined to a discrete standard of pricing. Would-be buyers and sellers of securities on exchanges must advertise prices over $1 in whole pennies, and trades must be executed at whole penny increments, with the exception of trades at the "midpoint" of advertised buy and sell prices, which sometimes lands on a half-penny increment. The rules are a bit looser for dark pools and internalizers, or for trades that are consummated by participants themselves outside of venues, but there are still constraints of this spirit.

The main goal of imposing a coarse discretization of price is to avoid participants snatching trades out of each other's hands by offering essentially meaningless improvement on price. For example, a would-be buyer who is willing to pay $10 a share for a stock being out-bid by someone willing to pay $10.0000001 a share for a stock is not very civilized. This sort of thing would undermine confidence in the health of the market system, and incentivize some rather silly behavior. Though this motivation of discretization is relatively clean and clear, there is a strange quirk of the way the rules for exchanges are currently written and enforced: participants often *trade* at half-penny increments on exchanges as mid-point trades, but can only *express interst to trade* at full-penny increments.

# Communication between Investors and Brokers

Once an investor has decided what they'd like to buy or sell, how much, and some goals in terms of price and/or timing, they need a way to communicate this to a broker who will manage the process of formulating individual orders, routing those orders to various trading venues, and keeping track of the trades that are produced. There is a lot of variety in exactly how granular the investor's instructions are and how they are communicated. Often the communication is accomplished or assisted through an OMS and/or EMS:

## EMS/OMS Providers

The current standard buy-side trading workflow incorporates certain core pieces of technology, including an order management system (OMS) and an execution management system (EMS).

An OMS tracks the life cycle and status of orders across the firm, often across multiple asset classes. The "life cycle" of an order refers to all of the relevant events that can occur over the "lifetime" of an order: events like the order being placed at a venue, partially filled, modified, routed away to another venue, fully filled, or canceled, etc.

An EMS is a tool that focuses on individual order execution and allows more fine-grained control, such as slicing a "parent" order into "child" pieces to be sent to a specific counterparty or venue. The EMS also generally has more context on what's happening in the market in real-time to allow the trader to make informed decisions.

A typical work flow is: a portfolio manager at an institutional investor decides to put on a position in US equities, creates an order in the OMS and assigns it to a trader. The trader then stages the order in their EMS which allows them to send pieces (child orders) externally to their brokers. As these child orders get executed at various trading venues, those fills (aka trades) flow back to the EMS, and subsequently to the OMS. The OMS will likely perform or

enable various compliance and risk functions, although in some cases, the EMS might have some built-in risk check logic as well.

These two sets of functionalities are quite related and overlapping, and some vendors try to combine them into a single offering. Overall, the general quality of these technology products is unimpressive, and like many other services in this industry, their pricing models and specific offerings are extremely opaque.

It's an extremely sticky business as the process to transition from one provider to another without losing any custom functionality is a painful one.

*Further Reading*: New Plateaus for OMS EMS Integration, EMS Consolidation, Liquidnet Trading Solutions

# High touch vs. low touch

Within a broker, there are different models for how investor orders are handled. One main axis of differentiation is "high touch" versus "low touch."

In this context, high touch refers to execution services managed by a human, whereas low touch refers to automated execution services, such as trading algorithms. High touch equities teams are sometimes called "cash trading" and low touch "electronic trading." High touch desks often provide additional services such as shopping orders to other clients to directly match a large amount (especially large trades are often called "blocks"), committing capital to take down some or all of the trade, and providing market color and commentary.

## Roles on a high touch desk

A high touch desk will typically have employees organized into the following roles:

1. **Trader** - a person who directly manages the execution of orders.
2. **Sales-trader** - a hybrid trader role and sales role that performs duties like shopping flow.
3. **Sales** - account management and business development.
4. **Desk analyst** - a person who conducts independent research and provides more in-depth color than a salesperson or sales trader.

### Roles on a low touch desk

A low touch desk will typically have employees organized into the following roles:

1. **Sales-trader/sales** - same as the analogous roles above.
2. **Algo developer/product management/client implementation** - there are several front-office technology roles that work together to build the actual trading platform, manage customizations, and connect to customers and the street.
3. **Operations/support** - the low touch desk's technology is often the backbone of the high touch desk, and when software/hardware/networking issues come up, people in these roles lead the effort to resolve them.

## Algo Customizations

One of the core features of most institutional trading platforms is the ability for individual clients to customize the behavior of their chosen trading algorithms. For example, a buy-side trader who is sending orders to be executed by a sell-side broker's electronic trading algorithm might have the option to specify that "the algo" (as it is often lovingly called) should exclude a subset of dark pools, and/or maintain a minimum percentage of volume, and/or always utilize a given minimum execution quantity. These customizations are generally configured or hard coded behind the scenes in the broker's system.

## Transaction Cost Analysis (TCA)

Buy-side investors may reasonably want to track and compare performance of brokers across many orders. They might use such data to inform decisions about what kinds of orders to give to which brokers, etc. Getting the relevant data in an actionable form can be challenging, as brokers are sometimes resistant to providing granular detail about their operations to the buy-side clients, and it is not an easy task to isolate the effects of broker decisions from the cummulative biases of data samples and broader market activity. Brokers providing their own analysis of their own trading is clearly problematic, as they are naturally incentivized to choose metrics and data sets that make them look good.

This particular category of communication from brokers back to investors is not very standardized, and the results are often not very satisfying. Getting enough relevant data into

a usable format alongside sufficiently granular market data in order to draw meaningful conclusions remains a non-trivial task. This is especially true for relatively smaller investors who have less resources and leverage to devote to this. As a result, several third-party providers have arisen to provide transaction cost analysis and help buy-side traders evaluate sell-side broker performance. For more discussion of transaction cost analysis and what it entails, see our related blog post.

# Communication between Brokers/Proprietary Traders and Trading Venues

There is a lot of information that flows back and forth between the trading venues and their members (the brokers and proprietary traders who are trading on their platform). There is the flow of orders from members into the venue, but there is also a lot of information that must travel the opposite way. The members want to receive timely information from the venue about the status of their own orders and trades. Members also want timely information on all quotes and trades happening on the venue, so they can base their own decisions on an up-to-date view of the market.

At a technical level, information passing from one party to another involves several different layers of physical and logical infrastructure. There is the physical equipment that carries data packets between machines, a protocol specification for the format and meaning of each individual packet, and a high level specification of what data will be communicated, when it will be communicated, and how it is organized. Below we'll give a brief overview of these components for the communications between trading venues and their members. It is also worth noting that these components in aggregate account for the latency of information between brokers and trading venues and between proprietary traders and trading venues. Though these brokers and proprietary traders often use similar means of communicating with venues, their actual needs in terms of latency might vary considerably. We discuss this further in our blog.

## Microwave/laser networks

The fastest way to send packets of information across miles of geographic distance is via microwave towers. As a result, market participants seeking the lowest latency in their connections to and from exchanges use microwave as their primary communication medium.

However, microwave can be unreliable (particularly in bad weather), so these participants typically also have fiber optic connections to fall back on.

## Extranet providers/dark fiber

Participants who are slightly less latency senstive might use fiber optic connections as their primary way of connecting to trading venues. There are many "extranet" providers who are in the business of leasing their extablished fiber optic connections to market participants.

## Data centers

Whether they are traveling via microwave or fiber, data packets orginate and terminate at machines that are housed in data centers. The time it takes a data packet to be formulated at its origin, travel from its origin to its destination, and be processed at its destination is referred to as "latency," and it depends upon the medium of transport (microwave vs. fiber), the distance traversed, any intermediary routing technology (switches etc.), congestion, and the hardware and software that handles the low-level details of packet processing.

The spaces in data centers nearest to where the machines belonging to a NYSE, Nasdaq, or CBOE family stock exchange reside are controlled by those same companies, and they allow market participants to purchase "colocation" - the right to place their machines very close to the exchange's machines. This gives participants an (expensive) option to minimize the distance that data packets must travel between their machines and the exchange's machines, thereby shortening the latency.

## FIX

FIX stands for the Financial Information Exchange protocol. It was invented in 1992 as a more reliable and unambiguous replacement for communication between investors and brokers that previously had taken place over the phone. FIX defines a format for each data packet that allows the recipient to interpret its contents, and is tailored to the kind of information that is relevant to financial transactions. At this point, FIX is a common organization of communication between investors and brokers as well as between stock exchanges and their members. Despite its ubiquity and sufficient expressivity, however, FIX

is not a universal choice among financial participants today because it is not particularly optimized for minimal latency.

# Proprietary protocols

Most stock exchanges offer their members multiple choices of protocols that control packet formatting. FIX is one choice, but typically the choice that is closest to the native implementation of the exchange software will produce the shortest processing times at the exchange, and hence the lowest latencies. Each exchange family offers the choice of access in a proprietary format that is designed to be highly compatiable with its internal implementation, and it is not surprising that use of these proprietary protocols tends to minimize the packet processing times, and hence the overall latency. As a side effect, this means exchange members have to customize their communication software to the various exchange families, and typically build translation modules to go between the various proprietary formats and FIX or any other formats they use internally.

# Orders and Market Data Content

Now that we've touched upon the infrastructure that physically transports packets, and the protocols that are used to format and process them, we'll zoom out a level and talk about what pieces of information go into data packets and how these pieces of information are commonly structured.

## Order Types

Brokers communicate their desires to buy/sell stocks to a venue through a menu of order types and associated parameters that the venue defines. A would-be buyer or seller of stock is typically torn between two goals: getting a favorable price, and getting the trade done quickly. Different order types allow participants to express different trade-offs and constraints relating to these competing goals.

There are a lot of important differences in the menus provided by different venues, but most are variations on a few common themes:

**Market Orders**

A market order is used to communicate an immediate desire to buy or sell, irrespective of price. When a market order to buy enters a venue, it can be immediately matched with any open orders to sell. Conversely, when a market order to sell enters a venue, it can be immediately matched with any open orders to buy. This order type represents an extremal point in the trade-off between price and time: full priority is given to executing immediately, completely insenstive to price.

### Limit Orders

Limit orders are used to express a constraint on price. The presumed goal is to execute a trade as soon as possible, within the band of acceptable prices. A limit order to buy, for example, will specify a ceiling for price, and will trade at the first opportunity to buy at or below its limit price. A limit order to sell will specify a floor for price, and will trade at the first opportunity to sell at or above its limit price.

The combination of limit orders and market orders is enough to build up some illustrative examples of trading dynamics. A market maker might have active limit orders to buy 200 shares of a stock for $10.00 a share and to sell 200 shares of the same stock for $10.02 a share. While waiting and available to be filled, these limit orders are referred to as "resting" on the order book kept by the trading venue. A buyer who wants to immediately buy 200 shares might enter a market order to buy, which will be matched against the limit order to sell at a price of $10.02.

### Midpoint Orders

A midpoint order is a way of delegating the determination of price to the broader market. Instead of declaring a specific limit for price, a midpoint order will execute as soon as possible at a price that is equal to or more favorable than the midpoint of the current spread. In other words, it behaves like a limit order, but where the limit is a dynamically adjusting price calculated by taking the midpoint between the highest open buy limit orders and the lowest open sell limit orders. This calculation is typically over the price limits being advertised across all exchanges (which does involve some latency as the relevant information travels from the origin exchange to the venue processing the midpoint order).

If we return to our toy example of a single market maker offering to buy 200 shares for $10.00 a share or to sell 200 shares for $10.02 dollars a share, and we assume there are no better prices available across all exchanges, the midpoint price is therefore $10.01. If a midpoint order to buy enters a venue under these circumstances, it will be willing to buy at $10.01 or lower, and therefore will not be matched against the limit order to sell at $10.02. It will wait to interact with a seller willing to sell at $10.01 or lower (which might be a market order to sell, a midpoint order to sell, etc.)

### Order Attributes

Many order types allow those submitting orders to specify additional parameters that control more of the fine-grained mechanics of how the orders behave. Limit orders can be "lit" (aka "displayed"), meaning that their sizes and limit prices are visible to other market participants who might want to take the other side of the trade, or "dark," which means they are not visible. Market makers typically use lit orders to advertise and generate trades.

Dark orders can often specify a minimum quantity, meaning that they will not execute for less than that specified number of shares. For example, an order for 1000 shares that has a minimum quantity of 500 shares can not be matched in a trade for 100 shares. It will wait until at least 500 shares can be traded at once.

Some orders can also have pushy designations like "immediate or cancel" (IOC), meaning that they must be executed immediately or not at all, or "fill or kill" (FOK), meaning that they must be executed in full or not all. There is a virtual alphabet soup of overall options, so we will not try to give a comprehensive accounting here. We will discuss more order types later to the extent that they are relevant to higher level issues of market interactions.

Now that we've given a brief overview of how brokers communicate their desires to venues, we'll move to discussing how venues communicate information about trades and prices to their members or other financial participants:

## Tapes

Publicly traded stocks in the US are organized into 3 groups called Tapes. Tape A has NYSE-listed securities, Tape C has Nasdaq-listed securities, and Tape B has everything else. As discussed above, a security is "NYSE-listed" if NYSE is the venue responsible for

running its official opening and closing auctions (among other duties), but it can trade on all venues at all times.

For a particular stock, these are the types of information that participants are most interested in:

## Top of book and depth of book data

"Top of book" and "depth of book" both refer to data about the best prices and associated quantities that would-be buyers and would-be sellers are currently advertising on a particular venue. The prices and quantities that buyers are advertising are called "bids," and the prices and quantities that sellers are advertising are called "asks" or "offers." For example, a bid might be "would buy 100 shares of security X for $40 a share." This is of course a wordier version than the shorthand formats that participants use, but this is what they mean. Similarly, a sample offer means: "would sell 200 shares of security X for $40.02 a share."

At a particular trading venue, if a buyer was advertising a price that was equal to or higher than a price being advertised by a seller, the buyer and the seller could just trade with each other (at least for the minimum of the two share quantities). Since the venue will catch this and match the trade, the highest bid price at a given venue should always be lower than the lowest offer price. The difference between these is called the "spread" - we've mentioned this before as a driver for compensation to market makers who are constantly willing to buy or sell. The collection of bids at the current highest bid price and the collection of offers at the current lowest offer price are referred to as the "top of book." Data feeds that are categorized as "top of book" include only information about these advertised orders, and not any bids at lower prices or offers at higher prices. Naturally, the current highest bid price and current lowest offer price are moving targets as trades happen and advertised bids and offers are canceled or adjusted. Top of book data feeds are typically event-driven, meaning that as soon as the top of book information changes, the new information is immediately disseminated. A top of book data feed might aggregate the sizes of bids or offers at a particular price, or it might give granular information about each bid/offer.

"Depth of book" refers to information about bids and offers that goes beyond the highest priced bids and the lowest priced offers. A typical depth of book data feed will include information about all bids and offers, but there are some variations on this. Some depth of

book data feeds aggregate sizes of the bids or offers at a particular price, some given granular information on each bid/offer, and some place a limit on the number of distinct prices (sometimes called "levels") that are included. For instance, the CBOE exchange family offers some data feeds that include information only about the 5 highest price levels for bids and the 5 lowest price levels for offers.

As an interesting side note, the highest bid price and the lowest offer price *across venues* can be equal, a state which is referred to as a "locked market." The highest bid price can even be strictly higher than the lowest offer price when we look across venus, which is referred to as a "crossed market." These situations do arise fairly frequently, but are typically quickly resolved, as someone recognizes an opportunity. You might expect that if a seller on one venue and a buyer on another have compatible desires, then both would likely consume the relevant data, recognize the compatability, and one of them would send an order to the other venue to trade. But even if both participants are tracking the bids and offers on the other venue and see each other, billing practices at venues can complicate this story. Many of the exchanges are "maker-taker" venues, meaning that they charge a fee to the entity who submits an order that executes against an advertised bid/offer (aka the "taker") and they give a rebate to the entity whose advertisement led to the trade (aka the "maker"). If a buyer is advertising a price P on a maker-taker venue and a seller is advertising that same price P on a different maker-taker venue, both entities might be banking on receiving the rebate instead of the fee to view that price P as worthwhile. This sometimes creates an amusing standoff until someone changes their view of the price they are willing to pay, or another participant intercedes. This scenario couldn't explain a crossed market though, since the range of rebates/fees is capped to be smaller than the 1 penny price increment allowed for bids and offers.

## Trades

In addition to quotations of what would-be buyers and sellers are advertising, market participants also want data about trades as they occur. Once a trade is executed at a venue, the venue sends confirmations to the involved parties, as well as disseminating some information to the larger market ecosystem about the trade. The broader population can learn what security was traded, how many shares, at what price, and at what time. The identities of the buyer and seller are not revealed. In fact, even the trading parties themselves will

typically not know each other's identities. If a trade happens on an exchange, the identity of the exchange will be publicly visible as well. If it happens on a dark pool, the trade will be identifiable as having occurred off exchange, but the identity of the specific pool will not be publicly visible.

### Auction Data and Alerts

Leading up to an auction, an exchange may choose to make some data available concerning the possible price or the imbalance between orders to buy and orders to sell. This kind of data may be used by participants to adjust their orders leading up to the auction, and may motivate new buyers and/or sellers to participate.

There are also alerts dissimentated through data feeds that participants might need or want to consume. For example, alerts may be sent out when trading is halted in a stock.

## Market Data Dissemination and Trade Reporting

Now that we've talked about what kind of market data participants typically seek and receive, we'll discuss the data feeds that disseminate these types of data and the entities that control them.

### SIPs

The SEC has mandated that some market data be made publicly available via Securities Information Processors (SIPs). Currently there are two SIPs, which are operated by NYSE and Nasdaq, respectively. The NYSE SIP (CTA) publishes trades and quotes for Tape A and Tape B securities, and the Nasdaq SIP (UTP) publishes trades and quotes for Tape C securities. In terms of quotes, the SIPs provide only top-of-book information from each exchange. A change in the top-of-book at a given exchange is first communicated by the exchange to the SIP and then from there disseminated to recipients of the SIP data feeds. This two hop approach is unsuprisingly slower than sending the new data directly from each exchange to the ultimate recipient.

Even though these are "public" data feeds that lack depth-of-book or auction imbalance information, they are still extremely expensive to connect to, consume, and redistribute,

especially for real-time data.

## Proprietary data feeds

Exchanges also offer multiple kinds and tiers of proprietary data feeds. These products run the gamut from trades and top-of-book quotations to full depth-of-book at the particular exchange. Because they can be consumed directly from the exchange without the data first traveling to a central processor as is the case for the SIPs, these proprietary feeds can deliver data at the lowest latencies (especially to recipients who are co-locating their equipment in data center space provided by the exchange). Some of the data products are free (e.g. all of the data feeds from IEX or NYSE National), some of them are a bit cheaper or comparable in pricing to the SIPs (e.g. Nasdaq Basic), but the full depth-of-book feeds at the main exchanges with the highest market share are an order of magnitude more expensive.

For more details on market data pricing for the real-time data products offered by exchanges, see our market data series of blog posts (part 1, part 2, part 3) and our market data pricing visualization tool.

## TRF

The FINRA Trade Reporting Facilities are where information about trades that happen off of exchanges is reported and then relayed to the appropriate SIP for public dissemination. There are currently two TRFs, also operated by NYSE and Nasdaq, but unlike the SIPs the TRFs are not partitioned by Tape. Instead, each dark pool may strike an agreement with the TRF of its choosing to report its trades. Such trades are ultimately visible to others via the SIPs, though with potentially greater latency than exchange trades and with the precise venue not specified. In other words, the TRFs allow other market participants to learn the price, size, and time of off-exchange trades, but only the binary attribute of "off-exchange" is visible/inferrable, not further granularity on which off-exchange venue produced the trade.

# Communications among Brokers/Proprietary Traders

In the realm of retail trading, some communication happens directly between brokers and proprietary traders. In fact, most retail trading occurs through direct relationships between retail brokers and proprietary trading firms. The process typically works like this: a retail trader submits an order to a retail broker (e.g. Robinhood, ETrade, etc.) The retail broker sends it to a proprietary trading firm that is acting as a retail wholesaler. The proprietary trading firm decides whether they want to execute the trade with themselves as the counterparty or route it elsewhere.

This is an attractive arrangement for the retail broker, who typically gets payment for order flow (PFOF) and/or price improvement from the retail wholesalers. In this context, "price improvement" means that the end customer (the retail trader) will get a slightly more favorable trade price than what market makers are currently advertising on the exchanges. The retail wholesaler is willing to do all of this (PFOF and/or price improvement) because retail traders typically have small orders (hence limited risk) and trade without significant short-term alpha. (This is a polite way of saying that on a short-term time scale, retail traders don't tend to buy low and sell high.) For this reason, market makers can make an excellent profit off of retail traders in aggregate, so they are willing to grant better prices and pay for order flow when they know they are trading against a less sophisticated counterparty.

# Regulatory Constraints and Surveillance

Here we will discuss the main regulations that serve as constraints on market activity, and the communications and data that regulators rely upon to enable their enforcement. These regulatory constraints should be thought of as part of the assumed common knowledge underlying market interactions. We did not discuss all of them above when we discussed the logistics of the trading day and discrete price increments only because some of them require a pretty deep in the weeds perspective to describe. Only now after we have discussed more of the details of communications and interactions between various participants do we have the semantic tools to summarize some of these regulations and their current effects on the market ecosystem.

## Major Regulatory Constraints

In the wake of the stock market crash of 1929, new federal laws were passed with the intent of protecting investors and the integrity of markets. Prior to this, stock sales were regulated only under state laws. The goal of the Securities Act of 1933 was to set consistent standards for public companies to protect investors from fraud. The Securities Exchange Act of 1934 created the SEC (the Securities and Exchange Commission) as the regulator of the public markets and established the concept of self-regulation, making exchanges share responsibility for policing their own markets.

Since then, many more regulations have been put in place with the intent of constraining potentially harmful trading activities, incentivizing healthy trading, and ensuring reliability and availability of market infrastructure. We'll address much of this in more detail in a market history section later, so for now we will just highlight the regulations that have a major impact on current trading mechanics.

### Securities Act Amendments

In 1975, the securities act amendments were passed, creating a more organized and centralized system to connect exchanges together. The SIPs were created to be a central source of the most relevant information across disparate exchanges.

## Reg NMS

In 2005, Regulation National Market System (aka Reg NMS) was introduced which established the Order Protection Rule, among other regulatory changes. The order protection rule mandates that trading venues cannot match trades at prices that are inferior to the best prices currently available on other venues. For example, if a market order to buy arrives at a venue V where the lowest limit a seller has set is $11.00 per share, but some other venue W has a displayed sell order available at $10.50 per share, the buyer and seller cannot be matched at venue V for $11.00 per share. The buy order must be routed to venue W instead to be matched with the seller there at the lower advertised price.

Obviously this is intended to protect the buy order from being executed at a sub-optimal price, but there are a lot of devils hidden in the details. For instance, what does "currently available on other venues" technically mean? For venue V to fulfill its obligations under Reg NMS, it must now track the best available prices on venue W (and everywhere else). Since information takes time to travel from venue W to venue V, the view that venue V has of the "best prices currently available on venue W" is necessarily delayed. The amount of this delay is a dynamic function of all of the communication infrastructure between W and V. V is this regard is similar to any broker connecting to venue W, and it must choose how much to invest in reducing latency.

## Reg SHO

Regulation SHO concerns the practice of short-selling, which is selling a stock without owning it. Reg SHO concerns "locate" and "close out" requirements which try to assure that stocks which are sold by short sellers can be reliably borrowed and delivered on time. Reg SHO also puts limits on short selling in the midst of a substantial intraday price decrease. This is intended to address concerns that short-selling can be used manipulatively to artifically depress the price of a stock.

## Reg SCI

Regulation Systems Compliance and Integrity (or Reg SCI as it is commonly called) was established in 2014 with the goal of reducing the likelihood of significant to catastrophic technology failures in the financial markets. This regulation created more formal requirements for the operations and testing of core components of trading technology and applies broadly to venues and clearing agents.

### LULD bands and circuit breakers

Limit Up Limit Down bands (LULD) prescribe a certain price window that a stock can trade in, as a function of the stock's prices over a preceding five minute window. The bands ensure that a price can't increase or decrease too suddenly by too large of a percentage. Trades outside the band are prohibited, and if prices do not revert back into the band within 15 seconds, a 5-minute trading pause in the stock is triggered.

While LULD bands function individually for each stock, there are also market-wide circuit breakers, which are triggered when the S&P Index decreases by certain percentages, as compared to the previous day's closing price. There are three levels at which circuit breakers can be triggered: Level 1 is triggered by a 7% relative decrease, and will halt trading across the market for 15 minutes (though only if it is before 3:25 pm). Level 2 is triggered by a 13% relative decrease and will also halt trading broadly for 15 minutes (again only before 3:25 pm). Level 3 is triggered by a 20% relative decrease, and halts trading across the market for the remainder of the day, no matter the time.

## Logistics of Regulatory Surveillance

To monitor for compliance with regulations, to investigate potential cases of market manipulation, and to more generally keep a finger on the pulse of the markets, regulators needs access to granular data about trading mechanics and behaviors. Here we discuss some of the primary data sources that FINRA and the SEC collect from market participants to enable regulatory surveillance.

### 605/606 reports

SEC Rules 605 and 606 require venues and brokers to make periodic public disclosures of basic aggregate information about executions. Rule 605 mandates that venues make

monthly reports of basic execution quality metrics on publicly accessible websites. Rule 606 mandates that brokers disclose information about which venues they route orders to and in what proportions.

## OATS

OATS stands for the "Order Audit Trail System" that is used by FINRA. It was established by FINRA rules in 1998 (approved by the SEC). All FINRA member firms are required to collect and submit the data according to the technical specficiations of OATS, and FINRA uses the data for market surveillance. The data spans orders, quotes, and trades for listed stocks as well as over-the-counter securities (securities that are not listed on a stock exchange). In aggregate, this data is intended to allow the re-creation of order lifecycles, so that the chain of events from order entry to order modification or cancellation or execution can be followed.

## CAT

In 2012, the SEC called for the creation of a "Consolidated Audit Trail" (CAT) to serve as a key data source for market surveillance and regulatory investigation. This is intended to replace OATS and to enable superior surveillance capabilities. In the longterm, all venues and brokers will report data into the CAT processor in prescribed format. The CAT processor will be responsible for assembling this data into easily query-able form and linking together the different pieces of an order life-cycle, so events like order submission, modification, routing, execution, and cancellation can be stitched together and in some cases attributed to the originating investor. Granular identifiers of venues, brokers, and end customers will be maintained so that ultimately regulators can identify actors in involved in any questionable activity. Regulators such as the SEC and FINRA will be given access to the CAT in prescribed ways to enable market surveillance.

The CAT is intended to represent a more complete and sophisticated tool for regulators (particularly at FINRA and the SEC) to get effective information about behaviors in the marketplace. However, it is a complex project with many different stakeholders who are often at odds about exactly what data should be collected, how quickly reporters should have to adapt their systems to report such data, and who should have access to what. Long delays and disagreements have plagued the CAT since its inception.

# Birds-eye View of Overall Market Activity

## How Much and How Often

Now that we've fleshed out the different kinds of entities in the US equities market and the logistics of communications between them, it's time to zoom out a bit and get a sense of scale and proportions. How much trading activity happens on a typical day? How it is distributed by size and time and place? How much variation is there intraday and interday? Here we will give you some basic answers to these questions. Obviously, there is a lot more structure and patterns underlying trading than the first order statistics here can capture, but our goal for now is just to get you started on building a basic intuition for the scale and nature of market activity.

# Venues

There currently 3 large parent companies which each operate several stock exchanges. ICE operates five (this is often also referred to as the "NYSE family"): NYSE, NYSE Arca, NYSE American, NYSE National, and NYSE Chicago. CBOE operates four: EDGX, EDGA, BZX, and BYX. Nasdaq operates three: Nasdaq, BX, and PSX. The only exchange outside of the three families is IEX, which is operated as an independent company.

There are roughly 30 ATS's, which is down from around 40 in 2014.

# Stocks and ETFs

Currently, there are roughly 8300 distinct securities traded in the US equities market. For each security, the market capitilization is calculated by taking the number of shares outstanding and multiplying by the current price per share. If we sum up the market capitalization in each security over all the publicly traded securities, we get the total market capitilization of the US public market, which passed 32 trillion in 2018. This represents about 38% of the global market.

Of the roughly 8300 securities, approximately 5300 of them are publicly listed companies. Most of the rest are ETFs, which stands for "Exchange-Traded Funds." More generally, there are ETPs ("Exchange-Traded Products"). An exchange-traded fund is typically designed to track the behavior of a particular collection of stocks and/or bonds, and it trades in the market throughout the trading day like an individual stock. When its price deviates from the underlying stocks/bonds it is referencing, arbitrageurs will use mechanisms that exchange shares of the fund for the underlying securities to bring the ETF's price back in line with what it is intended to track. ETNs ("Exchange-Traded Notes") similarly trade like individual stocks and relate to a collection of underlying securities, but in this case the underlying securities are unsecured debt securities.

In terms of scale, 8300 securities may not seem like a very big number, especially when we are talking about a market that functions electronically. 8300 things is certainly a number of things that computers can easily manage, right? Well, yes and no. 8300 rows in a database, or even an excel spreadsheet, is easy enough. But as a multiplier, 8300 can become a real problem. If a computation and data-intense analysis on one day's data for one stock takes 1 second to run, for example, running it on all stocks in sequence (one at a time) could take 8300 seconds, which is 138 minutes. And this is probably too optimistic - as the program is likely to save some information into memory as it goes, using up resources, and making each run of the computation potentially slower than the previous ones. Not to mention that the running time per stock is likely to vary widely based on how much data there is to churn through.

Though trading in different stocks may frequently exhibit correlation (e.g. multiple stocks impacted by common news or other market drivers), at a mechanical level trading in different stocks is separable, and can be implemented in parallel on different machines. To match a trade in Microsoft, for example, you only have to keep track of the orders and current price information for Microsoft, and you don't need to know anything about the other stock orders and prices. Stock exchanges will typically split the 8300 securities into several groups, and the trading for each group will be processed by its own machine. In the extreme case, you could imagine 8300 different machines, each acting in parallel to process the orders and produce trades for a single security. The same kind of trick can be used to speed up analytics or any other computational process that looks only at the data for one stock at a time. Parallelization is one tool that keeps things running quickly in the US equities market, mitigating the effect of the 8300 multiplier.

# Trades and Sizes

On any given day, around 6.7 billion shares are traded in the US equities market. ETFs represent about 18-20% of these volumes. Trades typically happen in round "lots," which for most stocks are multiples of 100 shares. The average trade size is currently very small - roughly 200 shares per trade. This is not to say that large trades aren't also common. Block trades (trades of at least 10,000 shares) happen frequently. For example, looking at the Nasdaq daily market summary of all trading in Nasdaq-listed securities for June 7, 2019, we see that there were over 8000 block trades, accounting for close to 20% of the day's total trading volume. The 200 share average trade size hides a much more varied distribution, comprised of many large trades, many small trades of 200 or 100 shares, and some even smaller "odd lot" trades of less than 100 shares.

Since there are roughly 6.7 billion shares traded each day with an average trade size of roughly 200 shares, we can infer that the number of trades per day is approximately 34 million. These are not spaced evenly throughout the trading day. There are often bursts of activity in one or many stocks, followed by stretches of relative inactivity. These bursts are often at speeds a human observer cannot keep pace with - frequently events are following within less than 1 millisecond of each other. Trading is also typically heavy in the closing auction: for NYSE-listed securities, for example, the closing auctions represent about 7% of the volume in those securities.

# Orders, Quotes, and Ticks

The number of orders that are transmitted to a venue is typically much higher than the number of resulting trades. This can be seen, for example, in the data that the SEC provides concerning the ratio of order cancelations to trades at each exchange.

The ratio of cancelations to trades varies widely by date and by venue, but something like 20/1 is somewhat typical. This gives us a general sense that there is on average a lot of nimble maneuvering of quotes around each trade.

Over the course of a trading day across all listed securities, there are tens of millions of times when an NBO price or NBB price changes. NBO stands for "National Best Offer," and it refers to the lowest price advertised by a would-be seller across exchanges. NBB stands for the "National Best Bid," and it refers to the highest price advertised by a would-be buyer across exchanges. The "NBBO" is a shorthand often used to refer to the pair of the NBB and NBO prices. Naturally, the NBB/NBO are time-dependent constructs, and practically speaking they are also latency-dependent constructs, as information about price changes takes time to travel from one market participant to another.

The number of NBBO changes per day is roughly on the same order of magnitude as the number of trades. Note that this does not count times when the total *size* available at the best bid or best offer price changes, but the prices themselves do not change. Nor does it count the times when displayed orders *away* from these best prices are changed or canceled.

A fairly high percentage of NBBO price changes occur in rapid succession: nearly 40% of them occur within 1 millisecond of another NBBO price change. To put this in context, we can calculate the number of milliseconds (ms) in the trading day as: 6.5 hours x 60 minutes/hr x 60 seconds/minute x 1000 ms/s = 23,400,000. If we approximate the total number of NBBO price changes across all roughly 8300 listed securities as 20 million, then we have on average about 2400 NBBO price changes per trading day per stock. So this is about 1 NBBO price change every 10,000 milliseconds. So the fact that they often ocurr within 1 ms of each

other reflects a rather extreme amount of clustering, suggesting that we should conceptualize price changes as isolated bursts of activity, punctuating long stretches of relative calm.

# Market Share

The market share of a venue is calculated as the amount of trading volume that occurs on that venue on a given day, divided by the total volume across the US equities market. Thus, the sum of the market share over all venues will be equal to 1. Roughly 62% of trading volume happens on exchanges, with the remaining 38% happening elsewhere (dark pools or internalizers). Within the exchange landscape, each of the 3 major exchange families, CBOE, Nasdaq, and NYSE, represents about 20% of trading volume (summing up all of the individual exchanges within each family). The remaining 2.5% market share belongs to IEX.

# Indices

There are several indices which are commonly used as heuristic indicators of overall market health. The behavior of indices is a frequent headline of financial reporting (e.g. "The Dow Crossed 20,000 Today!"), presumably because indices are quick and easy to report, whereas more detailed statements like "these 1345 stocks trended up by small amounts today, whereas these 1642 stocks trended down by small amounts today, these 587 stocks wiggled a lot but ultimately landed in roughly the same place, and these 763 stocks..." well, you get the idea.

So how do indices condense the more detailed activity into a single number? Each index represents a hypothetical portfolio of holdings that an investor could choose to maintain among all of the publicly listed stocks. To form a hypothetical portfolio, one must choose which stocks to include, and in what proportions. Once that is decided, the "value" of the portfolio can be calculated by multiplying the current prices of the stocks in the portfolio and summing these with weightings by the specified proportions. Naturally this gives you a number that changes dynamically as prices change. The underlying unit (whether you imagine you are holding 100 shares of Microsoft and 200 shares of Intel vs. 200 shares of Microsoft and 400 shares of Intel, etc.) doesn't really matter, as long as it is consistent. It is the relative behavior of the index over time that tells you something about the evolution of stock market over time.

Sometimes what it tells you is pretty limited or difficult to interpret. There are many degrees of freedom in defining an index, as there are many ways to decide which stocks to include, and many ways to decide the relative proportions. Some indices are clearly more meaningful than others. An "index" that is comprised of simply one stock, for example, presumably tells you fairly little about the behavior of other stocks. The price movements in a single stock will perhaps reflect general trends across the entire market and more granular trends across the relevant sector (e.g. energy, financials, utilities, etc.), but will also be driven by the specifics of that particular company. If we want the meaning of an index to transcend the details of any

one particular company or small subset of companies, we need to choose a reasonably "representative" sample of stocks and proportions to define our index.

## S&P 500

A very commonly used example of this is the S&P 500 index, which is comprised of the roughly 500 top companies, primarily chosen by market capitilization (recall: market cap = price per share times number of outstanding shares). The weighting for the S&P 500 Index corresponds to market cap, meaning that the relative weight given to a particular company is propotional to its market cap. This allows a rather clean interpretation of changes in the index: if the index increases by 10%, for example, that means the total market capitilization of all the included companies has increased by 10%.

The S&P 500 is often used as a proxy for general market health, which is quite reasonable, as the included companies cover about 80% of the total value of the US stock market. This tells us that the stock market is pretty top heavy: total market capitalization is fairly concentrated in the top 500 companies among the many thousands of total companies. Looking at the S&P 500 over time gives us a broad sense of how the market has historically trended up and down.

## DOW

The DOW Jones Industrial Average is comprised of 30 stocks of large companies. The included companies change over time, as individual companies may rise and fall in influence. When the index was created by Charles Dow and Edward Jones in the late 1800s, it included 12 companies, and it was calculated by simplying adding the price of 1 share of each of the companies and dividing by 12. It is still calculated by added the prices of 1 share of stock for each company in the index and dividing by a single number, but that number is now called the "divisor" and has historically been adjusted to make the index consistent across stock splits and mergers. In between adjustments of the divisor, we can think of the DOW as an average of the 30 stock prices with equal weight on each. This means, for example, that a 10% relative increase in the price of one of the higher priced stocks will cause a greater relative increase in the DOW than a 10% relative increase in the price of one

of the lower priced stocks. In contrast, a $1 absolute price increase in any one of the component stock prices will have an indentical effect on the index.

The companies included in the DOW index represent about a quarter of the value in the total US stock market. But because of the smaller sample size and the fact that the index is not weighted by market cap, the DOW is not really a great proxy for overall market health. It remains a popular data point mostly because of its historical longevity and broad familiarity.

## Russell Indices

The Russell Indices, e.g. the Russell 3000 Index, the Russell 2000 Index, and the Russell 1000 Index, etc., pick out different subsets of stocks based on market capitilization. For example, the Russell 3000 Index includes the largest 3000 companies by market capitalization, while the Russell 1000 includes only the top 1000 companies and the Russell 2000 includes the next 2000. Once the subset of stocks is determined, their individual contributions to the index are weighted proportionally to market cap. The Russell indices can be used as proxies for the cummulative performance of different categories of stocks: large cap, small cap, etc.

## Index Funds

Since an index corresponds to a hypothetical portfolio of stocks, it is possible to maintain a portfolio matching the hypothetical one, and turn the hypothetical behavior into realized returns. This process isn't quite perfect, since sometimes the stocks or weights in an index change, and adjusting a real portfolio to keep in sync with the hypothetical one involves buying and selling that incurs additional cost and risk. But to a reasonable extent, one can approximate the performance of an index with a real portfolio, and this is what index funds are designed to do. Investing in an index fund that is properly managed should allow you to achieve *approximately* the returns corresponding to the hypothetical index it is implementing. This is why index funds that simulate the S&P 500 index, for example, are popular investments for investors who just want to generally benefit from stock market growth in an overall sense, without entangling their returns too closely with the performance of a smaller number of individual stocks.

When the constitution of an index changes, a lot of market activity will result as index funds seek to adjust their holdings accordingly. The Russell indices, for example, are reconsituted on a scheduled yearly basis. This is to adjust for changes to market capitilization that have occurred over the preceding year. Trading on and around such scheduled events will be markedly heavier (and different in distribution) from ordinary trading activity.

# Volatility and Variation

Over time, there is a lot of variation in the kind of daily stats we have quoted above. Some variation in trading behavior is driven by structural things: e.g. index rebalancing or the fact that stock options have certain standardized expiration dates (following the third Friday of a given month). But even aside from these anticipated structural events, within a day or hour or minute, etc. there can be wild changes in activity and prices within a single stock, multiple stocks, or broadly across the market. The number of NBBO changes in a day for instance, can vary by a multiplicative factor of 3 or 4 (or more) when we compare a relatively calm trading day to a relatively turbulent one.

In a technical sense, "volatility" in the context of stock trading usually refers to the standard deviation or variance of prices for a particular security or set of securities. Roughly, this is a measure of how far prices tend to drift above and below their average over a specified unit of time. This gives additional color around price trends that mere averages or medians do not. For example, if we say "the average price of Microsoft over the trading day today was X," we mean that we took all the trades of Microsoft over the day and averaged their prices (presumably weighted by their sizes). But this alone tells us nothing about whether we had, say, many trades at price exactly X, or trades that swung widely between lows of X/2 and highs of 2X, averaging out to X overall. Volatility disambiguates these disparate situations. Low volatility means that we had trades mostly clustered tightly around the average price, while high volatility means that prices oscillated considerably.

Information about expected future volatility can be inferred by looking at option prices. Since the value of an option depends on the probability that a stock price will go above or go below a specified strike price in the future, the price of an option right now can tell us something about "the market's" estimate of the likelihood of particular future price movements. The CBOE Volatility Index, or VIX, measures this kind of "implied volatility" by using SPX option trades to infer the market's expection of 30-day volatility in the S&P 500 index.

*Further reading*:

SIFMA Insights: US equity Market Structure Primer, July 2018 (*this was a source for many of the stats above)

# External Resources

100 Years of Market Structure Improvement (KCG 2016, Page 27)

The New Stock Market: Sense and Nonsense (Fox, Glosten, Rauterberg 2015)

The Hidden Alpha in Equity Trading (Oliver Wyman 2014)

Will the Market Fix the Market? A Theory of Stock Exchange Competition and Innovation (Budish, Lee, Shim 2019)

Equity Market Structure Primer (Sifma 2018)

Yardeni Research US Flow of Funds: Equities (2019)

Incentivizing Trading Behavior Through Market Design (Aisen 2017)

# History of the Stock Market

Key points/trends in time to cover:

- Exchange Act
- creation of the SIPs
- creation of Nasdaq
- creation of BATS, Direct Edge, etc.
- floor trading replaced by electronic trading
- FINRA - history of its role?
- Reg NMS
- Finra rule 5310 on best ex
- consolidation into 3 exchange families
- rise of dark pools as taking on significant volume
- Dodd-Frank (fee filings becoming immediately effective)

# Current Issues

Conflicts of interests between participants and controversial practices:

- Brokers: conflicts of interest vs. their clients, bundling of services, SEC fines
- Exchanges: conflicts of interest vs. brokers and vs. investors, monopolistic practices
- Investors: fraud, ponzi schemes, concerns about (manipulative?) short-selling
- Issuers: stock buy-backs, issues around IPO pricing
- Market makers and proprietary trading firms: predatory trading practices, wasteful technology investment
- Regulators and academics: revolving door between regulators and the regulated, academic studies funded by stakeholders, data quality/availability issues
- Everyone: technology failures/outages

# Proposed Fixes

some implemented or partially implemented:

- speed bumps
- Reg SCI
- CAT
- increased disclosure requirements

some theoretical:

- frequent batch auctions
- SIP upgrades
- even more increased disclosure requirements
- MFID II

# Glossary

## Alpha

in the context of trading, when an order has alpha, it means it has inherent profitability that has not yet been realized. For example, if the market is about to move higher, a buy order placed at the current market price has alpha. Conversely, an order with negative alpha would

## Bid

## Offer

## BBO

## NBBO

## Passive (order)

when referring to an order, the term *passive* means priced conservatively relative to the prevailing market. For example, if the prevailing best offer for a stock is $10.00 (i.e. the lowest price at which someone is advertising publicly their desire to sell the stock is $10), a passive buy order would be one priced *lower* than $10, such that it would not immediately transact in the market.

## Aggressive (order)

when referring to an order, the term *aggressive* means the order is priced at a level such that it can immediately interact with a counterparty. For example, if the prevailing best offer for a stock is $10.00 (i.e. the lowest price at which someone is advertising publicly their desire to sell the stock is $10), an aggressive buy order would be one priced *at or higher* than $10, such that it can immediately transact in the market.

## Displayed (order)

a *displayed* order is one whose price and size are publicly disseminated on a given exchange or ECN's market data feed.

## Non-displayed (order)

a *non-displayed* or *hidden* order is one that is live and available to transact on a venue, but that is not advertised via the venue's market data feed.

## Direct bilateral relationship

a "bilateral relationship" is simply a relationship between two parties, in this case a retail broker and a retail wholesaler.

# History of the stock market/market structure

[Develop in the context of competing motivations/goals among the different participants, and evolving balance between them. This will lead nicely into the next section on controversies.]

Key points/trends in time to cover:

- Exchange Act
- creation of the SIPs
- creation of Nasdaq
- creation of BATS, direct edge, etc.
- floor trading replaced by electronic trading
- FINRA - history of its role?
- Reg NMS
- Finra rule 5310 on best ex
- consolidation into 3 exchange families
- rise of dark pools as taking on significant volume
- Dodd-Frank (fee filings becoming immediately effective)

# Hot Topics

## Overview

# Brokers

## Conflicts of interest

### Rebate harvesting / non-passthrough pricing

### Operating a dark pool

### Trading against customers / central risk book?

### IPO underpricing

### TCA

(i.e. grading their own report card)

### (Retail) Payment for order flow

## SEC fines

(maybe have a full page describing SEC fines of market participants - and link to those sections here?)

### Misleading customers

#### In marketing materials

(e.g. Barclays)

#### BAML passing through the wrong LastMarket tag

### Insufficient risk checks

https://www.wsj.com/amp/articles/wall-streets-dark-pools-face-new-transparency-rules-1531924509

https://www.sec.gov/news/press-release/2018-256

https://www.wsj.com/articles/credit-suisse-barclays-to-pay-about-150-million-to-settle-dark-pool-investigations-1454256877

ITG secret trading desk, misusing dark pool info $20.3mm - https://www.sec.gov/news/pressrelease/2015-164.html

Pipeline $1mm Oct 2011 https://www.sec.gov/news/press/2011/2011-220.htm

Merrill Lynch (BAML) $42mm for misleading customers about trading venues June 2018 https://www.sec.gov/news/press-release/2018-108

# Exchanges

## Monopolistic practices/price gouging

### Prop market data feeds

Allison's blog post part 1 - https://medium.com/prooftrading/market-data-pricing-part-1-of-many-f4213ac214f6

Adrian's paper - https://iextrading.com/docs/The Cost of Exchange Services.pdf

### Connectivity

### Co-location

### Pricing tiers

https://www.nytimes.com/2016/03/02/business/dealbook/stock-exchange-prices-grow-so-convoluted-even-traders-are-confused-study-finds.html

### Listings? Spotify IPO

## Conflicts of interest

### Operating the SIPs

### For-profit SRO

## Complex order types

## SEC fines

NYSE fined $14mm for July 2015 outage - https://nypost.com/2018/03/06/new-york-stock-exchange-fined-14m-for-2015-outage/

# Investors

## Fraud, ponzi schemes, etc.

## Activist short selling

(manipulative?)

# Issuers

## Stock buybacks?

# Market makers, prop firms

## Predatory trading

## Wasteful technology investment

(e.g. "racing")

## KCG collapse

(maybe a whole page of major technology outages/failures?)

## BATS failed IPO

## NYSE July 2015 3.5 hour outage

(July 2015)

## Nasdaq FB IPO

Nasdaq UTP SIP 15 min outage (January 2013)

Nasdaq SIP outage/total shut down for 3 hours (August 2013)

NYSE CTA SIP 27 min outage (October 2014)

## Other exchange outages

# Regulators/academics

## Revolving door between regulators and banks

## Academic studies funded by industry firms

(Citadel, Nasdaq, etc.)

# Proposed fixes

[maybe discuss proposed overhauls, as well as more middling suggestions, etc. and what the challenges to implementation are. E.g. MFID II, frequent batch auctions, distributed SIPs or augmented SIPs, increased disclosure requirements. For each, discuss what problem it is supposed to address, how it would or would not address it, how stakeholders view it - who is for and against and why, speculate main impediments to adoption or unintended consequences]