**BLOCKCHAIN AND THE LAW – DRAFT 1/29/2021**

## *I. Introduction*

Good morning. I'd like to take this opportunity to talk about blockchain and the law. Like many, I have grown increasingly fascinated by some of the ways that distributed ledger—or blockchain—technology might impact both our financial markets and our laws.

Now I bet that almost all of you have some familiarity with this topic because during the last few years, discussions about bitcoin, cryptocurrency, distributed ledgers and blockchain technology have captured the imagination of millions.

- Newspaper headlines breathlessly announce the latest developments, investment bankers abandon their jobs for fintech startups, and celebrities tout the latest new coin offering.
- Recently, I glanced over at my son's video game website, only to see a large advertisement banner broadcasting "click here to find the very best cryptocurrency investments."

- Here at UVA, I've even been told that our foundation has been approached by a donor seeking to give the law school some cryptocurrency, and they are now researching how to accept such a gift.

Of course, there wasn't always this much enthusiasm. About 9 or 10 years ago I first heard about bitcoin from one of my students here who mentioned this to me as a neat new currency without any government backing. [best intel on trending events always comes from students]. I promptly dismissed this as a ridiculous idea that clearly wouldn't go anywhere. This is why I remain a law professor…

Despite all the attention on blockchain, however, I would bet that many people don't really understand what it does or how it might impact our laws. Sometimes when I talk to very sophisticated people, they admit that they often nod wisely when the subject comes up but cannot really explain what it does. I don't think it's necessary for most people to understand the inner workings—think about the internet and how little most if us know about TCP/IP layers and all that. But I do think it is important helpful to have a high-level sense of what a blockchain really does to understand its potential. Because I do think there is a possibility for vast legal changes, especially with payment systems, corporate laws, and other areas that need to keep track of property ownership. I think the most interesting ideas relate to foundational infrastructure changes—not the hottest new crypto-coin. And while the exact path forward is uncertain, I'd argue that it is worth paying attention to this space.

It is difficult to discuss this in the abstract, so I thought I might address this topic by telling two quick stories: one about the law and history of stock clearing markets, and one about what blockchain technology actually does. After that, I will briefly try to connect these two narratives to suggest some possibilities for the future of corporate law. And then I'll expand to address a few other areas of law and invite a discussion on the topic.

II. ***The Path of Law: Financial Markets Example***

Let me begin with a story about the history of stock clearing markets. What happens after you buy a share of stock? How does the resulting settlement and transfer of your new ownership rights take place? For most people, these are especially uninteresting questions. Front-end trading strategies and hedge fund algorithms may be exciting, but the recessed plumbing of back-end stock clearing is not. Many just ignore the topic and assume that a share of stock will eventually "get" to its new owner. But I think there's a pretty neat story here.

If we go back to the first half of the 1900's, stock transfers looked a lot like buying a selling a used car.

When a corporation issued stock, it would typically create a specifically identified certificate designating ownership of some number of shares. These certificates were stashed in file cabinets, safe deposit boxes, or perhaps broker storerooms—similar to what you probably do now with the pink slip for your automobile.

When you wanted to sell some stock, you might call your broker to execute the financial trade. Then, to settle and clear everything, you (or your broker) would track down the certificate, get it signed by a notary, and send it over to the buyer (or their broker). The buyer would then send over the documentation to the firm's transfer agent (typically a large bank) who would formally record the change on the corporation's ledger (so future dividends and voting information would go to the right person) and issue a new share certificate. [With large trades, this might need to be done many times because it was uncommon to issues certificates for more than 100 shares.]

We might call this settlement version 1.0, and it usually worked out OK during the first half of the 1900's. But shares of stock can trade hands much more frequently than a used car. [slide] By the 1960's, trading volume rose, and the system was snowed under. During the height of this paperwork crisis, brokers found themselves running days or even weeks behind. Clerks made mistakes, and the number of "fails" increased dramatically. (For example, Lehman Brother ran an infamous analysis that concluded as follows: "Well, we have $475 M in securities whose owners we can't locate, and we owe clients $220 M in securities that we can't find!"). Later Congressional hearings estimated that the mafia took advantage of the chaos to steal more than $400 million in securities. Finally, at the height of the crisis, traders closed the stock markets every Wednesday—just so the unruly piles of certificates could be inspected for authenticity, organized for distribution, and routed to their new owners.

## WALL STREET IN THE 1960s

This clearly would not do, and Congress pressured the SEC to clear this logjam. Eventually Wall Street settled on a fix: share immobilization.

How does this version 2.0 of stock clearing work? A central entity emerged to replace the network of messengers scurrying across the back alleys of New York. Over time, this organization, now known as the Depository Trust and Clearing Corporation ("DTCC"), began to serve as the permanent record owner for most shares. Certificates (physical or electronic) are now immobilized by this central entity and a corporation and held in fungible bulk. This means that the same record holder persists as the formal owner of the stock, and the clearinghouse transfers beneficial ownership electronically from seller to buyer via bookkeeping adjustments.

Let me try to illustrate with an example. Suppose you run a modestly endowed investment portfolio, Rivanna Investments. Rivanna owns Snapchat but decides to sell it all off. The student will notify its broker—we'll call them Fidelity—to execute the trade. The economic effect of the sale will usually be benchmarked immediately, but it will take longer to work everything through settlement. If another client of Fidelity happens to take the other side of the trade, then DTCC, in theory, need not even be aware of the exchange. Fidelity can just adjust its internal records to account for the transfer of cash and the beneficial stock ownership between buyer to seller. In many cases, however, the buyer will be represented by different bank or brokers. If so, DTCC will need to record a bookkeeping adjustment to reduce the shares allocated to Rivanna's broker and increase the shares allocated to other bank custodians. Note, however, that DTCC persists as record holder of the shares throughout this process and that no specific shares are identified as "the ones" being traded. Snapchat might not even realize that a large trade has occurred.

The rise of DTCC and share immobilization has, on the whole, been a welcome development. Indeed, it is impossible to imagine how the old system could support the roughly 1 billion stock trades that now occur each day on the NYSE. But the use of intermediate agents greatly complicates business law and the legal mechanisms that were crafted during settlement version 1.0 to convey the vote and other important legal rights to beneficial shareholders. And there are problems—which should not be too surprising since many corporate laws were written under settlement version 1.0.

Let me mention just one example: Section 11 tracing. Section 11 of the Securities Act of 1933 imposes liability when shares are sold pursuant to a registration statement that contains materially misleading statements or omissions. It is meant to ensure the accuracy and integrity of registration filings and generally considered to impose hefty sanctions on a firm's officer, issuers, and underwriters for misleading registration statements. Section 11 lawsuits can grow complicated, however, when it becomes necessary to determine which shareholders may assert a claim or join a class action lawsuit. The statute is most clearly written to protect initial purchasers—who might be expected to rely on the registration statement (either explicitly or constructively) when buying stock directly from the issuing firm. These buyers have clear standing to bring a Section 11 claim when something goes awry.

But what about secondary market purchasers who transact after the initial issuance? Jurisdictions are split on whether these buyers can bring or join claims, and the U.S. Supreme Court has not weighed in on the topic. The most common approach is to allow secondary market claims if plaintiffs can "trace" their stock back to the specific shares that were issued in connection with the tainted registration statement. The logic seems to be that if the shares can be linked to the offending registration filing, then a Section 11 claim should not be severed, under the plain language of the statute, simply because there has been a transfer of ownership.

But what happens if a court cannot be 100 percent certain that a given plaintiff's shares were sold via the tainted registration statement? Consider the facts of Krim v. pcOrder.com. Several plaintiffs filed a Section 11 claim alleging that pcOrder had conducted an initial public offering with a fraudulent registration statement. For a few months after this offering, no other shares were available for trading on the secondary markets. Eventually, however, some shares that had not been part of the public offering (owned by insiders at the firm) began to trickle out into the market. Most of pcOrder's stock was isolated in fungible bulk by the DTCC, with Cede listed as the formal record owner.

Consider the fate of three different plaintiffs in the case. Plaintiff A bought 1000 shares on the secondary market soon after the IPO—when the only shares available for purchase were those released in the IPO. Accordingly, plaintiff A could successfully trace his shares, as a matter of simple logic, and gain standing for a Section 11 claim. By contrast, plaintiff C bought shares at a later date—when some of the insider shares had leaked into the market. The Fifth Circuit court ruled that this plaintiff was unable to definitively trace his shares back to the offering because almost 10 percent of the pool of available shares (at the moment of purchase) were not ones issued in connection with the registration statement. As we have seen, specific shares are not typically identified with a transfer, so plaintiff C was unable to pierce the fog of DTCC's fungible bulk holdings. Even plaintiff B—who bought stock after plaintiff A but before plaintiff C, from a pool where 99.85% of the total shares available could be attributed to the IPO —was blocked from joining the case. The probability that plaintiff B bought at least 1 of his 3000 shares from

those issued via the public offering was incredibly close to 100%. But most courts, including this one, insist on absolute certainty. This may or may not be problematic, depending upon one's views about the optimal level of deterrence under the Securities Act. But it does seem strange to establish a general principle that Section 11 claims are available for secondary buyers and then adopt a follow-on rule that effectively guts those same claims in most contexts.

The bottom line, then, of this first story is that we have cobbled together a functioning settlement system. But corporate law has paid a price from the resulting complexity.

## III. ***How does blockchain technology work?***

Now, we come to story #2 on the rise of blockchain technology.

What is a blockchain and what can it do? A distributed ledger is simply a sequential database of assets that is shared across a network of users. It is distributed in the sense that all participants in the network have their own copy of the ledger identifying historical transactions and the resulting ownership rights associated with the entire group of assets. By comparison, most economic entities currently use double entry bookkeeping ledgers to track the disposition of their assets; these ledgers are both private and fragmented. They only contain information related to the specific assets that each firm owns. Distributed ledgers, on the other hand, are more akin to government managed real property registration systems, where anyone might examine current and historical ownership claims on a given parcel of property. But while public property records are typically centralized and housed in a single location, with limited accessibility, a distributed ledger can be split into hundreds or thousands of identical copies and situated in the scattered computer systems of individual members or users.

The accuracy and security of a distributed ledger is maintained through blockchain technology,

which ensures that all copies match and that all modifications (reflecting new transactions) follow the same path. As the relevant assets are bought and sold, the details of each transaction are gathered and organized into a string of data according to an established formatting algorithm. Multiple transactions are then grouped and transferred into a data block. This block is linked (or "chained") to the earlier blocks of transactional data in the ledger, time stamped, and processed in a way that both refers to and verifies prior transactions. In other words, the cryptography is designed so that it becomes progressively more difficult for older blocks to be re-written, increasing the verifiability and security of prior transactions. The results are broadcast to members in a synchronization protocol, and the system starts a new cycle.

A slightly more extensive description might be helpful for understanding the transformative nature of distributed ledgers. First, consider a public distributed ledger. A member joining the distributed ledger network will receive unique public and private encryption keys. The private key is used to certify a transaction and can be verified by others in the network. If a member wants to settle a fresh transaction, she uses her private key to "sign" and transmit the relevant details (such as which assets are being sold and to whom they are being transferred) through a string of data. The public key is analogous to an email identifier, and it allows others—such as a buying party—to locate the selling member in order to send funds related to the exchange. Any such transfer to a member via their public key can only be unlocked with that member's private key. As more of these transactions occur, they are queued, processed, and eventually published throughout the system according to the system architecture (more on this in a moment). This makes the ledger verifiable to any participant with read permission, in a way that should render the history of each transaction secure and irreversible.

The right to publish changes to a distributed ledger may vary according to the system. With a public and open network, such as bitcoin, anyone who accomplishes a specific task might earn the right to add the next block of data—subject to confirmation of their success. The task will be arduous—such as solving the next iteration of a complex mathematical puzzle—but also easy to verify as correct once a solution is presented. With privately permissioned networks, the right to modify might not require solving a task; preauthorized, trusted agents (such as banks) can simply add the new block.

With bitcoin, for example, a member must find (or "mine") a numerical solution that takes data from the most recent solution and the valid transactions to be recorded as inputs. The goal is to combine this data with another number (called a nonce) in a way that provides an output, when run through the relevant hash function, that accomplishes some specific task—such as producing an output string that begins with 12 zeros. The first person to accomplish this "wins" the right to write the data to the distributed ledge, gains a "coin" or token of value, and starts the cycle anew. All others must begin from scratch with the new output string, and any leftover work from the prior iteration will probably be useless. The task is thus (1) difficult to accomplish, as the solution nonce can only arise through trial and error and massive computer processing power; and (2) easy to verify by looking at the solution output (are there really 12 zeros at the start of the output?). Moreover, the difficulty of the task might be modified on the

fly to render a solution slower or faster. If the goal of the system is to generate a new block approximately every 10 minutes, then the solution might require only 10 zeros at the start of the sting (if the goal is to expedite a solution) or require 14 zeros at the start of the string (if the goal is to delay the next solution).

Now, why would independent participants bother to record the details of exchange transactions between other parties? One answer is that they can be provided with incentives for doing so. For example, the solution to a mathematical puzzle might incorporate the details from completed transactions at the top of the recording queue. Winning the race to solve each iteration thus requires participants to process and memorialize the most recent transactions. Why participate at all? For one reason, the parties seeking to record their asset transfer details might pay a commission to the recording member who successfully publishes the transaction in the next block of data. Second, members might earn a "point" or "token" when they become the first person to solve the next iteration of the problem, offer proof of this work, and publish another block in the chain. These tokens might then be monetized (as with bitcoin and other cryptocurrency systems) and become valuable in their own right.

In a closed distribute ledger system (types 2 and 3 above), it may not be necessary to create cryptocurrency rewards via the generation of tokens. The payment of commissions should be sufficient to incentivize members to develop the next bock in the chain. Indeed, if the members with modification rights enjoy sufficient levels of trust, mining for the solution to a complex mathematical puzzle may not be needed at all. The modifying parties will simply add a new block whenever enough transactions are queued to justify the next link in the chain.

Embedded throughout the use of distributed ledgers is a processing "hash function," the heart of the cryptography. A hash function is simply an algorithm that takes a variable length string of data as an input, crunches the information, and spits out a fixed-length string of numbers and letters. The hash function used by bitcoin, for example, was initially devised by the National Security Agency ("NSA") and produces a 256-bit output. An identical output result is always obtained when the same string of input data is fed into the function. But it is exceedingly difficult to reverse engineer the input data by observing only the output data; changing just one character in the input string will completely transform the output of a hash function. In order to create the next block in the chain, a member will need to take the output hash string from the most recent block and compute a new output hash string that incorporates the details of the queued transactions for recording, adds a time-stamp for additional security, and meets any other predefined requirements.

## WHAT IS DISTRIBUTED LEDGER TECHNOLOGY?

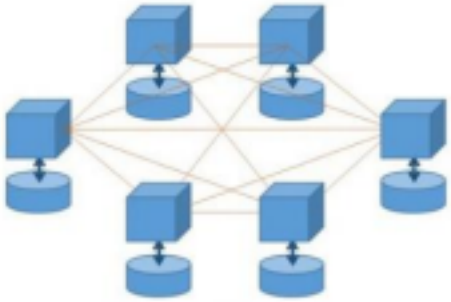

Distributed blockchain

In this way, the most recently created block intersects with older blocks in the chain. Because the only way to develop a new block is to begin with the previous block's output string—and because the only way to obtain the previous block's output string is to crunch the most recent transactional records with the hash function—there will be an ongoing confirmation of prior transactions that grows stronger as the length of the chain increases. Any effort to rewrite history, by going back several links to change the recorded owner of a given asset, will generate a different output string when run through the hash function. This alteration would not match the results obtained by other members and should be rejected. The only way to succeed in stealing assets through a modified distributed ledger would be to find an alternative solution to the hash algorithm and quickly permeate this alternative reality down through subsequent links in the chain in a way that complies with the synchronization protocol. This is thought to be exceedingly difficult (though not impossible ) because the correct chain continues to move forward as other members add links—making any attempt at theft a moving target. Finally, each member will need to keep an accurate version of the current ledger, or version control problems will proliferate. For this reason, distributed ledgers must incorporate a synchronization protocol. The details of alternative protocols can grow complicated, but, in a nutshell, a member seeking to add a new block must generally broadcast its request, and wait for a majority (or supermajority) of participants to agree that the solution works (if proof of work is required) and is timely (first in line to add a new block). Again, differences may arise between public and permissioned ledgers, but once the pending block is verified, the new chain of data should be irrevocably connected. The data is disseminated throughout the network, and the next cycle begins.

## IV. ***Connecting Blockchain Technology with New Legal Systems***

Now, let's try to put the two of these stories together to suggest some possibilities for the future of corporate law.

if we could snap our fingers and create an ideal stock clearing platform, how would it work? What would we want? Well, we would probably abandon the complicated and multi-layered distinction between record and beneficial owners. Instead, we might have golden ledger, real time clearing and settlement. When there's a trade, it's quickly processed on the blockchain and the transfer happens. So one possibility would just be to hook investors up, give us each an account, and go to town.

One of the more interesting questions involves the role, if any, that banks, brokers, and other intermediaries might play in a new settlement platform. It is possible to imagine a world of complete disintermediation, where individual investors join exchanges directly, downloading software to participate as direct members of a distributed ledger.

But I doubt that this is how it will all happen because many of us want to work with agents or advisors, and we should all be paranoid about a possible security breach. Some investors will continue to enjoy the financial advice that they receive from experts, and they might be loath to trade directly. Others may continue to seek the diversification and scale benefits that arise through mutual funds and alternative investments, such as hedge funds. Moreover, the security risks and capacity limitations associated with a public ledger stock exchange will probably be too great for most people to stomach in the near term.

For these and other reasons, it is likely that banks and brokers and other entities will continue to play a role in the creation and operation of new stock clearing platforms. One can imagine, for instance, a consortium of intermediaries establishing a private ledger where they are the only ones who retain viewing and modification rights. An investor seeking to buy or sell stock would contact a broker member to execute trades. The broker would locate a counterparty and then process and record each transaction on the distributed ledger. With just a small number of permissioned parties who trust each other, each broker would not need to solve an algorithmic puzzle; they could just write a new block as soon as a sufficient number of transactions are queued for processing.

From the investor's perspective, not much would seem to change. The brokers would continue to provide economic information and share positions. The trades would settle quicker, and there would be a detailed and traceable record of title for every single share of stock. Depending on the level of visibility offered by the ledger, we might be able to see ownership and trading data about other shareholders in a firm. It is also possible, however, that the platform might keep this information from individual investors.

Privacy – real concerns. Can also imagine anonymous accounts as a way to preserve ability to opt out of identification.

Governance - From a design perspective, we might think about 4 different systems differing along the key dimensions of who can see the information and who can change the information: (1) traditional ledgers where a single copy is privately retained by each user (which is not a

distributed ledger); (2) private distributed ledgers where there are multiple copies of the ledger that may only be viewed and changed by authorized participants; (3) public distributed ledgers that are viewable by many but modifiable only by a subset of trusted actors; and (4) publicly shared distributed ledgers that may be viewed and modified by any user under consensus protocols (some cryptocurrencies).

The challenge – huge coordination event. Like moving from driving on right side of road to left. Various Transition scenarios

o Pilots (underway)
o Redundant systems (Shadow processing)
o Roll out. IPO; With new trades – pulled share by share
o Firm recall from DTC
o Big Bang

Companies – desire to experiment driving change

Investor communities – recognize there might be a better way

Support from regulators: State law – seems supportive. To be honest, many of the corporate laws still on our books were designed with clearing v. 1.0, so the challenges have arisen with centralized fungible bulk. In some cases it could be easier to work through legal resolution with traceable shares.

At the federal level, there is probably going to be a need for SEC support and several regulations might need to be relaxed. More specifically, the SEC adopted or approved rules about 15 years ago that could make broader experimentation difficult.

One rule effectively requires issuers to make their shares eligible for centralized clearing as a condition of listing on an exchange. Another DTC rule, approved by the SEC, allows DTC to refuse to honor requests by issuers to withdraw securities, which makes it tough for a large firm to flip a switch and migrate to a new platform.

At the time these rules made some sense because the SEC didn't want to support a return to paper based exchanges. But with new technology, it fair to ask whether this is still the best approach. But the issues are complicated, nuanced, and important. Some have called for a formal comment process at the SEC, and we should all stay tuned.

## V. ***What other Legal Systems May be Impacted by Blockchain? [as time permits]***

Payment systems

Smart contracts

Tax regulation, reporting, and enforcement

Securities Regulation

Intellectual Property

Liability and Responsibility

In short, an enormous number of issues may be impacted—which is why I think this will be a terrific area to watch for lawyers over the next few years. Let me stop here and invite your comments and questions.